

THE YALE LAW JOURNAL

YAIR LISTOKIN

Learning Through Policy Variation

ABSTRACT. Rationalist analysis of policymaking, exemplified by cost-benefit analysis, ignores the variance in outcomes associated with policies and seeks to maximize expected outcomes. Burkeans, by contrast, view policy outcome uncertainty negatively. The Burkean approach is echoed in the precautionary principle, which argues that policies with hard-to-determine or high-variance outcomes should be avoided. Both approaches are the subject of vast literatures. This Article argues that both approaches are wrong. When policies can be reversed in future periods, variation in the outcomes associated with a policy is a good thing. Reversibility means that the downside risk of high-variance policies is limited; policies with unexpectedly bad outcomes can be changed in the next period. The upside of high-variance policies, by contrast, may last indefinitely, since policies with unexpectedly good outcomes will be retained. Thus, when policies are reversible, policymakers should deliberately choose policies with uncertain outcomes, other things equal. The Article also examines the assumption of policy reversibility. It shows that the most important source of irreversibility for policy analysis is irretrievable “sunk costs” rather than the potential for catastrophic outcomes or policy inertia. As a result, policies are more reversible than commonly appreciated. The Article then examines optimal policymaking under irreversibility. Under extreme irreversibility, conservatism of a particular sort, called the “real options” approach, constitutes the best policy. More generally, the Article argues that the appropriate attitude toward policy variance depends upon the reversibility of policy. This analysis illuminates many puzzles in constitutional law and institutional design, such as the puzzling difference between entrenched statutes, which are unconstitutional, and sunset clauses, which are permitted. The Article concludes with recommendations to encourage policymakers to use variance more effectively.

AUTHOR. Associate Professor of Law, Yale Law School. I thank Bruce Ackerman, Ian Ayres, Richard Brooks, Heather Gerken, Oona Hathaway, Alvin Klevorick, Daniel Markovits, Theodore Ruger, and participants in Yale Law School’s faculty seminar for many helpful suggestions and comments. All errors are my own.



ARTICLE CONTENTS

INTRODUCTION	483
I. THEORIES OF PUBLIC POLICYMAKING	486
II. CHOOSING POLICIES FROM AN OPTIMAL SEARCH PERSPECTIVE: THEORY AND APPLICATIONS	491
A. The Theory of Optimal Search	492
B. Factors Influencing the Optimal Degree of Policy Variance	495
C. The Optimal Search Approach: Applications	497
1. Choosing Between Reversible Regulations	497
2. The Optimal Search Approach and Penalty Default Rules in Contract Law	501
3. Increasing Shareholder Democracy in Corporate Law	505
4. Federalism and the Optimal Search Approach	512
III. THE OPTIMAL SEARCH APPROACH: OBJECTIONS, RESPONSES, AND MODIFICATIONS	518
A. Burkean Objections	519
B. The Costs of Changing Policies	522
C. Irreversibility, Real Options, and the Optimal Search Approach	524
1. The Real Options Approach to Policymaking	524
2. Burkeanism, the Precautionary Principle, and the Real Options Approach	526
3. Sticky but Reversible Policies	527
D. Optimal Variance in Different Policymaking Contexts	529
E. Reversibility and Institutional Design	533
1. Sunset Clauses and Legislative Entrenchment	534
2. Separation of Powers	536
3. Stare Decisis	538
IV. PUBLIC POLICYMAKING INCENTIVES AND THE OPTIMAL SEARCH APPROACH	539
A. High-Variance Policies and Reelection/Reappointment Incentives	540
B. High-Variance Policies and Incentives for Political Advancement	542

C. Federalism and Incentives To Innovate	544
V. IRREVERSIBILITY AND POLITICAL INCENTIVES: APPLICATIONS AND RECOMMENDATIONS	546
A. Reversible Regulations	546
B. Contract Default Rules and Increased Judicial Policymaking	547
C. Shareholder Power	549
D. Federalism and Preemption	551
E. Other Sources of Variation: Direct Experimentation	552
CONCLUSION	553

LEARNING THROUGH POLICY VARIATION

INTRODUCTION

How should policymakers choose laws and regulations when outcomes are uncertain? The answer initially seems simple: they should choose the best policies—the ones with the highest average payoffs along some metric. Burkeans have a different view. They are skeptical of human ability to divine the best policies. Instead of encouraging policymakers to choose the policies that seem best, Burkeans urge policymakers to choose policies that change the status quo incrementally rather than drastically. When policymakers can learn from the results of their laws and make changes, however, both the common sense position—choose the best policy—and the Burkean position—choose new policies cautiously and incrementally—are often wrong.

When learning is possible,¹ innovative high-risk policies with lower average outcomes but the potential for greater outcomes become preferable.² If a high-risk policy proves a failure, then the policy can be changed, and the policy with the highest average payoff can be pursued. If the policy succeeds, then policymakers will have achieved an ideal outcome and will no longer need to search for alternatives. Learning allows policymakers to limit the downside of a high-risk policy but still enjoy the upside, making a high-risk policy with a lower average payoff a better initial choice in many cases than a low-risk policy with a higher average payoff.

In other words, policies serve two functions. Their primary function is to achieve some outcome in the current period. But the information provided by observing a policy's outcome also assists the search for better policies for the

-
1. "Learning" about policies refers to the process through which the variability in outcomes associated with a policy is reduced. Before implementing a policy, policymakers may have only a dim idea about the effects of the policy. After implementing the policy and observing its effects, policymakers will often have a much better sense of the outcomes associated with the policy in current and future periods. Thus, policymakers will have "learned" about the policy and have a much greater ability to predict the policy's impacts. It may also be possible to "learn" about a policy without trying it. For example, one can learn about untried Policy A by extrapolating the effects of a similar but tested Policy B to Policy A. See *infra* Part III. This form of learning necessarily will be inferior to directly learning about Policy A by trying it.
 2. See Martin L. Weitzman, *Optimal Search for the Best Alternative*, 47 *ECONOMETRICA* 641 (1979) (developing the mathematical basis of the optimal search approach); see also Philippe Aghion et al., *Optimal Learning by Experimentation*, 58 *REV. ECON. STUD.* 621, 642-43 (1991) (same); Volker Wieland, *Learning by Doing and the Value of Optimal Experimentation*, 24 *J. ECON. DYNAMICS & CONTROL* 501, 513-19 (2000) (characterizing the process of experimentation for economic agents, such as firms with unknown demand curves). This Article examines public policymaking as an optimal search process—a novel perspective in both the economics and legal literatures.

future. And the best policy from a search perspective often differs from the best policy for the current period. The “optimal search” for a policy seeks an excellent policy that will enable policymakers to end the policy search. Thus, optimal search theory favors high-variance policies, because variance increases the probability of finding excellent policies. The average outcome of a policy matters less from an optimal search perspective than the upside of a policy because it is unlikely that a reasonable but suboptimal outcome will end the search for a good policy. A bad policy, moreover, can be changed in the next period.

The economics literature on optimal search focuses on the choice between two new possibilities. Policymaking, however, often involves the choice between a new policy and the status quo. Generally, new policies have higher variance in outcomes than existing policies. The optimal search approach, therefore, suggests that new policies should be implemented frequently. Even when choosing between alternative untried policies, the policy with greater variance is the better choice, other things equal.

The implications for public policymakers are wide ranging. Take contract law, for example. A number of scholars have proposed complex rules that aim to improve the status quo.³ Critics argue, however, that the efficiency of the rules depends on questionable assumptions about individual behavior.⁴ Given these defects, most of these proposals have never been tried. The optimal search approach, however, suggests that these policies should be tested. If the alternative default rules fail to improve social welfare,⁵ then the policies can be discarded. If the default rules succeed, however, then policymakers will have achieved a significant, long-lasting improvement in the quality of law.

Corporate law involves similar disputes. It typically changes little, and every proposed change is met by critics who praise the status quo. The critics may be right that the status quo is better on average than are proposed reforms, but the optimal search approach suggests that even if the critics are right, new policies should be tried so long as they have some chance of constituting a significant improvement over the status quo.

Reversible regulations provide a third area where the optimal search approach illuminates policymaking. At present, debates on regulation are between those who favor the cost-benefit approach, which advocates choosing

3. See, e.g., Ian Ayres & Eric Talley, *Solomonic Bargaining: Dividing a Legal Entitlement To Facilitate Coasean Trade*, 104 YALE L.J. 1027 (1995).

4. See, e.g., Louis Kaplow & Steven Shavell, *Do Liability Rules Facilitate Bargaining? A Reply to Ayres and Talley*, 105 YALE L.J. 221 (1995).

5. Default rules may fail because the assumptions underlying their supposed desirability prove faulty, or they may fail because of an unforeseen hindrance.

LEARNING THROUGH POLICY VARIATION

the best law on average, and those who favor the precautionary principle, which advocates caution in many regulatory choices. The best risk regulation from an optimal search perspective constitutes a modification of cost-benefit analysis that contrasts with the precautionary principle even more sharply than conventional cost-benefit analysis.

The optimal search approach also buttresses the argument of those advocating for policy experimentation at the state level and against wide-ranging federal preemption of state laws. In a federalist system, policy variance becomes even more desirable than in a national system as the learning benefits of variance are shared through space as well as through time. While nationally applicable policies will often maximize per-period outcomes, the preemption that often accompanies these laws stifles learning through variation to a degree underappreciated even by those who argue that states are the laboratories of democracy. As a result, the optimal search approach favors extremely limited preemption of state law.

While illuminating, the optimal search approach, which favors high-variance policies, relies on a number of assumptions. It assumes that policymakers can learn from their laws and that they can change these laws in response to their learning. When policies have irreversible effects, the benefits of variance in policies are greatly reduced. Indeed, when new policies are irreversible, the dynamic analysis emphasized by the optimal search approach indicates that variance is no longer positive, or even neutral, but rather negative.⁶ Burkean approaches thus have continued salience for policymakers when policy is examined in a dynamic context because Burkean approaches are optimal when policies are irreversible. Similarly, expected-value maximization rationalist approaches, which ignore variance, become more attractive for policymakers when policies are sticky but reversible, as the learning benefits of high variance and the flexibility benefits of low variance partially offset each other. In total, the choice of optimal policies depends critically on policy reversibility.

Policy reversibility has two sources. Some degree of irreversibility is inherent in all policies, while other sources of irreversibility arise from policymaking institutions. To gain the maximum benefits of the optimal search approach, this Article recommends institutional mechanisms that maximize

6. This is generally known as the “real options” approach to decisionmaking under uncertainty. For a comprehensive economic account, see AVINASH K. DIXIT & ROBERT S. PINDYCK, *INVESTMENT UNDER UNCERTAINTY* (1994). For legal applications, see IAN AYRES, *OPTIONAL LAW: THE STRUCTURE OF LEGAL ENTITLEMENTS* (2005); and Cass R. Sunstein, *Irreversible and Catastrophic*, 91 *CORNELL L. REV.* 841, 862-64 (2006). Irreversibility and the real options approach are discussed in detail in Part III.

reversibility, such as sunset clauses, unicameral legislatures, and a reduced emphasis on *stare decisis*.

In both federalist and uni-jurisdictional settings, the optimal search approach assumes that policymakers aim to maximize social welfare. The Article later relaxes this assumption and examines public policymakers' incentives to innovate in uni-jurisdictional and federalist contexts. Because policymakers' incentives to innovate are often lower than optimal in each context, several recommendations, such as subsidizing federalist innovation or emphasizing innovation in contexts with electoral insulation, may be justified.

This Article is organized as follows. Part I summarizes the Burkean, classical liberal/rationalist, and "experimentalist" approaches to public policymaking. Part II develops the optimal search perspective and demonstrates that high-risk policies with relatively low average outcomes often should be instituted before policies that other policymaking rationales classify as superior. Additionally, Part II examines how this optimal search insight relates to different variables, such as the choice of discount rates and the time required to evaluate a policy. It also develops the optimal search approach through idealized applications to risk regulation, contract law, corporate law, and federalism.

The remaining Parts of the Article address impediments to learning through policy variation via the optimal search process. Part III evaluates Burkean and rationalist objections to the optimal search perspective and discusses how the optimal search approach applies to policymaking and institutional design when the effects of policy changes are irreversible. Part IV relaxes the assumption that policymakers pursue socially beneficial policies and examines the optimal search approach from a public choice perspective. Part V reexamines the examples of the optimal search approach developed in Part II, taking account of irreversible policy effects and imperfect political incentives and modifying the approach's recommendations accordingly.

I. THEORIES OF PUBLIC POLICYMAKING

Classical liberals⁷ demonstrated a "faith in the capacity for human beings to rationally reconstruct the laws and institutions of social life [and] to reform

7. This Part does not provide a full review of the political philosophies of classical liberals or Burkeans. Rather, it briefly sketches their attitudes toward policymaking to demonstrate how the optimal search approach relates to these two schools of thought. For an article applying Burkean and liberal political philosophies to methods of legal interpretation, see Ernest Young, *Rediscovering Conservatism: Burkean Political Theory and Constitutional Interpretation*, 72 N.C. L. REV. 619 (1994).

LEARNING THROUGH POLICY VARIATION

and improve upon the ways of the past.”⁸ With this faith in the ability of reason to predict the effects of policies,⁹ policymaking for classical liberals follows a straightforward logic: choose the best policy.¹⁰ For utilitarians, the best policy is that which yields the greatest aggregate happiness or well-being.¹¹ Other metrics for the best policy are of course possible. Individuals concerned with distribution, for example, might seek a policy that prioritizes the welfare of the least well-off member of society.¹² Whatever the metric, the approach to choosing policy remains the same: choose the policy that will produce the best outcome under the given metric.

Edmund Burke rejected the liberals’ faith in their ability to predict the outcomes of policies.¹³ He stated, “We are afraid to put men to live and trade each on his own private stock of reason; because we suspect that this stock in each man is small”¹⁴ Instead of relying on reason and abstract theory, Burke emphasized tradition and past practice. Past practice has been tested. It is a “deliberate election of ages and of generations; it is a Constitution made by what is ten thousand times better than choice, it is made by the peculiar circumstances, occasions, tempers, dispositions, and moral, civil, and social habitudes of the people, which disclose themselves only in a long space of

-
8. RICHARD HUDELSON, *MODERN POLITICAL PHILOSOPHY* 41 (1999).
 9. If the impacts of a policy are uncertain, then reason should be able to predict the probabilities of certain outcomes.
 10. See John C. Harsanyi, *Cardinal Utility in Welfare Economics and in the Theory of Risk-Taking*, 61 *J. POL. ECON.* 304, 334-35 (1953) (suggesting “a close affinity between the cardinal utility concept of welfare economics and the cardinal utility concept of the theory of choices involving risk [expected utility maximization],” thereby implying that policy decisions from a welfarist perspective should maximize expected value or utility).
 11. See generally JEREMY BENTHAM, *AN INTRODUCTION TO THE PRINCIPLES OF MORALS AND LEGISLATION* 11-12 (J.H. Burns & H.L.A. Hart eds., Oxford Univ. Press 1996) (1789) (describing the “principle of utility” as “that principle which approves or disapproves of every action whatsoever, according to the tendency which it appears to have to augment or diminish the happiness of the party whose interest is in question”).
 12. See, e.g., JOHN RAWLS, *A THEORY OF JUSTICE* (1971).
 13. Some scholars argue that Burke was himself a classical liberal along many dimensions. Others claim that his reverence for tradition precluded his being a liberal. For a discussion of Burke’s work, see CONOR CRUISE O’BRIEN, *THE GREAT MELODY: A THEMATIC BIOGRAPHY AND COMMENTED ANTHOLOGY OF EDMUND BURKE* 605-18 (1992). For an application of Burke’s views to judicial decisionmaking, see Cass R. Sunstein, *Burkean Minimalism*, 105 *MICH. L. REV.* 353 (2006); and Young, *supra* note 7. The discussion of Burke’s thoughts in these articles influenced this Part.
 14. EDMUND BURKE, *REFLECTIONS ON THE REVOLUTION IN FRANCE* 74 (Frank M. Turner ed., Yale Univ. Press 2003) (1790).

time.”¹⁵ Burke was therefore an empiricist, preferring data about the effectiveness of policies rather than theoretical predictions.

Burke disliked radical change, which in his view is based on imperfect reason rather than experience: “Men little think how immorally they act in rashly meddling with what they do not understand.”¹⁶ Burke did not reject all change, however. He favored incremental change, with each step evaluated empirically before the next step is taken. Burke wrote,

By a slow but well sustained progress, the effect of each step is watched; the good or ill success of the first gives light to us in the second; and so, from light to light, we are conducted with safety through the whole series. We see that the parts of the system do not clash. The evils latent in the most promising contrivances are provided for as they arise. . . . We compensate, we reconcile, we balance. We are enabled to unite into a consistent whole the various anomalies and contending principles that are found in the minds and affairs of men.¹⁷

Burke’s recommendation for choosing policies therefore differs from the classical liberal approach. Burke argues that a new policy that is expected to yield benefits to society should not be implemented in its entirety immediately. Instead, Burke argues that a portion of the policy should be implemented and then evaluated. If the reform appears successful, then an additional element of the policy should be implemented. This process should continue until some element of the policy either fails or the entire policy reform is implemented.¹⁸

The differences between the classical liberal/rationalist approach to policymaking and the Burkean/conservative approach to policymaking can be seen below in Example 1. Suppose that policymakers must choose between existing Policy A and new Policy B. Policy A costs nothing to implement and

15. Edmund Burke, Speech on the Reform of the Representation of the Commons in Parliament (May 7, 1782), in 4 SELECT WORKS OF EDMUND BURKE: MISCELLANEOUS WRITINGS 15, 21 (E.J. Payne ed., 1990).

16. EDMUND BURKE, AN APPEAL FROM THE NEW TO THE OLD WHIGS 113 (London, J. Dodsley 1791). In this account, the value of past practice is very much empirical. Past practice has worked and is therefore owed deference. There are other values associated with Burke’s fealty to tradition, of course. For a comprehensive discussion, see Anthony T. Kronman, *Precedent and Tradition*, 99 YALE L.J. 1029 (1990).

17. BURKE, *supra* note 14, at 143-44.

18. If some part of a policy reform fails, then policymakers should either return to the old status quo, or—if a partial reform yields better outcomes than the old status quo—maintain the partially reformed policy as a new status quo.

LEARNING THROUGH POLICY VARIATION

yields social benefits of \$50 per period. Policy *B* is more innovative and costs \$20 to implement. It yields benefits of \$120 per period with a probability of 0.6 and benefits of \$0 with a probability of 0.4. While these examples use dollars to measure benefits for ease of exposition, it is important to remember that the true benefits are measured in terms of aggregate welfare.¹⁹ Only if social outcomes can be measured according to Kaldor-Hicks utility terms can dollar values across persons be aggregated by merely summing changes in wealth across individuals.²⁰

Example 1.

POLICY	A	B
Cost	\$0	\$20
Expected Benefit per Period	\$50	$\$72 = 0.6(\$120) + 0.4(\$0)$

A rationalist prefers *B* to *A* in this example. Policy *B* provides an expected benefit of \$72 and an expected cost of \$20, for a net benefit of \$52. This is greater than Policy *A*'s net benefit of \$50. Policy *B* produces a greater expected benefit than Policy *A*, so *B* should be chosen.

A Burkean, however, would probably prefer *A* to *B*. The Burkean would suppose that the expected benefits of the untested Policy *B* are nearly impossible to estimate through rational analysis. Furthermore, the downside risk of untested Policy *B* is likely to be underestimated. Past practice (Policy *A*) may illustrate some ill-defined but nonetheless important public preferences that policymakers may be unable to pinpoint. Therefore, the negative

-
19. These benefits and costs can be quantified by any social welfare metric. For example, the benefits in dollars may measure the dollar benefits to the least well-off member of society (a social welfare function advocated by John Rawls) rather than measuring total wealth. More generally, the benefits and costs could account for risk aversion in a rationalist framework, as expected utility theory does within economic science. This would make the Burkean and rationalist frameworks more similar to each other, although Burke would presumably doubt the ability of the rationalist to appropriately measure risk aversion. Throughout this Article, I generally adopt a utilitarian framework.
20. Kaldor-Hicks efficiency proposes that a more socially efficient outcome can be reached by moving from the Pareto optimal level if those who would be made better off by the outcome compensate those who would be made worse off by the outcome. For a useful discussion of Kaldor-Hicks efficiency, see Guido Calabresi, *The Pointlessness of Pareto: Carrying Coase Further*, 100 YALE L.J. 1211 (1991).

consequences of changing *A*—and no longer reflecting public preferences for past practice—may be more significant than the rationalist would suppose. In light of this overlooked downside, the potential net gain from choosing Policy *B* (\$2 in this period) is simply not worth the risk.²¹

The Burkean conservative would not necessarily shun every new policy. Consider Policy *A'* that is very similar to *A* but one step in the direction of Policy *B*. *A'* has no cost of implementation, a 0.6 chance of yielding benefits of \$52, and a 0.4 chance of yielding benefits of \$49. *A'* would probably be preferred to *A* under the Burkean framework. Policy *A'* would be implemented. If it succeeded, then Burkeans would contend that *A''*, which is another step in the direction of *B*, should also be implemented. In this way, Burkeans hope to attain the possible benefits of Policy *B*, without incurring the downside risks that accompany drastic policy changes.

Rationalist/liberals have identified several critiques of the Burkean position.²² They question the assumption that Policy *A* is necessarily wise because it has been tried. If experts believe that an untried position like *B* is better, then the rationalists see no reason to delay implementation out of respect for the traditional Policy *A*. First, there is no guarantee that incrementally different policies like *A'* and *A''* exist—some policies simply cannot be adopted piecemeal. When this is the case, Burkean conservatism will often lead to inferior outcomes. Policy *B* has a higher expected value than Policy *A*, but the Burkean will continue to favor *A*. Second, even if such policies exist, the Burkean approach is inferior because it takes a considerable number of periods to yield *B*'s benefits; Burkean reform is purposefully deliberate, which is detrimental if reform is warranted.

Another strain of policymaking theory that is closely related to the approach developed in this Article emphasizes pragmatic learning, experimentalism, and dynamic decisionmaking.²³ This emphasis occurs most

21. An extremely risk-averse classical liberal might advocate cautious policymaking approaches that conform with the Burkean approach. Nevertheless, there is a difference between the sources of caution. For the liberal, caution is a product of a risk-averse response to a probability distribution of different policy outcomes. For the Burkean, caution is warranted by the tendency to underestimate the probability of bad outcomes.

22. See, e.g., Adrian Vermeule, *Many-Minds Arguments in Legal Theory* 11-12 (Harvard Law Sch. Pub. Law & Theory, Working Paper No. 08-02, 2008), available at <http://ssrn.com/abstract=1087017>.

23. See, e.g., Michael C. Dorf & Charles F. Sabel, *A Constitution of Democratic Experimentalism*, 98 COLUM. L. REV. 267 (1998); Daniel A. Farber, *Environmental Protection as a Learning Experience*, 27 LOY. L.A. L. REV. 791 (1994); Bradley C. Karkkainen, *Environmental Lawyering in the Age of Collaboration*, 2002 WIS. L. REV. 555. For a review of this literature, see Orly

LEARNING THROUGH POLICY VARIATION

frequently in environmental policy analysis. Daniel Farber sums up this attitude as follows: “Rather than viewing [environmental] policy making as a one-shot exercise, in which the goal is to adopt the optimum solution based on current information, we might do better to think of a continuous process of learning and experimentation.”²⁴

“Experimentalists” such as Farber, Charles Sabel, and Michael Dorf typically extol the virtues of federalism (and of other forms of decentralized decisionmaking) because of its learning benefits.²⁵ In addition, scholars who examine environmental decisionmaking from a dynamic context often stress the value of delaying high-variance innovative decisionmaking to await future knowledge.²⁶ Thus, the dynamic decision-making context is often used to justify Burkean-type precautions in environmental contexts.²⁷

With this extremely brief sketch of various approaches to policymaking as a background, I present an approach to policymaking—the optimal search approach—that combines elements of rationalism, Burkean conservatism, and experimentalism to generate some initially counterintuitive results about the desirability of certain policies.

II. CHOOSING POLICIES FROM AN OPTIMAL SEARCH PERSPECTIVE: THEORY AND APPLICATIONS

Policies and their effects are not static.²⁸ A policy’s performance in one period yields information about its probable performance in the next period. In fact, it would be surprising if this information had no bearing on the probability of the policy continuing in future periods. In other words, policies

Lobel, *The Renew Deal: The Fall of Regulation and the Rise of Governance in Contemporary Legal Thought*, 89 MINN. L. REV. 342 (2004).

24. Farber, *supra* note 23, at 791.

25. For a thorough examination of experimentalism in a federalist system, see *infra* Subsection II.C.4.

26. See, e.g., Farber, *supra* note 23, at 803 (“If a decision has irreparable consequences, then it may be worth delaying the decision in order to obtain new information. Taking an irreversible step forecloses the possibility of future learning, and therefore incurs an extra cost that does not show up in the usual cost-benefit analysis. In a formal sense waiting is equivalent to purchasing an option contract, and under many circumstances that option has positive value.” (footnote omitted)). For a detailed discussion, see *infra* Section III.C.

27. For example, when new policies are irreversible and the outcome of *B* can be determined without actually implementing *B*, dynamic analysis suggests that policy *A* is preferred to policy *B*. See Sunstein, *supra* note 6; *infra* Part III.

28. Even if there is policy inertia, policies are not set in stone. For a detailed discussion, see *infra* Part III.

are not chosen and then simply retained. Picking an optimal policy entails choosing a policy in one period and then reoptimizing the policy in a future period given the additional information obtained since the prior policy decision. When accounting for reoptimization, policy choices might be very different than they would be absent reoptimization. Instead of choosing the best policy for the present, policymakers must balance the choice between good policies in the current period with the ability of this period's policy choice to improve future policy choices. If a current policy is expected to be bad, but potentially improves the future policy choice set, then it may be a better choice than a policy that has high expected value in the current period but adds little to the menu of options in the future.

A. The Theory of Optimal Search

The experimentalists described in Part I develop elaborate systems of decentralized decisionmaking to foster dynamic learning.²⁹ Here, I take a different approach. I aim to normatively characterize optimal policymaking under a number of different conditions. I develop the optimal search approach to be applied when policies have reversible effects and learning is possible. The optimal search approach values high-variance policies to a much greater degree than static rationalism or Burkean conservatism.

In other words, static rationalism advocates choosing the policy with the highest expected value, ignoring the variance of the policy. Burkean conservatism warns against policies with considerable variance and encourages incremental policy change. Dynamic experimentalists provide few normative recommendations about the content of policy, focusing instead on the organizational setting in which policy is made.

The optimal search approach, by contrast, insists that high variance is a valuable feature of a policy. High-variance policies offer the potential to find excellent policies for future periods. Therefore, high-variance policies will be better choices than low-variance policies of the same or (even higher) expected value whenever changing policies is a feasible possibility.

The optimal search approach constitutes a modification of the rationalist approach. When policies can be changed without imposing irreversible effects and learning is possible, the optimal search approach produces the best policy

29. See, e.g., Dorf & Sabel, *supra* note 23, at 314 (advocating the “construct[ion of] the organizational rudiments of local, or, rather, subnational, pragmatist government, by transposing to the public sphere the institutions of benchmarking, simultaneous engineering, and error detection”).

LEARNING THROUGH POLICY VARIATION

for rationalists. Put simply, the optimal search approach emphasizes how the rational policy calculus changes with learning and reoptimization.

These insights are illustrated by modifying the hypothetical policy decision presented in Example 1. Suppose that a policymaker faces a choice between two policies with uncertain values but with known distributions of potential payoffs.³⁰ The policies must be implemented to resolve the uncertainty regarding their efficacy and to determine the policy payoff per period, which can be represented by v . Suppose that the policy must be tested for a year (the length of a period) and that the annual discount rate is ten percent.³¹ Any policy has a cost, c , of implementation. This cost represents the cost to society of adjusting and implementing a policy. Suppose further that once a policy has been implemented, it is costless to return to that policy in a future period.³² In other words, in each period a policymaker has the option of choosing any previously tested policy or a new policy at some positive switching cost.

Consider the choice between two policies, A and B , summarized in the table for Example 2. Policy A is a relatively safe policy because it has been attempted before.³³ Policy A costs nothing to implement ($c = 0$), and yields social benefits of \$50 per period ($v = 50$). Policy B is a more innovative policy and costs \$20 to implement.³⁴ It yields benefits of \$120 per period with a probability of 0.2 and benefits of \$0 with a probability of 0.8.

30. This example parallels the one found in Weitzman, *supra* note 2, at 641-43.

31. None of the results in this Article depends upon these assumptions, and they will be relaxed or altered in subsequent Sections.

32. This assumption is obviously too strong. For many reasons, including the political system and status quo bias, the existing policy is cheapest and easiest to implement in the next period. Previously implemented policies are the next cheapest, and brand new policies are presumably the most expensive. The consequences of the policy inertia that is imposed by this distribution of implementation costs are discussed in detail in Part III.

33. Alternatively, Policy A may be a very low-risk policy that has not been implemented before.

34. One might ask why Policy B has never been attempted before. There are several responses. One answer may be that Policy B had never been thought of before. A second may be that there has been some potentially permanent change in the environment that changes the payoff distribution of Policy B , making Policy B an essentially new policy.

Example 2.

POLICY	A	B
Cost	\$0	\$20
Expected Benefit per Period	\$50	\$24 = 0.2(\$120)+0.8(\$0)

The choice may seem obvious: Policy *A* is better than Policy *B*. Policy *A* has a considerably higher expected social benefit per period ($\$50 > \24) and lower cost ($\$0 < \20) than Policy *B*, making *A* preferable to the rationalist/liberal. Policy *A* also entails no risk, while Policy *B* is quite risky, making *A* preferable from the Burkean/conservative position. Indeed, one might be skeptical of any legal theory that advocates Policy *B* when the choice between *A* and *B* appears so lopsided. Yet the optimal search theory of policy decisions indicates that *B* should be tried before *A*.

Consider first the payoff for choosing Policy *A*. If one chooses Policy *A* in this period, one should choose it in every period. Because there is no uncertainty resolved by choosing *A*, the information set of the policymaker who chooses *A* in this period will be the same in the next period. Therefore, if *A* is preferred in this period, then it must also be preferred in the next period.

The payoff from Policy *A* is \$50 per period. Policy *A* thus has a present discounted value of

$$\sum_{t=1}^{\infty} \frac{\$50}{(1+0.1)^t} = \$500$$

Now consider the expected payoff in all periods in the future from trying *B*. The expected payoff is as follows. By choosing Policy *B*, society incurs a cost of \$20. With a probability of 0.2, the policy succeeds and yields its benefits into the indefinite future. With a probability of 0.8, Policy *B* fails and yields no benefits this period. If Policy *B* fails, the policymaker should change policies to Policy *A*. Policy *A*'s value is as given above, though it is discounted because it arrives one period later. The expected lifetime payoff from choosing Policy *B* in this period is therefore

$$-20 + 0.2 \left[\$120 + \sum_{t=2}^{\infty} \frac{\$120}{(1+0.1)^t} \right] + 0.8 \left[0 + \sum_{t=2}^{\infty} \frac{\$50}{(1+0.1)^t} \right] = \$585.82$$

LEARNING THROUGH POLICY VARIATION

The expected payoff from Policy *B* exceeds the benefit of Policy *A*, even though *B* appears inferior to Policy *A* along every possible dimension.

So why does *B* “beat” *A*? There is a difference between determining the best policy and choosing the policy to attempt first. When choosing the policy to attempt first, policymakers must consider not only the expected benefits and costs of the policy in the current period, but also the potential for realizing an extremely good policy that can be continued in future periods. Even though *A* performs better than *B* on average, *B* offers the potential for a much better outcome than *A*. Therefore, *B* should be tried first even though there is a high likelihood that *B* will fail.³⁵ If *B* works, then policymakers have found an excellent policy and will retain *B*. If *B* fails, then policymakers choose *A* for the remaining periods.³⁶

B. Factors Influencing the Optimal Degree of Policy Variance

The relative desirability of *B* as compared to *A* depends upon the values of a number of variables. First, it should be emphasized that the average outcomes of *B* and *A* matter, even though they are not determinative. While the values of *B* and *A* in the search program for an optimal policy depend primarily on their upsides, the opportunity costs of implementing *B* and *A* for the current period depend upon their average outcomes. As *B*'s average outcome deteriorates, *B* becomes less desirable relative to *A*, even when holding *B*'s upside constant.

Second, high implementation costs reduce the desirability of high-variance policy innovations such as *B*. For example, suppose the costs of implementing *B* are \$200 rather than \$20 because the costs of educating the population about the new policy are extremely high. In this case, Policy *A* becomes more desirable than Policy *B*—some policy changes are simply not worth the switching costs. Note, however, that this altered example requires switching costs to be over forty percent of the infinite real effects of the policy. While some policies might have such high switching costs, other policies' switching costs do not rise to such a high proportion of total effects. In particular, important policies are likely to have proportionately lower switching costs.

35. Throughout the remainder of this Article, terms like “*B* being better than *A*” or “*B* beating *A*” will mean that *B* should be chosen before *A*. At no point do I mean to suggest that *B* is better than *A* in every period; I suggest only that *B* should be chosen first under many circumstances in spite of its inferiority in expected-value terms.

36. For a rigorous proof of this argument, see Weitzman, *supra* note 2.

Third, high-risk policies such as *B* become more desirable relative to low-risk policies like *A* when the discount rate is lower—that is, when the future becomes more important relative to the present. Policy *B*'s value lies primarily in the upside policy potential that might be enjoyed in future periods. When discount rates are low, this value is discounted to a lesser extent. Therefore, Policy *B*'s value rises relative to Policy *A*. Suppose that the discount rate were 0.05 in our example rather than 0.1. In this case, the value of choosing Policy *A* becomes \$1000, while the value of choosing Policy *B* becomes \$1225.82. The difference in value between choosing *B* and choosing *A* grows from about \$86 to \$226.

Fourth, Policy *B* becomes more attractive when the probability of its success is correlated with the probability of other policies' success. If *B*'s success or failure tells us information about policies *C*, *D*, and *E*, then *B* becomes increasingly attractive. Not only does implementing *B* first provide information about *B*'s potential in future periods, it also provides information about the value of using *C*, *D*, and *E* in the future. In other words, *B* provides substantial information to facilitate the policy search. The value of this information should be added to the value of choosing *B* first when deciding whether to implement *B*.

Fifth, Policy *B* becomes less attractive when it takes longer to evaluate. If Policy *B* takes two periods to evaluate, then its low expected value will be felt for more than one period. Consider the value of Policy *B* if it must be in effect for four periods before it can be evaluated. The expected value of choosing Policy *B* first thereby becomes

$$-20 + 0.2 \left[\sum_{t=1}^{\infty} \frac{\$120}{(1+0.1)^t} \right] + 0.8 \left[0 + 0 + 0 + 0 + \sum_{t=5}^{\infty} \frac{\$50}{(1+0.1)^t} \right] = \$495.06$$

Policy *B* should not be tried before Policy *A* when Policy *B* requires four or more periods to evaluate, in spite of the potential for an excellent outcome under Policy *B*. The low expected value of Policy *B* for the first four periods trumps the high upside of Policy *B*.³⁷

Sixth, the optimal search approach does not imply that innovative policies should always be chosen. Suppose Policy *B* is chosen in this period and proves to work. In this case, the search for an effective policy comes to an end; the

37. This hypothetical is further examined in Section III.C.

LEARNING THROUGH POLICY VARIATION

costs of continuing to search are not worth the benefits because an effective policy has been found. Thus, the optimal search approach does not require policy innovation. It merely points out that innovation has long-run benefits that are not captured by the expected outcome of the innovation.

The superiority of *B* to *A* is subject to a number of critiques that will be explored in detail in Sections III.A and III.B. Before examining these complications, the next Section will examine the usefulness of the optimal search approach in debates about risk regulation, contract law, corporate law, and federalism.

C. The Optimal Search Approach: Applications

This Section examines how the optimal search approach might be applied in some important contexts. It is important to emphasize, however, that none of the following examples constitutes a comprehensive case study. Indeed, each application developed in this Section does not address many impediments to a full implementation of the optimal search perspective.³⁸ Instead of exhaustively describing how policymaking decisions should be made, these applications are intended to illustrate the potential value of adding an optimal search calculus to the existing criteria for choosing policies.

1. Choosing Between Reversible Regulations

The optimal search approach illuminates some of the debates regarding reversible risk regulation.³⁹ Reversible risks are risks that, if and when they materialize, are confined to one period. Two approaches—cost-benefit analysis and the precautionary principle—frame most discussions of risk regulation.⁴⁰

It should be emphasized at the outset of this Subsection that the debate between the precautionary principle and cost-benefit analysis is most salient when risks are irreversible—a situation that will be discussed below. Nevertheless, many risks are not irreversible and thus are conducive to analysis via the optimal search approach.⁴¹

38. Some of the most important complications are examined in Part III.

39. The problem of irreversible risks is discussed in detail in Part III.

40. See Douglas A. Kysar, *It Might Have Been: Risk, Precaution, and Opportunity Costs* 3-4 (Cornell Law Sch. Legal Studies Research Paper Series, Research Paper No. 06-023, 2006), available at <http://ssrn.com/abstract=927995>.

41. Cost-benefit analysis does not obviously distinguish between reversible and irreversible risks. In addition, “some formulations [of the precautionary principle] . . . apply to even

Cost-benefit analysis sums the costs and benefits of different policies and instructs that the policy with the greatest net benefit should be chosen. In this regard, the intellectual underpinnings of cost-benefit analysis lie in the classical liberal/rationalist approach to policy. Cost-benefit analysis recognizes that the actual effects of many policies are uncertain. The response of cost-benefit analysis is to use expected values. The expected value of a policy is calculated by multiplying the value of an outcome by the probability of its occurrence for each possible outcome, and then summing up the resulting products.⁴² The policy with the highest net expected value is the one preferred by cost-benefit analysis.

The precautionary principle “seeks to trigger an incremental process of risk regulation through the simple admonition, ‘When an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause-and-effect relationships are not fully established.’”⁴³ The precautionary principle’s intellectual foundations are very similar to those of Burkean conservatism in that they emphasize the inability of rational science to properly anticipate the effects of changes in the environment.⁴⁴

theoretically reversible risks.” David A. Dana, *A Behavioral Economic Defense of the Precautionary Principle*, 97 NW. U. L. REV. 1315, 1315-16 (2003).

42. See Amartya Sen, *Discipline of Cost-Benefit Analysis*, 29 J. LEGAL STUD. 931 (2000).

43. Kysar, *supra* note 40, at 3-4 (quoting Peter Montague, *The Precautionary Principle*, RACHEL’S ENV’T & HEALTH WKLY., Feb. 19, 1998, at 1, but mistranscribing “cause and effect” as “cause-and-effect”). Note that there are many different articulations of the precautionary principle. See, e.g., Sunstein, *supra* note 6, at 848-50. The one stated here is merely one oft-cited statement of the principle. A full characterization of the precautionary principle or of cost-benefit analysis is beyond the scope of this paper.

44. See C.A. BOWERS, MINDFUL CONSERVATISM: RETHINKING THE IDEOLOGICAL AND EDUCATIONAL BASIS OF AN ECOLOGICALLY SUSTAINABLE FUTURE 91-94 (2003) (drawing a link between environmental conservation and Burke’s approach). Curiously, while legal scholars have spilled considerable ink on the precautionary principle (a search in Westlaw’s Journals and Law Reviews database for [“precautionary principle”] on October 31, 2008 yielded over 2100 documents) and on Burke (a search in Westlaw’s Journals and Law Reviews database for [Edmund /5 Burke!] on October 31, 2008 yielded over 1400 documents), they have not explored the seemingly obvious linkage between Burkean political philosophy and the precautionary principle. A search in Westlaw’s Journals and Law Reviews database for [“precautionary principle” /50 Burke!] on October 31, 2008 yielded only six documents, five of which are related to a different Burke and the last of which makes only a casual allusion to the similarity between the precautionary principle and Burkean conservatism. See Daniel A. Farber, *From Here to Eternity: Environmental Law and Future Generations*, 2003 U. ILL. L. REV. 289, 305-06 n.90. It is important to note that Westlaw’s Journals and Law Reviews database neither catalogues all volumes of currently published law journals, nor includes scholarship in non-law disciplines.

LEARNING THROUGH POLICY VARIATION

Consider how these two approaches would examine the choice between Policies *A* and *B* described in Example 2. Suppose that Policy *A* would implement an existing technology approved by regulators and that Policy *B* would implement a new technology for reducing a pollutant. Regulators must consider whether to permit the new technology implemented by Policy *B*. The *B* technology may do a better job of eliminating the pollutant than the existing *A* technology, but it will probably do an inferior job. Permitting the *B* technology is expensive because it requires issuance of new regulations and installation of the new technology. Suppose further that the pollutant's harms are felt in this period only, that the harms of the pollutant affected by the technology are well known, and that the technology can be properly tested only by approving it for use.

Simple cost-benefit analysis—ignoring informational or learning value and simply taking the expected per-period value of each technology—makes the choice between *A* and *B* an easy one. Permitting the *B* technology entails a higher cost (\$20 for *B* versus \$0 for *A*) than the old technology⁴⁵ and a lower expected benefit⁴⁶ (\$24 of pollution reduction for *B* versus \$50 of pollution reduction for *A*). The net benefits of *A* are higher than *B*. Therefore Policy *A*—favoring the status quo technology—should be chosen.

The precautionary principle also favors Policy *A*. Policy *B* “raises threats of harm to human health”⁴⁷—in 80% of cases Policy *B* does a poorer job of reducing pollutants than Policy *A*. Policy *B* is less incremental than Policy *A*. Policy *B* entails trying a new technology, a nonincremental maneuver, while Policy *A* involves no alteration of population risk. Therefore, the precautionary principle dictates that Policy *A* be chosen, particularly in light of the fact that the Policy *B* technology appears inferior along a number of dimensions, including expected benefits. When learning is possible, both per-period cost-benefit analysis and the precautionary principle produce inferior outcomes relative to the optimal search approach. Regulators should choose Policy *B* because of its potential to produce an excellent outcome over the long run.

The optimal search approach constitutes a modification of traditional cost-benefit analysis. Instead of taking expected values of various regulations and choosing the policy with the highest expected value (implicitly assuming that the policy will continue indefinitely), policymakers should conduct cost-benefit analysis in a dynamic context where learning is possible. In this context, not

45. Both issuing new regulations and allowing some polluters to change their technology for pollution reduction entail costs.

46. The benefits of either technology are realized in the form of reduced pollution.

47. See *supra* note 43 and accompanying text.

only the per-period expected value, but also its distribution across periods, matters. A long-run cost-benefit analysis, including all periods and the possibility of switching policies after learning about the impact of policies, may favor a policy that has low expected value but a wide distribution (such as Policy *B*) over a policy that appears better along many dimensions.

Executive Order 12,866 requires cost-benefit assessment of agencies' regulatory actions.⁴⁸ The Office of Management and Budget's *Circular No. A-4* sets guidelines governing this cost-benefit analysis.⁴⁹ The benefit estimates required by the Office of Management and Budget (OMB) do not place a value on the information provided by a particular regulation.⁵⁰ Instead, they measure only the direct benefits of a regulation, such as the amount of pollution reduction. The informational value of a regulation might be large in relation to the direct benefits. For example, Policy *B* provides expected benefits of only \$24 per period, but learning that *B* works better than expected allows us to gain *B*'s benefits in all future periods. Policy *A*, by contrast, provides no valuable information. This informational value can be estimated as the difference between *B*'s upside and *A*'s upside. Over the long run, this difference in value is considerable. Indeed, in our example it is enough to overcome *A*'s superiority in the average case (see Example 2). *B* makes better use of resources than *A*, despite *A*'s seeming superiority. By requiring that agencies place a value on the information provided by a regulation through the use of dynamic reoptimization models rather than assuming that the proposed regulation will be enacted indefinitely, the OMB could achieve better "resource allocation."⁵¹ The optimal search approach turns the precautionary principle applied to reversible risks on its head.⁵² The precautionary principle advocates the avoidance of risk; the optimal search approach seeks it out. The ability to change policies limits the costs of high-variance policies in the optimal search approach. If a regulation proves to be a bad one (the downside risk is realized),

48. Exec. Order No. 12,866 § 6(a)(3)(c), 3 C.F.R. 638 (1993), *reprinted in* 5 U.S.C. § 601 (2000).

49. OFFICE OF MGMT. & BUDGET, CIRCULAR A-4, REGULATORY ANALYSIS (2003), *available at* <http://www.whitehouse.gov/omb/circulars/a004/a-4.pdf>.

50. The OMB Circular makes no mention of measuring the information benefit estimates of a particular regulation under its guidelines. *Id.*

51. *Id.* at 1.

52. The precautionary principle often assumes that probabilities cannot be estimated, a condition called uncertainty. *See* Kysar, *supra* note 40, at 10-12. So long as the risks involved are not irreversible, uncertainty does not require drastic changes in the lessons of the optimal search approach. With uncertain but reversible risks, one can assume the worst—that is, place heavy probability on bad outcomes—and still favor risky policies, rather than safe ones.

LEARNING THROUGH POLICY VARIATION

it can be changed. Therefore, reversible but risky regulations have a downside limited to one period but an upside that can be realized indefinitely, making risk relatively attractive in this context.

One critique of standard cost-benefit analysis focuses on its use of high discount rates to discount the future and its sensitivity to the choice of discount rate.⁵³ At least one critic has considered the possibility of a discount rate of zero.⁵⁴ These critics' arguments complicate the case for ordinary cost-benefit analysis but add heft to the argument for choosing high-variance reversible regulations under the optimal search approach. If the discount rate is zero, for example, then the value of the upside of Policy *B*'s technology grows. This high upside will be realized in all future periods in which *B* proves to be a success. The less these future periods are discounted, the more emphasis should be placed on the option value of the technology. If there is no discounting, then Policy *B* becomes an almost necessary choice. The possibility of enjoying lower risks to health in an infinite number of undiscounted future periods trumps the small increase in risk associated with the *B* technology in the current period.

2. *The Optimal Search Approach and Penalty Default Rules in Contract Law*

Economic analysis of contract law seeks to provide default rules that maximize contractual surplus.⁵⁵ One purpose of contractual defaults is to save "parties the time, trouble, and risk of error implicit in crafting their own contract terms."⁵⁶ Thus, many scholars advocate "majoritarian" default rules, in which policymakers choose the default rule that would be chosen by most parties considering an issue. These rules save the greatest number of parties the cost of drafting a precise contract.

In separate articles, Ian Ayres and Robert Gertner⁵⁷ and Lucian Bebchuk and Steven Shavell⁵⁸ have developed a different rationale for policymakers' choice of default rules. In some cases, default rules should not be majoritarian but rather information forcing. An information-forcing rule compels parties

53. See Lisa Heinzerling, *Regulatory Costs of Mythic Proportions*, 107 YALE L.J. 1981, 2043-56 (1998).

54. See *id.*

55. See ROBERT E. SCOTT & JODY S. KRAUS, *CONTRACT LAW AND THEORY* 26-27 (4th ed. 2007).

56. *Id.* at 90.

57. See Ian Ayres & Robert Gertner, *Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules*, 99 YALE L.J. 87 (1989).

58. See Lucian Ayre Bebchuk & Steven Shavell, *Information and the Scope of Liability for Breach of Contract: The Rule of Hadley v. Baxendale*, 7 J.L. ECON. & ORG. 284 (1991).

with superior information to divulge this information because the default rule, which would apply if the party did not divulge this information, is crafted to work against the party with the superior information.

Much of the development of the perspective favoring information-forcing rules concerns the celebrated case of *Hadley v. Baxendale*.⁵⁹ *Hadley* developed the rule that promisors are not liable for consequential damages caused by their breaches of contract. Ayres and Gertner and Bebhuk and Shavell argue that there is no indication that most parties would prefer the *Hadley* rule, making *Hadley* nonmajoritarian.⁶⁰ Instead, they claim that *Hadley* constitutes an information-forcing rule.⁶¹ Promisees possess better information than promisors regarding the probability of consequential damages. The *Hadley* rule limiting consequential damages forces a promisee facing substantial consequential damages (a “high type”) to reveal this information to a promisor. If the promisee does not reveal this information, then the promisor will take ordinary precautions, secure in the knowledge that the *Hadley* rule protects against unusually large consequential damages. Under certain conditions, high consequential damages promisees reveal their “type,” while low consequential damages promisees (“low types”) remain silent. Promisors take the right precautions for each type, without having to go through the trouble of investigating the potential consequential damages for every promisee. If there are more low types than high types and communication costs are sufficiently low, the *Hadley* information-forcing rule produces the most efficient outcome.

By demonstrating that the “majoritarian” default rule paradigm is sometimes inferior, the information-forcing or “penalty” default rule literature constituted a major advance in contract law scholarship. From a policy perspective, however, information-forcing rules have fared less well.⁶² The problem is simple: information-forcing rules trump majoritarian rules under

59. *Hadley v. Baxendale*, (1854) 9 Ex. 341, 156 Eng. Rep. 145; see Barry E. Adler, *The Questionable Ascent of Hadley v. Baxendale*, 51 STAN. L. REV. 1547, 1547-48 (1999). Following the literature, this discussion will emphasize the *Hadley* rule. The same conclusions apply to other information-forcing or penalty default rules.

60. See Ayres & Gertner, *supra* note 57, at 112-15; Bebhuk & Shavell, *supra* note 58, at 308. Even if the *Hadley* rule would be the choice of most parties considering the rule, the rule was not chosen for this reason.

61. See Ayres & Gertner, *supra* note 57, at 106-07.

62. See Eric A. Posner, *There Are No Penalty Default Rules in Contract Law*, 33 FLA. ST. U. L. REV. 563 (2006) (arguing that penalty default rules do not exist in practice). *But see* Ian Ayres, *Ya-Huh: There Are and Should Be Penalty Defaults*, 33 FLA. ST. U. L. REV. 589 (2006) (contradicting Posner’s assertions and claiming that there are many examples of penalty default rules).

LEARNING THROUGH POLICY VARIATION

some assumptions, but they fall short under others. First, if communication costs prevent promisors from communicating their types to promisees, then the *Hadley* rule falters—it induces promisors to take precautions assuming that damages will be low, even though there are some high-damage promisees who cannot indicate their type.⁶³ Similarly, if no one knows about a penalty default rule, then the rule cannot serve its information-forcing function. In the case of ignorance, a majoritarian rule that gives most parties their desired rule is preferable.

More generally, the *Hadley* rule’s efficiency depends on its ability to induce a “separating equilibrium,” wherein promisors know each promisee’s type.⁶⁴ An alternative equilibrium is a “pooling equilibrium,” wherein promisees are undifferentiated.⁶⁵ In a pooling equilibrium, the information-forcing rule fails to force information. The *Hadley* rule results in a pooling equilibrium under many conditions, including high communication costs as well as a continuous distribution of types, rather than the discrete “high type/low type” framework of the seminal *Hadley* models. Whenever the *Hadley* rule fails to induce separation, the majoritarian rule eclipses the information-forcing rule. Unlike an information-forcing rule, the majoritarian rule ensures that most parties in an undifferentiated pooling equilibrium enjoy their preferred rule.⁶⁶

Second, the *Hadley* rule’s efficiency hinges on the relative percentages of high and low types. If there are more high types than low types, then the *Hadley* rule forces considerable amounts of costly information sharing. By contrast, a consequential damages rule might induce low types to agree to a contractual consequential damages limitation. When there are fewer low types than high types, the consequential damages rule achieves the same result as the *Hadley* rule—the right precautions for each type—while economizing on communication costs. In sum, the *Hadley* rule’s efficiency depends upon a number of hard-to-verify assumptions, making “an accurate evaluation of a penalty-default rule’s efficacy in the *Hadley* setting . . . a heroic task.”⁶⁷

In the face of uncertainty about whether a penalty rule will serve its desired information-forcing function, a policymaker considering a penalty default rule faces the following calculus. Under some conditions, the penalty default rule enhances efficiency by forcing information sharing. Under other conditions,

63. See Bebchuk & Shavell, *supra* note 58, at 302 proposition 6(b) (demonstrating the inefficiency of the *Hadley* rule when communication between parties is undesirable).

64. See Adler, *supra* note 59, at 1547-48.

65. *Id.* at 1556.

66. See *id.* at 1559-70.

67. *Id.* at 1552.

such as ignorance of the penalty, the penalty default rule leaves many parties with an undesirable rule. A majoritarian rule, by contrast, has a more limited downside and upside. Even though the majoritarian rule does not force efficient information sharing, the policymaker can at least be certain that the rule gives most parties their desired rule. A majoritarian rule entails less risk than a penalty rule and may well provide a better average outcome than a penalty rule.

The majoritarian rule entails even less risk when it is the existing default rule.⁶⁸ Not only does the majoritarian rule reduce risk by granting most parties what they want, the majoritarian rule also has the advantage of being tried and true—its hidden deficiencies have been revealed by practice. Given these realities, it may therefore come as no surprise that some scholars argue that “there is no such thing as a penalty default rule.”⁶⁹

The optimal search approach, however, suggests that penalty default rules may be more desirable than generally realized. Returning to the example above, a majoritarian rule corresponds to Policy *A* in the example—it is low risk, but with no potential for information sharing. A penalty default rule corresponds to Policy *B*. Penalty default rules are riskier than majoritarian rules but provide some chance of genuinely superior outcomes in at least some contexts.⁷⁰ The example demonstrates that, in some cases, penalty default rules (corresponding to Policy *B*) should be implemented before majoritarian rules, even if these majoritarian rules provide better average outcomes per period. Penalty defaults provide the potential for continuing with an effective information-forcing Policy *B* in future periods. The value of this upside potential to future periods may well be great enough to indicate that the penalty rule should be tried first, even if there is considerable uncertainty about whether the penalty rule actually enhances efficiency.⁷¹

68. Since most contract default rules have been chosen using the majoritarian paradigm, see SCOTT & KRAUS, *supra* note 55, at 89-90, the majoritarian rule is the existing default law in most cases.

69. Posner, *supra* note 62, at 565, 573-86.

70. Intriguingly, penalty default rules may entail less variance than majoritarian rules when the penalty rule is the existing default rule—as in the case of the *Hadley* damages rule. When the penalty rule is the status quo, majoritarian rules become more attractive, all else equal. When neither the majoritarian rule nor the penalty rule is the default rule, however, the majoritarian rule will generally entail less variance.

71. In Example 2, Policy *B* has a small probability of producing a favorable outcome. Penalty defaults probably have higher probabilities of success, but success is not quite as advantageous as in the example. If necessary, the example could be altered. The basic principle remains the same. So long as a penalty default rule provides some possibility of a

LEARNING THROUGH POLICY VARIATION

Trying some penalty default rules in contract law may be especially attractive because the success of one penalty default rule may be correlated with success in other areas. Consider a hypothetical rule forcing car insurers to choose between prominently stating their history of rate changes or being restricted from changing rates. This rule may yield valuable information about another hypothetical rule forcing credit card issuers to make the same choice.⁷² If the car insurance rule facilitates more informed and different consumer decisionmaking, then the credit card rule is also likely to have an impact. If the car insurance rule changes nothing, then it is less likely that the credit card rule is worth its costs. The potential to extrapolate the results of the car insurance rule increases the desirability of the rule because the car insurance rule provides information about the credit card rule that is valuable for finding the best policy with respect to credit cards in the future. Policymakers should therefore be more likely to adopt the car insurance information-forcing rule, all else equal.

The optimal search approach demonstrates that penalty default rules have an informational value that makes them more attractive than previously realized. The next Subsection demonstrates that a similar calculus applies to the debate about shareholder democracy.

3. *Increasing Shareholder Democracy in Corporate Law*

Corporate law scholars and practitioners intensely debate the appropriate degree of shareholder democracy and power in corporations.⁷³ Some argue that

better outcome while involving more risk than the majoritarian rule, then the optimal search approach dictates that the penalty default rule should be tried first in some cases.

72. Current law requires credit card issuers to prominently state their rates, but the issuers reserve the right to change the rates unilaterally. See Truth in Lending Act, 15 U.S.C. §§ 1601-1693r (2000).
73. Critics of enhanced shareholder democracy may argue that policy variance is unnecessary in this context because federalism in corporate law already produces sufficient policy variance, see, e.g., Roberta Romano, *Empowering Investors: A Market Approach to Securities Regulation*, 107 YALE L.J. 2359, 2392 (1998), with the “race to the top” leading all states to adopt value-enhancing policy variants. If one accepts this proposition, then high-variance policies are unhelpful because the variance they provide is unlikely to find innovative policies missed by the states. Others, however, argue that state competition for corporate charters produces relatively little variance. See, e.g., Lucian Arye Bebchuk & Assaf Hamdani, *Vigorous Race or Leisurely Walk: Reconsidering the Competition over Corporate Charters*, 112 YALE L.J. 553, 605-06 (2002). These arguments echo a broader debate about whether state competition for corporate charters produces a “race to the top” – with all states providing efficient and ever-improving laws – or a “race to the bottom” – with states undercutting each other by offering laws that favor management or other interested parties. Compare Roberta Romano, *Law As a*

shareholders should be entrusted with greater ability to dictate corporate decisions and governance mechanisms. Others argue that enhanced shareholder rights may undermine successful U.S. corporations by adding costs without corresponding benefits. To date, the critics of enhanced shareholder democracy have derailed significant expansions of shareholder democracy—there have been very few changes to the corporate decision-making process. The optimal search approach, by contrast, suggests that enhanced shareholder democracy is the higher variance approach, and may therefore be worthwhile even if its expected value is negative.

a. Increasing Shareholder Power

Most corporate governance arrangements derive from two sources—corporate charters and the state of incorporation’s corporate laws. To amend a corporation’s charter, state law requires board initiation and board submission to a shareholder vote of any proposed amendments. To alter the state of incorporation, a company commonly merges with a shell corporation incorporated in the desired state. To complete the merger, the board must initiate a shareholder vote.⁷⁴

As a result, shareholders have little say over the basic governance arrangements of a corporation. Any major changes in these arrangements require board approval. Critics contend that this is inefficient. If a governance arrangement is inefficient but suits the board of directors, virtually no means

Product: Some Pieces of the Incorporation Puzzle, 1 J.L. ECON. & ORG. 225, 280-81 (1985) (arguing that corporate law competition creates a race to the top), and Ralph K. Winter, Jr., *State Law, Shareholder Protection, and the Theory of the Corporation*, 6 J. LEGAL STUD. 251 (1977) (characterizing corporate law competition as a race to the top), with Lucian Arye Bebchuk, *Federalism and the Corporation: The Desirable Limits on State Competition in Corporate Law*, 105 HARV. L. REV. 1435, 1441 (1992) (describing how state competition in corporate law does not invariably lead to a race to the top), and William L. Cary, *Federalism and Corporate Law: Reflections upon Delaware*, 83 YALE L.J. 663 (1974) (describing state-based corporate law as race to the bottom). The optimal search approach to corporate law described in this Article may be valuable to this debate for two reasons. First, if Bebchuk and Hamdani’s lack-of-innovation argument is the better one, then high-variance optimal search policies at the federal level offer the potential for great benefits. Second, attempting new policies offers the possibility of resolving the seemingly never-ending debate about the value of state competition in corporate law. If innovative corporate law policies are attempted but consistently prove to be failures, then this provides strong evidence of a “race to the top.” For a thorough discussion of the relation between federalism and the optimal search approach, see *infra* Subsection II.C.4 and Section IV.C.

74. See Lucian Arye Bebchuk, *The Case for Increasing Shareholder Power*, 118 HARV. L. REV. 833, 843-46 (2005).

LEARNING THROUGH POLICY VARIATION

exist to change it because any change initiated by shareholders must be approved by the board.⁷⁵

Critics of the status quo want to expand shareholder power so that inefficient arrangements can be altered by shareholders without management approval. They advocate expanding the scope of bylaw amendments. Bylaw amendments may be unilaterally initiated by shareholders, but at present the ability of bylaws to alter governance arrangements is quite limited. If the scope of bylaw amendments were to expand, then shareholders would be able to alter governance arrangements more easily.⁷⁶

Critics of the status quo also advocate allowing shareholders to unilaterally initiate and approve charter amendments. Because charter amendments enable corporations to specify their own governance arrangements and alter state default arrangements, allowing unilateral shareholder charter amendments would similarly enable shareholders to alter governance arrangements.⁷⁷

Stephen Bainbridge and others dispute these recommendations,⁷⁸ arguing that director primacy has “stood the test of time”⁷⁹ and that

[a]ctive investor involvement in corporate decision-making seems likely to disrupt the very mechanism that makes the widely held public corporation practicable: namely, the centralization of essentially nonreviewable decision-making authority in the board of directors. The chief economic virtue of the public corporation is not that it permits the aggregation of large capital pools, as some have suggested, but rather that it provides a hierarchical decision-making structure well-suited to the problem of operating a large business enterprise

75. *See id.*

76. *See id.* at 844-46.

77. *See id.*

78. *See, e.g.,* Stephen M. Bainbridge, *Director Primacy and Shareholder Disempowerment*, 119 HARV. L. REV. 1735, 1758 (2006) [hereinafter Bainbridge, *Shareholder Disempowerment*]; Stephen M. Bainbridge, *Director Primacy: The Means and Ends of Corporate Governance*, 97 NW. U. L. REV. 547, 605-06 (2003); Iman Anabtawi, *Some Skepticism About Increasing Shareholder Power*, 53 UCLA L. REV. 561, 598 (2006); A. Gilchrist Sparks, III, *Corporate Democracy: What It Is, What It Isn't, and What It Should Be*, INSIGHTS, Mar. 2006, at 20, 23 (noting the “overwhelming historical success of the existing governance system”); *see also* Leo E. Strine, Jr., *Toward a True Corporate Republic: A Traditionalist Response to Bebchuk's Solution for Improving Corporate America*, 119 HARV. L. REV. 1759, 1759 (2006) (critiquing Bebchuk's proposal to increase shareholder power from the perspective of “traditionalist investors”).

79. Bainbridge, *Shareholder Disempowerment*, *supra* note 78, at 1758. Note the Burkean allusions in Bainbridge's argument.

with numerous employees, managers, shareholders, creditors, and other constituencies. In such an enterprise, someone must be in charge: “Under conditions of widely dispersed information and the need for speed in decisions, authoritative control at the tactical level is essential for success.”⁸⁰

Given these views, it is no surprise that Bainbridge and others support the current distribution of power between shareholders and managers.

b. Changing Shareholder Voting Procedures

Unlike the “shareholder power” debate, the shareholder franchise debate focuses narrowly on one method of corporate governance—the ability of shareholders to elect directors and approve mergers. Even if one believes that directors should have considerable power to make corporate decisions, one might still argue that shareholders should be free to choose their directors. As one Delaware Chancery Court opinion states, “The shareholder franchise is the ideological underpinning upon which the legitimacy of directorial power rests.”⁸¹

Scholars and courts differ regarding the effectiveness of the shareholder franchise.⁸² Delaware courts state axiomatically that shareholder voting enables shareholders to unseat directors.⁸³ A variety of evidence, however, suggests that the shareholder franchise does not work as well as the Delaware courts suggest. For example, the stock market responds favorably to close directorial proxy contests with dissident slate victors and unfavorably to close contests resulting in a management victory.⁸⁴ This suggests that shareholder voting systematically favors management to a degree not justified by stock market value maximization.⁸⁵ In addition, proxy contests are few and far between,

80. *Id.* at 1749 (quoting KENNETH J. ARROW, *THE LIMITS OF ORGANIZATION* 69 (1974)).

81. *Blasius Indus., Inc. v. Atlas Corp.*, 564 A.2d 651, 659 (Del. Ch. 1988).

82. Compare Lucian A. Bebchuk, *The Myth of the Shareholder Franchise*, 93 VA. L. REV. 675 (2007) (arguing that the shareholder franchise is illusory), with Martin Lipton & William Savitt, *The Many Myths of Lucian Bebchuk*, 93 VA. L. REV. 733, 734 (2007) (arguing “that the ‘myth of the shareholder franchise’ is no myth at all”).

83. See *Unocal Corp. v. Mesa Petroleum Co.*, 493 A.2d 946, 959 (Del. 1985) (“If the stockholders are displeased with the action of their elected representatives, the powers of corporate democracy are at their disposal to turn the board out.”).

84. See Yair Listokin, *Corporate Voting vs. Market Price Setting* (July 3, 2008) (unpublished manuscript, on file with author), available at <http://ssrn.com/abstract=1112671>.

85. See *id.*

LEARNING THROUGH POLICY VARIATION

potentially suggesting some handicaps to the ability of dissidents to wage proxy contests.⁸⁶ Others, however, believe that the shareholder franchise works well, noting that “under the existing rules, running an election contest through separate proxy materials is already a viable alternative and a viable threat. . . . [S]hareholders do run election contests on a regular basis under the existing rules.”⁸⁷

There are several reasons why the shareholder franchise may not function efficiently.⁸⁸ First, management enjoys discretion over the timing of a vote—within a certain range management can choose the day most likely to deliver a victory.⁸⁹ If a previously scheduled date looks like it will result in a loss, management enjoys the ability to change the date of the proxy vote. Second, management benefits from superior relationships with and contact information about shareholders.⁹⁰ Dissidents must sue merely to obtain a list of shareholders, while management often has longstanding relationships with shareholders that may make shareholders receptive to management mailings and arguments. Third, management bears less of the cost of proxy contests than dissidents.⁹¹ Management time and effort in resisting dissident proxy challenges is paid for by the corporation. In addition, management may expend unlimited funds from corporate coffers for soliciting proxies. Dissidents, by contrast, are only reimbursed for proxy expenses when they defeat management. Fourth, the prevalence of staggered boards means that a potential dissident must win two or more proxy contests over a long period to gain control over a corporation.⁹² The critics argue that these disadvantages combine to discourage shareholders from undertaking proxy contests and to

86. See Bebchuk, *supra* note 82, at 682-88 (arguing that the small number of contested elections indicates some problems with the shareholder franchise). *But see* Jonathan R. Macey, *Too Many Notes and Not Enough Votes: Lucian Bebchuk and Emperor Joseph II Kvetch About Contested Director Elections and Mozart's Seraglio*, 93 VA. L. REV. 759 (2007) (arguing that Bebchuk has no baseline for determining the appropriate number of contested elections).

87. Letter from Wachtell, Lipton, Rosen & Katz to Jonathan G. Katz, Sec'y, SEC (June 11, 2003), available at <http://www.sec.gov/rules/other/s71003/wachtello61103.htm>; see also Lipton & Savitt, *supra* note 82 (defending shareholder franchise).

88. See Bebchuk, *supra* note 82, at 688-94.

89. See generally Yair Listokin, *Management Always Wins the Close Ones*, 10 AM. L. & ECON. REV. 159 (2008) (discussing management advantages in corporate voting).

90. Bebchuk, *supra* note 82, at 688-94.

91. *Id.*

92. *Id.*

reduce the probability of winning for those shareholders that ultimately decide to pursue proxy contests.⁹³

In response to these perceived problems, many have suggested changes to the shareholder franchise. Bebchuk, for example, proposes that “reimbursement of expenses to challengers receiving a sufficiently significant number of votes (for example, one-third of the votes cast), and shareholder power to replace all directors” be enacted as corporate law default rules.⁹⁴ Bebchuk argues, “Furthermore, confidential voting and majority voting should be required in all elections.”⁹⁵ Many scholars and policymakers advocate shareholder access to the corporate ballot,⁹⁶ which would enable shareholders to propose alternative slates of directors on the proxy statements that a corporation provides to shareholders before its annual meetings.⁹⁷

As with increasing shareholder power, many dispute proposed reforms to the shareholder franchise.⁹⁸ Two critics of Bebchuk’s proposed reforms argue that

[i]gnoring decades of salutary historical development and the overwhelming lessons of observed boardroom behavior, Bebchuk advocates the abandonment of the traditional process for selecting and retaining directors of U.S. public corporations. In its stead, Bebchuk offers a novel electoral system of his own recent invention—a regime specifically designed to encourage costly proxy contests Bebchuk has utterly failed to carry the burden of justifying the radical reform he proposes. . . . Bebchuk has systematically failed to account for the likely and severe negative consequences [including waste and disruption of managerial efforts, enhanced power for special interests, excessive management short-termism, difficulty recruiting board personnel, and adverse impact on board processes] of his proposal.⁹⁹

93. See, e.g., *id.*

94. *Id.* at 677.

95. *Id.*

96. See, e.g., Lucian Arye Bebchuk, *The Case for Shareholder Access to the Ballot*, 59 *BUS. LAW.* 43 (2003).

97. At present, dissident slates of directors must mail their own proxy forms to shareholders. See Bebchuk, *supra* note 82, at 696-98.

98. See, e.g., Lipton & Savitt, *supra* note 82; see also Macey, *supra* note 86 (arguing that Bebchuk has no baseline for determining the appropriate number of contested elections).

99. Lipton & Savitt, *supra* note 82, at 733-34. Again, note the strong Burkean tone to the arguments in favor of the status quo.

LEARNING THROUGH POLICY VARIATION

This is no mere academic debate. For a number of years culminating in 2007, the SEC has considered shareholder ballot access regulations and received more than thirty-four thousand comment letters from members of the public about the proposed regulations—a record number of comments.¹⁰⁰ In total, the shareholder franchise debate resembles the shareholder power debate. On the one hand, critics of the status quo identify imperfections in the shareholder franchise and argue for various policy reforms to correct these faults. On the other hand, a group exists that believes that the reform proposals will add more costs than benefits and that the status quo works relatively well. That group demands that reformers show a compelling case for change before a change is made.

c. *Optimal Search and Reforms to Shareholder Power and the Shareholder Franchise*

The academic and professional debate about the optimal degree of shareholder power and shareholder voting demonstrates that the impact of proposed reforms to these institutions is uncertain. In other words, the effect of the reforms is subject to high variance. Some think the proposals will be a home run; others fear they will be a disaster.¹⁰¹

Suppose that the defenders of the status quo have the better argument; the proposed reforms are more likely to do harm than good. Does this mean that the proposed reforms should not be enacted? The optimal search approach says “not necessarily”—so long as the reforms are reversible. Because the proposed reforms have high variance, they have a high search value, corresponding to Policy *B* in the above example. If enhanced shareholder power and shareholder voting mitigates agency costs and ensures that corporations are run more efficiently, then these reforms can be continued, leading to a permanently better outcome than the status quo. If the proposed reforms prove expensive

100. See Annette L. Nazareth, Comm’r, SEC, Speech by SEC Commissioner: Remarks Before the International Corporate Governance Network (Oct. 29, 2007) (transcript available at <http://www.sec.gov/news/speech/2007/spch102907aln.htm>) [hereinafter Nazareth Speech].

101. See, e.g., Macey, *supra* note 86; John F. Olson, *Professor Bebchuk’s Brave New World: A Reply to “The Myth of the Shareholder Franchise,”* 93 VA. L. REV. 773 (2007); Lynn A. Stout, *The Mythical Benefits of Shareholder Control*, 93 VA. L. REV. 789 (2007); E. Norman Veasey, *The Stockholder Franchise Is Not a Myth: A Response to Professor Bebchuk*, 93 VA. L. REV. 811 (2007). It is extremely unlikely that the issue would have engendered so many comment letters if its proponents did not think it was likely to have significantly positive effects and its detractors did not think the reform would have negative implications.

and make corporate decisionmaking impossible, then the reforms will not be continued and the old status quo of director primacy and few proxy contests will be reenacted. The status quo, by contrast, has much lower variance because its effects on voting and corporate governance are relatively well understood. Even if the status quo is better on average, like Policy A, it has a lower aggregate value because it does not include the option to use the reforms should they prove effective.

The optimal search approach refutes the argument of critics that any reform must meet a heavy burden to justify itself. When policies can be changed in subsequent periods, this argument is incorrect. Reforms that have higher variance than the status quo are made desirable, rather than undesirable, when there is more uncertainty about their effects.

4. *Federalism and the Optimal Search Approach*

The previous examples assume that the policy universe consists of one jurisdiction and that policies can be changed each period. The examples also place no obvious value on settled expectations for law. Under these conditions, the optimal search approach offered a number of counterintuitive recommendations for policymaking, emphasizing the option value of policies rather than their expected value. But what if the policy universe contains multiple government units? The optimal search approach applies with greater force in such circumstances—policies with high variance and low expected value become increasingly desirable from a systemwide perspective as the number of jurisdictions increases.

Federalism is one example of a system of government with multiple jurisdictions. A system of government is federal if “(1) two levels of government rule the same land and people, (2) each level has at least one area of action in which it is autonomous, and (3) there is some guarantee . . . of the autonomy of each government in its own sphere.”¹⁰² The implications of the optimal search approach derived in this Part, however, do not apply exclusively to federalism. The analysis below applies whenever (1) a policy has correlated effects in two or more locations, (2) these locations are separately governed, (3) the effects of a policy are observable, and (4) policymakers care from a normative perspective about the effects of a policy in both locations.

¹⁰² WILLIAM H. RIKER, FEDERALISM: ORIGIN, OPERATION, SIGNIFICANCE 11 (1964).

LEARNING THROUGH POLICY VARIATION

a. “Experimentalism” and Federalism

Federalism has long had an association with experimentation.¹⁰³ Justice Brandeis’s oft-cited description of federalism reads, “It is one of the happy incidents of the federal system that a single courageous State may, if its citizens choose, serve as a laboratory; and try novel social and economic experiments without risk to the rest of the country.”¹⁰⁴

It is therefore not surprising that federalism and devolution are the lynchpins of experimentation in the “democratic experimentalist” framework discussed in Part I. Dorf and Sabel praise

a public sector model of problem solving adapted to a polity in which omnibus, national measures can rarely address the particularities of local experience, yet locales in isolation from one another are unable to explore and evaluate even the most immediately promising solutions to their problems. The model requires linked systems of local and inter-local or federal pooling of information, each applying in its sphere the principles of benchmarking, simultaneous engineering, and error correction, so that actors scrutinize their initial understandings of problems and feasible solutions. These principles enable the actors to learn from one another’s successes and failures while reducing the vulnerability created by the decentralized search for solutions.¹⁰⁵

The criteria for choosing experiments in an “experimentalist” federal framework are somewhat unclear. Experimentalists suggest that localities will naturally experiment because different populations will have different policy goals. In other words, different jurisdictions will pursue new policies because the new policies have higher expected value for that particular jurisdiction than do existing policies¹⁰⁶—that is, policies offering the “most immediately

103. Of course, there are many other arguments for federalism. For a comprehensive review of these arguments, see DAVID L. SHAPIRO, *FEDERALISM: A DIALOGUE* 58-106 (1995). This Article does not examine these powerful alternative justifications for federalism.

104. *New State Ice Co. v. Liebmann*, 285 U.S. 262, 311 (1932) (Brandeis, J., dissenting). Justice Brandeis’s argument for federalism has been adopted by many scholars examining a wide variety of different contexts. A Westlaw search for law review articles quoting Justice Brandeis’s federalism-as-laboratory quotation retrieved almost seven hundred articles.

105. Dorf & Sabel, *supra* note 23, at 287-88.

106. See generally *id.* at 314-24 (discussing how diversity can lead to local differences in policy and how policy should be “benchmarked” to produce new policies that are better than previous policies).

promising solutions to [the jurisdiction's] problems."¹⁰⁷ These are the "courageous" states—the ones that overcome inertia to attempt a new policy with positive expected value. Other jurisdictions can costlessly observe the outcomes of these high-variance/high-expected-value policies and adopt the policies if they are successful while avoiding their negative effects if they are failures.¹⁰⁸

According to the experimentalist, the primary obstacles to such experimentation are, at present, national government obstruction and local inability to plan and evaluate different experiments.¹⁰⁹ Once these obstacles are removed, federalism naturally produces significant amounts of experimentation. In the experimentalist context, federalism and devolution thereby facilitate the adoption of high-variance/high-expected-value policies, with benefits for all jurisdictions.

b. Optimal Policy Variation Under Federalism

The most salient distinction between the experimentalist and optimal search approaches to federalism is the degree of desirable experimentation. Experimentalists believe that federalism facilitates passage of policies that are risky but have positive expected value and encourages such "experiments." The optimal search approach, by contrast, suggests that even some risky policies that have *negative* expected value should be attempted, and that federalism makes high-variance/low-expected-value policies favored by the optimal search approach even more attractive.

When one jurisdiction tries a new policy with high variance, it opens the possibility of continuing with this policy in future periods. But the state trying the policy is not the only state with the option to continue with the successful policy. When policies have similar effects across states, other states may implement the policy that has been tried by one state. These "externalities" raise the return to policy variance, because the upside of the risky policy can be shared by all jurisdictions, while the low expected value of the policy is imposed on only one jurisdiction.

¹⁰⁷. *Id.* at 287.

¹⁰⁸. Susan Rose-Ackerman questions the argument that federalism promotes innovation. See Susan Rose-Ackerman, *Risk Taking and Reelection: Does Federalism Promote Innovation?*, 9 J. LEGAL STUD. 593 (1980). Public choice considerations, such as public policymakers' desire for reelection, may diminish incentives for risk taking by policymakers in a federal system. After a thorough analysis of these incentives, Rose-Ackerman concludes that "few useful experiments will be carried out in [state and local governments]." *Id.* at 594.

¹⁰⁹. See Dorf & Sabel, *supra* note 23.

LEARNING THROUGH POLICY VARIATION

Consider the choice between Policy *A* and Policy *B*, as demonstrated in Example 3. Policy *A* costs nothing to implement and yields social benefits of \$50 per period within the implementing jurisdiction. Policy *B* costs \$20 to implement and has a probability of 0.08 of yielding benefits to the implementing jurisdiction of \$120 per period and a probability of 0.92 of yielding benefits of \$0.

Example 3.

POLICY	A	B
Cost	\$0	\$20
Expected Benefit per Period	\$50	$\$9.60 = 0.8(\$120) + 0.92(\$0)$

In a uni-jurisdictional world, *A* is a better choice than *B*. The payoff for *A* is

$$\sum_{t=1}^{\infty} \frac{\$50}{(1+0.1)^t} = \$500$$

while the payoff from *B* is

$$-20 + 0.08 \left[\$120 + \sum_{t=2}^{\infty} \frac{\$120}{(1+0.1)^t} \right] + 0.92 \left[0 + \sum_{t=2}^{\infty} \frac{\$50}{(1+0.1)^t} \right] = \$495.06$$

In spite of the optimal search benefits of *B*, it is an inferior policy because its expected value is so low relative to *A* that the search benefits of *B* cannot compensate.

Now, however, suppose that there are two jurisdictions of equal size and importance, and that only the first jurisdiction can implement Policy *B*. In this case, Policy *B*, with an expected value considerably inferior to Policy *A*, becomes a socially optimal choice. If both jurisdictions choose Policy *A*, then the benefit to each jurisdiction is

$$\sum_{t=1}^{\infty} \frac{\$50}{(1+0.1)^t} = \$500$$

This makes the average benefit per jurisdiction equal to \$500. If one jurisdiction tries Policy *B* but the other jurisdiction can use Policy *B* if it succeeds, then the payoff to the jurisdiction that tries *B* is

$$-20 + 0.08 \left[\$120 + \sum_{t=2}^{\infty} \frac{\$120}{(1+0.1)^t} \right] + 0.92 \left[0 + \sum_{t=2}^{\infty} \frac{\$50}{(1+0.1)^t} \right] = \$495.06$$

The payoff to the jurisdiction that does not try *B*, but implements *B* if *B* proves successful is

$$0.08 \left[\$50 + \sum_{t=2}^{\infty} \frac{\$120}{(1+0.1)^t} \right] + 0.92 \left[\$50 + \sum_{t=2}^{\infty} \frac{\$50}{(1+0.1)^t} \right] = \$554.46$$

The average payoff per jurisdiction if one jurisdiction tries Policy *B* is therefore

$$\frac{\$495.06 + \$554.46}{2} = \$525.26$$

Thus, the average payoff when one jurisdiction tries Policy *B* is greater than the average payoff when both jurisdictions choose *A*. From a social perspective, *B* should be tried, in spite of its inferiority in both the uni-jurisdictional case and in expected-value terms. This occurs because federalism allows the low-expected-value costs of high-variance policies to be limited to one jurisdiction, while the upside benefits of searching can be shared across jurisdictions. Federalism thereby increases the variance-loving aspects of the optimal search perspective. The experimentalist framework recognizes the potential for learning from local policy variation but fails to emphasize the degree to which variance for even policies with negative expected value is optimal.

LEARNING THROUGH POLICY VARIATION

c. Preemption and Optimal Policy Variation

State and federal law often overlap. The Constitution and Supreme Court precedent both establish that state law is preempted in many cases.¹¹⁰ For example, state law is preempted “where state law stands as an obstacle to the accomplishment and execution of the full purposes and objectives of Congress.”¹¹¹

While many have critiqued this preemption standard,¹¹² the optimal search approach reorients and strengthens the critique. Most scholars who critique the preemption standard praise federalism for allowing states to choose policies that strive to maximize the welfare of their citizens.¹¹³ This critique, however, is incomplete. A policy may maximize the welfare of the citizens of a state but cause sufficient harm to citizens in other states that the state law should be preempted on efficiency grounds.¹¹⁴ Indeed, on efficiency grounds courts should strike down any state law that Congress preempts because Congress is able to internalize the externalities imposed by state law on citizens of other states.

The argument for preemption is significantly weakened by the optimal search approach. Because learning is possible across both time and space in a multi-jurisdictional world, policies that have extremely negative expected values (such as Policy *B* in Example 3 above) prove to be good choices from an optimal search perspective. To the extent that courts measure “the full purposes and objectives of Congress” for preemption purposes by the expected

110. For a thorough examination of preemption, see Caleb Nelson, *Preemption*, 86 VA. L. REV. 225 (2000).

111. *Boggs v. Boggs*, 520 U.S. 833, 844 (1997) (quoting *Gade v. Nat’l Solid Wastes Mgmt. Ass’n*, 505 U.S. 88, 98 (1921)).

112. See Nelson, *supra* note 110, at 229 n.16 (cataloguing and summarizing the primary critiques of preemption, including critiques found in KENNETH STARR ET AL., *THE LAW OF PREEMPTION* 47-56 (1991); Betsy J. Grey, *Make Congress Speak Clearly: Federal Preemption of State Tort Remedies*, 77 B.U. L. REV. 559, 561 (1997); S. Candice Hoke, *Preemption Pathologies and Civic Republican Values*, 71 B.U. L. REV. 685, 687-88 (1991); Donald P. Rothschild, *A Proposed “Tonic” with Florida Lime To Celebrate Our New Federalism: How To Deal with the “Headache” of Preemption*, 38 U. MIAMI L. REV. 829, 830 n.3 (1984); Paul Wolfson, *Preemption and Federalism: The Missing Link*, 16 HASTINGS CONST. L.Q. 69, 114 (1988)).

113. See *supra* note 112.

114. Electric power production provides an excellent example. Coal-fired plants are considerably cheaper than gas-fired plants for producing electricity. Burning coal, however, produces more externalities than does gas. A national regulator might decide that these externalities outweigh the benefit of cheaper power production. If the acid rain harms other states but does not harm the state where the power is produced, however, then a state regulator may decide to approve coal power production.

effect of Congress's national policy, they will ignore the informational benefits of federalism and thereby cause overpreemption.¹¹⁵ Moreover, state laws provide the greatest information when they compare in substance to the national policies, because such policies are comprehensive enough to apply to all locations, rather than being narrowly tailored to local idiosyncrasies. Current law, however, places such state policies in the greatest danger of preemption. Thus, by suppressing the optimal search value of policy variation, preemption is likely to be overapplied considerably. Indeed, as shown in Part IV, states pursuing the interests of their own citizens actually will provide insufficient policy variation under many circumstances.¹¹⁶ Grants to states, therefore, should replace preemption in many cases of conflict between state and national laws.¹¹⁷

At this point, the benefits of choosing public policies with high variance due to the optimal search perspective have been demonstrated both theoretically and practically in several different contexts. The remainder of the Article addresses a number of theoretical and practical objections to the optimal search approach and develops some recommendations for optimally garnering the "learning" benefits of policy variance.

III. THE OPTIMAL SEARCH APPROACH: OBJECTIONS, RESPONSES, AND MODIFICATIONS

The optimal search approach places positive value on policies with high variance in outcomes. New policies tend to have high variance because they have not been tested. The range of possible outcomes from a new policy is wider than the range in outcomes from a well-tested policy. Therefore, the optimal search approach places greater value on policy innovation than do other approaches.

Compared to risk-averse rationalists, risk-neutral rationalists do not dislike variance *per se*.¹¹⁸ Differences between the optimal search approach and the traditional rationalist approach stem from the value attached to variance in the optimal search approach: the optimal search approach may choose policies

115. *Hines v. Davidowitz*, 312 U.S. 52, 67 (1941).

116. This conclusion arises from the optimal search approach rather than the experimentalist approach to federalism. See *infra* Section IV.C.

117. This does not mean that the benefits of standardization should be ignored. Instead, they should be compared with the underappreciated informational benefits of variation.

118. Again, the argument will be presented in risk-neutral terms for the sake of clarity. Parallel arguments could be made in utility terms for a risk-averse utilitarian rationalist.

LEARNING THROUGH POLICY VARIATION

with low expected value when the policies have high variance, while the traditional rationalist seeks to maximize expected value. That said, the rationalist and optimal search approaches are not contradictory. The optimal search approach merely constitutes a modification of the rationalist approach that accounts for policy learning. In a world with reversible policies and learning, rationalists should favor the optimal search approach, which has a greater long-run expected benefit than a traditional rationalist approach given the possibility of learning and changeable policies.

A. Burkean Objections

The optimal search approach's emphasis on variance brings it into direct conflict with Burkean political theory, which stigmatizes "innovation."¹¹⁹ Indeed, Burkean political theory cautions against innovation even when innovation aims to produce the policies with the highest expected values. Certainly, then, the Burkean critique of innovation applies with even greater force to the optimal search approach, which actively seeks policies with high variance, even when they do not produce the best expected outcome. I will argue, however, that while there are undoubtedly points of disagreement between the optimal search approach and Burkean political theory, these disagreements are primarily confined to the question of policy reversibility. On other issues, the optimal search approach and Burkean theory share an empirical inclination that makes them far from inconsistent.

Burkeans emphasize the limitations of human reason.¹²⁰ They argue that humans are unable to predict the outcomes of policies. Instead of prediction, Burkeans stress "empiricism."¹²¹ One strand of empiricism is respect for past practice. Past practice has been proven empirically to work in previous periods. Therefore, it has an empirical validity absent in any new policy. A second strand comes in the formulation of new policies. Burkeans favor incremental change that is regularly evaluated empirically. Only when the changes pass empirical muster should the reform process continue.

The optimal search approach conforms to these Burkean ideas. The optimal search approach derives its effectiveness from its empiricism. Policies are

119. While Burke was opposed to "innovation" in the form of abrupt change in policy, he was more receptive to extremely gradual policy modifications. See Ian Harris, *Edmund Burke*, in *STANFORD ENCYCLOPEDIA OF PHILOSOPHY* §8 (Edward N. Zalta ed., 2004), <http://plato.stanford.edu/entries/burke/> (last visited Oct. 27, 2008).

120. See *supra* Part I.

121. Young, *supra* note 7, at 646.

tested, and they are continued only if they are effective. Rather than asserting that one policy is better than another, the optimal search approach encourages empirical testing of policies until a successful policy is found. While there is inevitably some rational judgment attached to choice of policies in the optimal search approach—one must be able to guess at the policies' variance and expected value¹²²—one can be skeptical of human reason and support the optimal search approach. Indeed, so long as the overlooked limitations of human reason do not inject bias into the process but only greater variance,¹²³ the limitations may increase the argument for change in the optimal search approach where greater variance is a positive rather than a negative trait.

The optimal search approach also partly satisfies another Burkean virtue that is often absent from traditional rationalist discourse. Anthony Kronman stresses Burke's vision of trusteeship—that present-day individuals have obligations to both future and past generations.¹²⁴ As Burke stated,

[O]ne of the first and most leading principles on which the commonwealth and the laws are consecrated, is lest the temporary possessors and life-renters in it, unmindful of what they have received from their ancestors, or of what is due to their posterity, should act as if they were the entire masters; that they should not think it among their rights to cut off the entail, or commit waste on the inheritance, by destroying at their pleasure the whole original fabric of their society
...¹²⁵

The optimal search approach recommends considerable innovation. I will not claim that this represents deference to the choices of past generations; the policies of past generations will frequently be changed. To the extent that deference to past generations is tantamount to the maintenance of past policies,

122. There is some rational judgment required for the Burkean approach as well. Unless Burkeans want to have rigid laws, which Burke (at least) did not, there must be some rational evaluation of new policies and debate about the scope by which they change existing policy.

123. That is, suppose that the range of possible outcomes is much wider than an expert thinks, but the increase in variability occurs on both the positive side and the negative side. If the negative consequences of a policy change increase but the positive consequences do not, then the expected value of a policy change goes down and the policy change becomes less attractive under the optimal search approach.

124. See Kronman, *supra* note 16, at 1066-68.

125. EDMUND BURKE, REFLECTIONS ON THE REVOLUTION IN FRANCE 81 (Frank M. Turner ed., 2003).

LEARNING THROUGH POLICY VARIATION

the optimal search approach is unfaithful to the past. This does not mean that the optimal search approach rejects the policies of past generations, however. In the optimal search approach, it is critical for policymakers to revert to previous policies if innovative policies fail. Policies are not changed because they are bad; they are changed because of the benefit of finding something better.

While the optimal search approach places the present in a rejectionist relationship with the past, it also places the present generation in a more harmonious relationship with future generations. The optimal search approach dictates that on some occasions a high-variance/low-expected-value policy is desirable even though it reduces present well-being, because of the value of the information to future periods. In other words, the present generation is sacrificing some benefit today in exchange for future benefits. The future generation enjoys all the benefits of the optimal search approach without any of the costs, enabling the present generation to confer a considerable benefit on future generations.¹²⁶

It is also important to note that the optimal search approach does not command innovation. If a successful policy has been found, then the optimal search approach does not contradict the Burkean suggestion. Both approaches will retain the current policy—the optimal search approach retains the policy because the search has ended, and the Burkean approach retains the policy because of the risks of innovation. It is true, however, that by emphasizing the informational benefits of innovation, the optimal search approach creates additional situations where innovation is worthwhile. In addition, the two approaches sharply differ in their recommendations when both agree that change is desirable. The optimal search approach favors high-variance policies; the Burkean approach favors incremental changes.

Given these similarities between the Burkean approach and the optimal search approach, why are the policy recommendations of the two so different? The simple answer concerns irreversibility. The optimal search approach assumes that the effects of policies can be observed and policies can change to reflect policy learning. When these assumptions fail, the optimal search approach no longer recommends innovation.¹²⁷ Indeed, when policies are

^{126.} This does not mean that the present generation is behaving altruistically. The optimal search approach benefits all members of the present generation who will enjoy the potential benefits of policy innovation in future periods. Alternatively, the present generation could consume some of the surplus created by the optimal search approach for future generations by leaving less for future generations along other dimensions. Because the optimal search approach creates value, such gain sharing could leave all generations better off.

^{127.} This will be discussed in detail in Section III.C.

irreversible or extremely sticky, the optimal search approach recommends conservative policymaking that recognizes the “real option” value of these policies.¹²⁸

Burkeans, however, must reconcile their faith in empiricism with claims of irreversibility. If policies cannot be changed, then the Burkean approach—slow changes only when the existing policies cease to be effective—cannot be carried out. So Burkeans at least recognize the potential for policy learning and change that is the linchpin of the optimal search approach. To favor cautious policy change, Burkeans must argue that the reversibility of policies depends on their variance—that drastic changes in policy are irreversible, while incremental changes in policy are easy to reverse. Variance may be related to reversibility, as discussed in the next Section. But high variance does not imply irreversibility. Some policies may be high variance but easily reversible, while other policies may involve little variance but prove hard to reverse. Burkean thought does not emphasize the harms of variance because of irreversibility. Instead, it emphasizes the potential for bad outcomes, rather than irreversible outcomes, from changing policies.

B. *The Costs of Changing Policies*

Part II demonstrated that when policies can be changed without cost, high-variance policies become increasingly attractive. But what if policy changes are costly? This Part examines the consequences of policy irreversibility and inertia for optimal policymaking.

First, irreversibility must be defined. Some aspects of all policies—such as the costs of learning about them or implementing them—are always irreversible, but this does not mean that the policy as a whole is irreversible. Instead, a policy is irreversible if the policy itself or a significant portion of its effects can never be undone.¹²⁹

There are several sources of irreversibility. This Article focuses on two of those sources. The first source entails costs involved in implementing a new

128. For a detailed description of the “real options” approach, see *infra* Subsection III.C.1.

129. See Sunstein, *supra* note 6, at 860-64 (distinguishing between irreversibility as “seriousness” and irreversibility as “sunk costs”). The definition provided here resembles the sunk cost definition. A cost is sunk if it cannot be retrieved. A policy with long-lasting effects can therefore be represented as having a high initial cost—measured as the present discounted value of the future effects—or it can be represented as having costs distributed over a long time period. In the numerical examples, I choose to equate irreversibility with long-lasting effects, but the two methods are mathematically equivalent.

LEARNING THROUGH POLICY VARIATION

policy. The second source involves policy effects that linger, even if the formal policy is changed in a subsequent period.¹³⁰

A new policy must be formulated and implemented. Some of the costs of policy formulation and implementation are functions of institutional design and will be discussed below, but other costs remain in any institutional setting. Policymaking requires considerable amounts of policymakers' time and effort—an important and irretrievable cost. Policymakers can only enact so many new policies in any given period.¹³¹ In addition, individuals must learn about and evaluate the new policy, adding additional costs. Furthermore, many individuals exhibit “status quo bias,” creating yet another barrier to policy change.¹³²

Many of these costs can never be retrieved. These are sunk costs and therefore irreversible.¹³³ For example, the costs of learning about a new law or overcoming cognitive discomfort with change are not reimbursed or eliminated if the law is subsequently changed. Irreversible sunk costs do not prevent policy changes in the future. From the future's perspective, the costs are sunk and therefore irrelevant. From an ex ante perspective, however, these costs represent a disincentive to change laws.

Some reforms require considerable irreversible costs for implementation in addition to drafting and learning costs. Consider a strong defense policy. Defense infrastructure and weaponry is expensive, and the costs of building a defense system cannot be retrieved if a future policymaker decides that the policy is ineffective and that the system should be destroyed or sold. Similarly, a regulation that requires installation of new manufacturing equipment that provides no value other than compliance with the regulation entails significant irreversible costs.

Inertia-causing sunk costs characterize all new policies. Indeed, if inertia is total, then a model of optimal policy changes is moot because there will be no

130. Almost all sources of irreversibility can be placed within these two broad categories. For example, inertia can be described as resulting from high formulation and implementation costs.

131. There are many policies that may have positive expected value, but cannot be implemented because of the scarcity of legislative time and effort. These time and effort expenditures are irretrievable opportunity costs of formulating a particular policy.

132. Daniel Kahneman, Jack L. Knetsch & Richard H. Thaler, *Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias*, J. ECON. PERSP., Winter 1991, at 193, 197-99.

133. See DIXIT & PINDYCK, *supra* note 6, at 8-9. The effects of a policy in the current period are obviously irreversible. If a policy is a failure, these effects cannot be recouped by changing the policy in the future period. These costs are accounted for in both the optimal search approach and the real options approach, which have countervailing benefits that surpass these per-period costs in many cases.

policy changes. So inertia must be partial rather than total. The following analysis of the effect of irreversibility on optimal policy assumes that inertia makes policies more long lasting, but not infinite. The greater the sunk costs described here, the stronger the power of policy inertia.

Some policies will have additional and ongoing irreversible costs. If a policy causes the extinction of a species, for example, then the effects of the policy linger even if the policy is subsequently changed. Similarly, greenhouse gas abatement policies have long, intergenerational time horizons given that some of these gases linger in the atmosphere for centuries.¹³⁴ These policies are more irreversible than the typical policy. Although, the typical policy demands irreversible upfront costs, the effects of the policy last only as long as the policy is enacted. “Irreversible effect” policies, by contrast, entail not only irretrievable upfront costs but also irreversible effects in future periods.

C. Irreversibility, Real Options, and the Optimal Search Approach

1. The Real Options Approach to Policymaking

Suppose that Policy *B* in Example 2 of Part II (the high-variance/low-expected-value policy) were irreversible. That is, if *B* were chosen initially, the policymaker could not then change to Policy *A* in a future period. This situation could arise if Policy *B* had irreversible effects or if changing Policy *B* once it had been enacted would entail infinite irretrievable switching costs—for example, if people would become so attached to Policy *B* that they could not bear the cognitive costs of switching. Under these circumstances, Policy *A* (the policy with higher expected value) would be preferred on rationalist, Burkean, and optimal search grounds. *A* has higher average value per period and therefore *A* would be preferable to *B* when change is impossible; irreversibility eliminates the value of a search for a highly effective policy through the choice of Policy *B*. Note that this conclusion depends upon an asymmetry of irreversibility: *A* can be changed—otherwise Policy *B* would not be under consideration—but *B* is unchangeable.

Indeed, when policies are irreversible, a counterpoint to the optimal search approach—the real options approach—applies.¹³⁵ The real options approach demonstrates that when new policies are irreversible and uncertainty about the

¹³⁴. See, e.g., NICHOLAS STERN, THE ECONOMICS OF CLIMATE CHANGE: THE STERN REVIEW 28 (2007) (“The impacts of climate change are persistent and develop over time. Once in the atmosphere, some [greenhouse gases] stay there for hundreds of years.”).

¹³⁵. See Sunstein, *supra* note 6, at 860–69.

LEARNING THROUGH POLICY VARIATION

policies' value can be resolved without implementing the policies, delayed implementation of policies with positive expected value is often more efficient than immediate implementation. By implementing an irreversible but attractive policy today, policymakers lose the ability to learn more about the policy in this period and to use that knowledge to determine whether to implement the policy in the next period. The loss of this "real option" to learn more and decide in the future may be more expensive than the cost of losing the use of an effective policy for another period. Therefore, delayed implementation of seemingly attractive policies may be efficient when policies are irreversible.

As an illustration, see Example 4 as follows:

Example 4.

POLICY	A	B
Cost	\$0	\$20
Expected Benefit per Period	\$50	$\$72 = 0.6(\$120) + 0.4(\$0)$

The single period utilitarian calculus now prefers *B* to *A*, since *B* has higher expected value. Suppose, however, that once Policy *B* is chosen, *A* can no longer be chosen (Policy *B* is irreversible, but not Policy *A*) and that the true value of *B* will be revealed in the next period regardless of whether Policy *B* is implemented. Under these conditions, the value of choosing *B* in the current period is

$$-20 + 0.6 \left[\sum_{t=1}^{\infty} \frac{\$120}{(1+0.1)^t} \right] + 0.4 \left[\sum_{t=1}^{\infty} \frac{\$0}{(1+0.1)^t} \right] = \$700$$

The value of waiting to learn *B*'s true payoff, implementing *B* if *B* proves to be effective, and retaining *A* if *B* is a failure is

$$50 + 0.6 \left[-20 + \sum_{t=2}^{\infty} \frac{\$120}{(1+0.1)^t} \right] + 0.4 \left[\sum_{t=2}^{\infty} \frac{\$50}{(1+0.1)^t} \right] = \$874.36$$

which is greater than the value of choosing *B* immediately. By waiting to implement *B* until the true effects of *B* are known, the policymaker gives up *B*'s higher expected value in one period, but gains the benefit of avoiding *B* should *B* prove to be a failure. In this example, the benefits of avoiding *B* are greater than the benefits of enjoying *B* for an extra period, so the policymaker should delay implementation of *B* until the next period.

The real options approach demonstrated in this example is the complement of the optimal search approach. When policies are reversible, searching for excellent policies is cheap, so innovative policies with high variance become preferred to some policies with greater expected value but lower upside. When new policies are irreversible, old policies are reversible, and learning about new policies is possible without implementing them, policymakers should delay implementation to learn more about proposed new policies.

2. *Burkeanism, the Precautionary Principle, and the Real Options Approach*

Whatever the real options approach says about entrenchment, one way to interpret the Burkean critique and the precautionary principle in risk regulation is to say that they refer to irreversible risks.¹³⁶ As the example shows, when policies are irreversible, Burkean conservatism is preferred to a choice of the policy with the highest expected value from a rationalist perspective. Note that some environmental policies, such as carbon dioxide emissions, may have nearly irreversible effects, implying that a precautionary "real options" approach applies to them.

The distinction between irreversible and catastrophic should also be emphasized. Some harms, such as the loss of a species, may be irreversible, but many would argue that such harms are not catastrophic. By contrast, other harms, such as large, one-time financial costs, can be characterized as catastrophic but are not necessarily irreversible. For example, once a catastrophic economic harm is experienced, an economy may return to its

¹³⁶ Sunstein makes exactly this argument in his article, *Irreversible and Catastrophic*. See Sunstein, *supra* note 6, at 855-64. His article assumes that cost-benefit analysis is superior to the precautionary principle when risks are reversible and does not consider the possibility that reversibility may imply a deviation from expected-value maximization that is similar in logic but very different in outcome from the irreversible harm precautionary principle.

LEARNING THROUGH POLICY VARIATION

previous steady state.¹³⁷ The real options approach and Example 4 above justify the precautionary principle in the face of irreversibility, but not in the face of catastrophic harm. If catastrophes are painful but sufficiently short-lived, then the logic of the optimal search approach described in Part II still applies. Catastrophic losses are not irrelevant to the optimal search approach. They lower the expected value of a policy, and this in turn lowers the desirability of a policy in the optimal search approach. The possibility of reversible catastrophic losses, however, does not undermine the learning value of experimental policies in the optimal search approach.

Thus, risk regulation should hinge on the question of reversibility as well as the expected impacts of a regulation. Subsection II.C.1 demonstrated that reversible regulations should be chosen according to the variance-loving optimal search approach. This conclusion often applies even when there are high sunk costs associated with implementing a new policy—the new regulation in that Subsection cost more than the old one. This Subsection, by contrast, explains that asymmetric irreversible policies should be chosen according to the real options approach. Most risk regulations, and indeed most policies, are neither completely reversible nor asymmetrically irreversible. The next Subsection examines the implications of partial reversibility for optimal policymaking.

3. *Sticky but Reversible Policies*

When policies are reversible, the optimal search approach indicates that high-variance policies become attractive relative to low-variance policies with lower expected value. When policies are asymmetrically irreversible, the real options approach indicates that cautious implementation becomes attractive relative to expected-value maximization. Policies, of course, are neither fully reversible nor irreversible. Interestingly, partial reversibility makes naïve per-period expected-value maximization more plausible than the previous analysis has indicated.

Suppose that all policies are sticky but not irreversible. For example, suppose that policies can be reversed after four periods.¹³⁸ Now consider the

¹³⁷. This discussion is not meant to minimize the harms caused by catastrophic occurrences, but to develop an analytical distinction between catastrophic and irreversible.

¹³⁸. The choice of four periods is arbitrary, but this choice demonstrates how partial reversibility brings optimal policy choices closer to naïve expected-value maximization. Partial reversibility could also be denoted by imposing a significant switching cost on any change in policy. The two methods can be made mathematically equivalent, as the longevity of a policy that should be changed imposes a cost that is directly analogous to a high switching cost.

examples used to illustrate the optimal search (Example 2) and real options (Example 4) approaches. First consider the optimal search example (Example 2), which had the following parameters:

POLICY	A	B
Cost	\$0	\$20
Upside Benefit	\$50	\$120
Downside Benefit	\$50	\$0
Upside Probability	1	0.2
Expected Benefit per Period	\$50	$\$24 = 0.2(\$120) + 0.8(\$0)$

Policy *A* surpasses Policy *B* from an expected-value perspective. The discussion of optimal search in Part II demonstrated that when policies are fully reversible, Policy *B* should be tried before *A*. When policies must be maintained for four periods, however, Policy *A* becomes a better choice than Policy *B*. Thus, the introduction of some policy stickiness brings optimal choices closer to the naïve expected-value maximization approach rather than the variance-loving optimal search approach.

Now consider Example 4, which was used to illustrate the real options approach. As demonstrated in Subsection III.C.1, although *B* trumps *A* in expected-value terms, the real options approach suggests that *A* should be chosen before *B* when *B* is irreversible. Now, however, modify Example 4 to suppose that policies are sticky but not irreversible—all policies (*A* and *B*) can be reversed after four periods but no sooner.¹³⁹ The payoff for choosing *A* in the current period¹⁴⁰ and retaining it for at least four periods is

¹³⁹. As in Example 4, the true impact of Policy *B* can be learned without actually implementing Policy *B*.

¹⁴⁰. Policy *A*, like Policy *B*, can only be changed after four periods in this example. Policy *A* has no variance in outcomes, so the delay is not caused by the time needed to observe the effects of Policy *A*. The delay could be caused by a delay in the ability to evaluate Policy *B* when Policy *B* is not being implemented. (Recall that the only reason to change Policy *A* under the real options approach is after Policy *B* has been proven effective.)

LEARNING THROUGH POLICY VARIATION

$$\sum_{t=1}^4 \$50 + 0.6 \left[-20 + \sum_{t=5}^{\infty} \frac{\$120}{(1+0.1)^t} \right] + 0.4 \left[\sum_{t=5}^{\infty} \frac{\$50}{(1+0.1)^t} \right] = \$774.87$$

while the payoff for choosing Policy *B* in the current period and retaining it for at least four periods is

$$-20 + 0.6 \left[\sum_{t=1}^{\infty} \frac{\$120}{(1+0.1)^t} \right] + 0.4 \left[\sum_{t=5}^{\infty} \frac{\$50}{(1+0.1)^t} \right] = \$836.60$$

When policies are symmetrically and partially sticky, the optimal policy is *B* in the current period. Policy *B* is the choice with the highest expected value, so introducing the possibility of reversibility—even sticky reversibility—moves the optimal policy from the cautious real options approach toward naïve expected-value maximization. Thus, when policies are partially sticky, the low per-period expected-value choices prescribed by the optimal search approach and the real options approach become less attractive. Instead, choices that maximize per-period expected value become increasingly desirable.

The optimal policy choices under a range of reversibility conditions have now been characterized. Under the real options approach, the more irreversible a new policy is, the more policymakers should favor caution as opposed to expected-value maximization. At some intermediate level of reversibility, policymakers should jettison the real options approach and simply choose the policy that has the greatest expected value, regardless of the policy's variance. As reversibility becomes relatively cheap, policymakers should favor innovation and variance relative to expected-value maximization under the optimal search approach.

The degree of policy reversibility is therefore a critical determinant of the appropriate policy choice. The next Section briefly identifies high reversibility settings conducive to high-variance policies under the optimal search approach.

D. Optimal Variance in Different Policymaking Contexts

Some policymaking settings have much greater costs associated with implementing new policies than do others. As a general matter, the lower the costs associated with policy change, the greater the benefit of variance in policy selection. Before analyzing different policymaking contexts, several points are

worth noting. The discussion assumes that the present offers the opportunity for policy change—perhaps because a policy has become highly salient—but that future opportunities for policy change may be limited. The discussion, moreover, assumes that the mechanisms of policymaking in each context are fixed—that is, that statutory lawmaking procedures cannot be altered.¹⁴¹ Finally, this Section does not aim to provide a thorough analysis of different policymaking contexts. Instead, it aims to propose how the costs and benefits of policy variance should be informed by the institutional policymaking environment.

Constitutional policymaking in the United States is characterized by extreme inertia. The “transaction costs” of changing constitutional policy are extremely high. All amendments to the Constitution must be proposed by supermajorities of both houses of Congress or by a supermajority of the states’ legislatures, followed by ratification by three-fourths of the states.¹⁴² Therefore, a constitutional policy is likely to be extremely sticky. If policymakers choose a high-variance constitutional policy for the information it provides and the policy proves to be a failure, the policy might be impossible to change, making the information of little value. Even if a majority of decisionmakers views the policy as a failure, the onerous supermajority requirement makes changing the failed policy extremely difficult and costly so long as some minority of decisionmakers favors the policy.¹⁴³ As a result, constitutional policymaking should maximize expected value per period rather than other measures. If one type of change is possible in the future but not another—for instance, a constitutional amendment can be implemented but not reversed—then constitutional policymaking should be made according to the real options approach. Although the enactment of a new constitutional policy appears more efficient on average, policymakers should delay implementing irreversible changes in order to first obtain more information about these changes.¹⁴⁴

Statutory policymaking requires multiple levels of approval. A statute must be approved by the House and Senate and then signed by the President or

141. This assumption is relaxed in later Sections.

142. U.S. CONST. art. V.

143. See Matthew C. Stephenson, *The Price of Public Action: Constitutional Doctrine and the Judicial Manipulation of Legislative Enactment Costs*, 118 YALE L.J. 2, 15-16 (2008).

144. It should be acknowledged, of course, that even constitutional policymaking is not impossible to reverse. The Eighteenth Amendment implementing Prohibition can be seen as an innovative constitutional policy—there are compelling theoretical arguments in favor of Prohibition—that was judged a failure and then repealed. See U.S. CONST. amends. XVIII, XXI. The long lag between Prohibition and repeal, however, demonstrates the inertia of constitutional policymaking.

LEARNING THROUGH POLICY VARIATION

supported by enough legislators to overcome a veto.¹⁴⁵ Each stage involves considerable costs, including the costs of placing an item on the agenda, learning about the relevant issues, and reconciling competing visions about the optimal policy. Given these costs, statutory policies exhibit considerable inertia. The costs of statutory policymaking, however, are considerably smaller than those incurred in constitutional policymaking—statutes generally require only ordinary majorities rather than supermajorities for approval.

Statutory policymaking therefore does not constitute an ideal setting for high-variance policies via the optimal search approach. If a high-variance/low-expected-value policy is enacted and proves a failure, the costs of policy change will discourage the implementation of policy change. The informational value of a high-variance policy is reduced in this context, since it will be difficult to act upon the knowledge provided by a new policy. It should be emphasized, however, that statutes, unlike constitutional amendments, are enacted frequently. Therefore, information produced by policy experimentation has positive value and should not be ignored when deciding upon statutory policies.

Judicial policymaking, particularly at the appellate level, entails lower costs than constitutional or statutory policymaking.¹⁴⁶ Courts generally are much smaller than legislatures, reducing the absolute costs of judicial agenda-setting and learning. Judicial policymaking also requires fewer approvals to become official policy. The policy embodied in a constitutional or statutory interpretation by the Supreme Court, for example, is instantly enacted. The interpretation can be overruled by statute or constitutional amendment, but such reversals require considerable time and effort, as discussed above. A judicial policy can be reversed cheaply, by contrast, through a subsequent Supreme Court decision.

Stare decisis raises the cost of judicial policymaking above these direct costs by encouraging or requiring judges to follow precedents.¹⁴⁷ If judges internalize stare decisis and violation of these doctrines imposes costs on judges, then high-variance policy innovations become less attractive. These innovations will be costly to introduce because they violate stare decisis. Moreover, once introduced, the policies will be costly to reverse because they will become the new frame of reference for stare decisis.

¹⁴⁵. U.S. CONST. art. I, § 7, cl. 2.

¹⁴⁶. Judges are policymakers in the sense that they create rules affecting future behavior. See *Republican Party of Minn. v. White*, 536 U.S. 765, 784 (2002) (noting the power of state judges to “make” common law).

¹⁴⁷. See *Planned Parenthood of Se. Pa. v. Casey*, 505 U.S. 833, 854-69 (1992).

In total, the formal constitutional status of judicial policymaking appears to provide an ideal setting for policy experimentation via the optimal search approach. Formally, judicial policies can be enacted and changed quickly relative to constitutional and statutory policies, raising the informational value of innovative policies. Experimentation should be reduced in settings where *stare decisis* is accorded more weight. Adherence to these doctrines raises the costs of policy change, and thereby raises the expected costs of policy innovation with low expected value but high upside. Subsection III.E.3 examines the costs and benefits of *stare decisis* and judicial minimalism while considering the optimal search approach.

Administrative policymaking, via rulemaking, entails intermediate policymaking costs. On the one hand, administrative agencies have institutional coherence lacking in legislative bodies—employees of administrative agencies are “employees” rather than members of a legislative body. The employees’ superior—the head of the agency or the President, can impose policymaking direction, reducing the costs of enacting new policies for agencies as compared to legislative bodies.¹⁴⁸

On the other hand, the Administrative Procedure Act (APA) establishes several procedural requirements for administrative policymaking.¹⁴⁹ Administrative “rulemaking” constitutes the primary mechanism whereby administrators “prescribe policy.”¹⁵⁰ While the costs of these requirements are real, they are not particularly onerous relative to the requirements for constitutional or statutory policymaking. Notice-and-comment rulemaking, for example, requires public notice of intent to formulate a policy and the opportunity for the public to respond to the intended rule, followed by agency consideration of the comments and possible (but not obligatory) revision of the policy.¹⁵¹ While the costs of noticing a policy change to the public and receiving and responding to comments are real, they are not prohibitive. Indeed, efficient policymaking may well benefit from public feedback, so that

148. This is not to say, however, that administrative agencies are inclined to change policies. Indeed, they may have incentives not to make frequent policy changes. These incentives will be examined in Part IV. See also JERRY L. MASHAW & DAVID L. HAREFT, *THE STRUGGLE FOR AUTO SAFETY* (1990) (discussing inertia in administrative agencies).

149. Administrative Procedure Act of 1946, 5 U.S.C. §§ 551-559 (2000).

150. Section 551 of the APA defines a rule as “the whole or a part of an agency statement of general or particular applicability and future effect designed to implement, interpret, or prescribe law or policy.” *Id.* § 551(4).

151. See Anne Joseph O’Connell, *Political Cycles of Rulemaking: An Empirical Portrait of the Modern Administrative State*, 94 VA. L. REV. 889, 901-02 (2008) (providing a comprehensive description of notice-and-comment rulemaking).

LEARNING THROUGH POLICY VARIATION

the costs of notice-and-comment are outweighed by the benefits. Neither formal rulemaking nor hybrid rulemaking appears to present overwhelming costs. While the recordkeeping requirements of formal rulemaking undoubtedly add to costs, such costs again appear small relative to the potential impacts of policy. Administrative policymaking appears to have greater reversibility than constitutional or statutory policymaking. Accordingly, administrative policymakers should pay greater attention to the optimal search benefits of innovative high-variance policies. If such policies fail, the administrative process should not provide an insuperable obstacle to changing them. If they succeed, then policymakers will have made a potentially important policy advancement.

In total, the optimal degree of policy variance depends critically on the institutional setting in which policymaking occurs. Learning through policy variation only becomes important when policymakers can apply the lessons learned by observing the impacts of different policies. When policymakers applying these lessons and changing policy face important institutional hurdles, such as in the constitutional and statutory policymaking contexts, then high-variance/low-expected-value policies become relatively undesirable. The transaction costs associated with policymaking change thus preclude the benefits of the optimal search approach. Many of these transaction costs are the result of deliberate choices rather than unavoidable realities. The next Section examines aspects of institutional design that hinder reversibility through the optimal search perspective. Applying the optimal search lens to these design questions reveals a hitherto unexamined cost of principles such as stare decisis and separation of powers.

E. Reversibility and Institutional Design

While some degree of policymaking inertia is inevitable, institutional choices contribute greatly to the degree of inertia. The optimal search approach highlights a cost of this inertia; inertia discourages high-variance/low-expected-value policies that improve social welfare in the long run only if the lessons learned from these policies can be implemented in future policy. This Section examines how the value of reversibility and optimal search informs some specific questions of institutional design.

Most advocates of policymaking transaction costs are well aware that these costs hinder the development of some efficient policies.¹⁵² They advocate

152. See, e.g., Stephenson, *supra* note 143, at 15-16 (acknowledging that judicially imposed costs to the enactment of new legislation may lead to the prevention of “socially desirable,

institutional principles such as separation of powers because they believe that these institutions reduce the risk of excessive concentration of power.¹⁵³ The purpose of the discussion in this Section is not to contradict these claims. Instead, this Section argues that these advocates' analysis of policy transaction costs is incomplete. While everyone agrees that transaction costs prevent the enactment of some policies with positive expected value and accepts this as a cost worth bearing, most overlook the potential learning benefits of high-variance/low-expected-value policies. The previous Sections have shown that such policies are often extremely desirable because of their potential to provide better outcomes and better knowledge. Thus, policy transaction costs do not merely delay some good policies, but also prevent the dynamic benefits of learning through variation. Preventing such learning may be a greater cost of institutional transaction costs than the simple prevention of some appealing policies.

Indeed, the importance of reversibility in institutional design may have some positive salience in addition to the normative approach taken through most of this Article.

1. *Sunset Clauses and Legislative Entrenchment*

The difference between legislative entrenchment and sunset clauses demonstrates the distinctions between the real options approach and the optimal search approach.¹⁵⁴ Legislative entrenchment, which is unconstitutional,¹⁵⁵ is "the enactment of either statutes or internal legislative rules that are binding against subsequent legislative action in the same form."¹⁵⁶ One implication of the rule against legislative entrenchment is that legislatures cannot make irreversible policies. Sunset clauses, by contrast, "cause a statute to lapse, by operation of law, after a defined period [and] are

constitutionally justifiable legislation" if, for example, the judiciary has incomplete information).

153. See, e.g., *Myers v. United States*, 272 U.S. 52, 293 (1926) (Brandeis, J., dissenting) (recognizing that the separation of powers aims "not to promote efficiency but to preclude the exercise of arbitrary power"); Martin S. Flaherty, *The Most Dangerous Branch*, 105 YALE L.J. 1725, 1826-28 (1996) (praising separation of powers but recognizing that it hinders "action").

154. For an examination of legislative entrenchment doctrines, see Eric A. Posner & Adrian Vermeule, *Legislative Entrenchment: A Reappraisal*, 111 YALE L.J. 1665, 1665-66 (2002).

155. See *Newton v. Comm'rs*, 100 U.S. 548, 559 (1879); 1 WILLIAM BLACKSTONE, COMMENTARIES *90.

156. Posner & Vermeule, *supra* note 154, at 1667.

LEARNING THROUGH POLICY VARIATION

the mirror image of entrenching clauses.”¹⁵⁷ In some sense, almost all policies have sunset clauses—policies automatically lapse when new policies on the same subject are instituted. Sunset clauses simply reduce the cost of changing policies in the future. Instead of requiring costly effort to change a policy, a sunset clause reverses the policy by default. Sunset clauses therefore enhance the reversibility of policies.

Posner and Vermeule call for the abolition of the restriction on “legislative entrenchment.”¹⁵⁸ They argue that legislative entrenchment enables government to commit itself to a certain course of action, thereby increasing the ability of individuals and other government actors to take actions that rely on the government continuing with its course of action.¹⁵⁹ They also note that “[e]ntrenchment is no more objectionable in terms of constitutional, political, or economic theory than are sunset clauses.”¹⁶⁰

Contradicting Posner and Vermeule, the optimal search and real options perspectives developed here explain why entrenchment is barred while sunset clauses are legitimate from an economic perspective. The real options approach demonstrates why irreversible policies with positive expected value—exactly the type of laws that Posner and Vermeule suggest should be entrenched—are often suboptimal choices because of the degree to which they restrict future policymaking options. The real options approach demonstrates that flexibility, which is impeded by legislative entrenchment, often has greater value than a policy with a higher expected value in a dynamic setting. If legislators appropriately value future flexibility, then no restriction on legislative entrenchment is necessary—the legislators will appropriately weigh the commitment benefits of entrenchment against the loss of flexibility. Of course, if legislators are so forward-looking, then the benefits of entrenchment—primarily, the promise that future legislators will not renege on a policymaking commitment—are negligible. Forward-looking policymakers will realize the loss of credibility incurred by reneging on a policy and will only renege on the policy if the benefits outweigh the costs, making entrenchment unnecessary. If legislators cannot value the future appropriately, then there is a good possibility that they will not place an appropriate value on future flexibility, which the real options approach demonstrates can be considerable. To protect the considerable benefits of future flexibility, entrenchment should be restricted.

157. *Id.* at 1676.

158. *Id.* at 1666.

159. *See id.* at 1670-73.

160. *Id.* at 1666.

Sunsetting, by contrast, presents the opposite implications. It makes policies more reversible. In the optimal search approach, this is unambiguously positive, as it enhances the search for excellent policies. Indeed, the optimal search approach suggests that policy sunseting justifies the adoption of policies with negative expected value. Sunsetting therefore enhances efficient policymaking, while legislative entrenchment hinders efficiency.

As a result, enhanced use of sunset clauses should be encouraged to enable public policymakers to gain more of the learning benefits of the optimal search approach in the face of irreversible costs. Suppose that there are multiple policies that should be tried in a given order under the optimal search approach. Passing each of these policies in succession would be costly. It would require policymakers to learn and relearn the benefits of each policy alternative and also would require that each policy alternative make its way onto the legislative agenda. These costs may well prevent policymakers from choosing policies according to the optimal search approach's prescriptions.

The use of sunset clauses can reduce these costs. Policymakers can pass a law that specifies that each policy should be tried for a given amount of time, to be followed by the next policy. If any experimental policy were deemed particularly successful, then future policymakers would be free to enshrine that policy as the permanent law at any point. Inertia, however, would lead to policy change rather than policy stagnancy. The final sunset provision could revert back to the law that existed before the reforms, or it could use a "penalty sunset" analogous to the penalty default rules described above. A penalty sunset would introduce an unpleasant final law that would strongly encourage future legislators to overcome policymaking inertia. Once legislators overcome this inertia, it is likely (though far from guaranteed) that the lessons learned through this statutory optimal search would be heeded.

2. *Separation of Powers*

Policy reversibility is partially determined by the degree to which powers are separated. Separation of powers refers to a government division into separate branches, with each branch holding the ability to check other branches from making policy.¹⁶¹ Policy change when policymaking power is divided is much more difficult and expensive than policy change when policymaking power is concentrated. For example, statutory policymaking in the United States—an institutional design marked by separation of powers—requires the approval of two houses of a legislative body as well as the executive. Statutes

161. See THE FEDERALIST NOS. 48, 51 (James Madison).

LEARNING THROUGH POLICY VARIATION

are then subject to judicial review. Approval from these disparate bodies requires political movements to win several consecutive elections before obtaining unobstructed policymaking authority.¹⁶² Unicameral parliamentary systems, by contrast, require only one such vote of approval. Once a political movement gains control of parliament in a single election, they may enact policy changes with relatively few institutional hurdles.

Separation of powers offers important costs and benefits. Because each branch of government checks the others, separation of powers prevents the accumulation of centralized power and limits the opportunity for abuse of power.¹⁶³ Separation of powers also produces greater deliberation and specialization within government.¹⁶⁴ These benefits of separation of powers do not come without a cost, as even advocates of separation of powers acknowledge.¹⁶⁵ “The price of separation is that it makes it more difficult for the federal government to act – whether for good or bad purposes.”¹⁶⁶ In other words, separation of powers makes all policies less reversible by raising the transaction costs of making and changing policy.

This discussion does not intend to dispute any of the purported benefits of separation of powers, nor does it attempt to conclude that power in the U.S. government is currently overseparated or underseparated. Instead, the optimal search approach shows an underestimated cost of making it more difficult for the government to act. Not only does such a restriction slow down the passage of some good laws, as is well appreciated, but high policymaking costs also alter the optimal choice of policies. If it is easy for the government to act, the government can choose high-variance/low-expected-value policies and learn from these policies, changing them when they fail. When reversing policies becomes difficult, as is the case with separation of powers, the benefits of policy variance are reduced. Policies with low expected value become increasingly unattractive when they cannot be changed easily. As a result, separation of powers hinders learning through policy variation in addition to simply slowing down the process of government.

162. See Bruce Ackerman, *The New Separation of Powers*, 113 HARV. L. REV. 633, 650 (2000).

163. See THE FEDERALIST NO. 47 (James Madison).

164. See Douglas W. Kmiec, *Of Balkanized Empires and Cooperative Allies: A Bicentennial Essay on the Separation of Powers*, 37 CATH. U. L. REV. 73, 75-76 (1987).

165. See, e.g., *supra* note 153 (quoting Justice Brandeis).

166. Rachel E. Barkow, *Separation of Powers and the Criminal Law*, 58 STAN. L. REV. 989, 991 (2006).

3. *Stare Decisis*

The reversibility of judicial policies is partially determined by the degree to which judges adhere to *stare decisis*.¹⁶⁷ Absolute *stare decisis*, when a court never overrules its own previous decisions, implies irreversibility of judicial policy. Less stringent forms of *stare decisis*, in which a court avoids overruling precedent absent an unusually strong justification, limit but do not prevent policy reversals. *Stare decisis* raises the transaction costs of policy change. In addition to the ordinary costs of changing policy, such as informing the relevant parties regarding the change, *stare decisis* requires judges to violate a principle when changing policy. Other things equal, raising the costs of policy change will reduce the number of changes.

Stare decisis is the subject of a vast literature.¹⁶⁸ Some praise *stare decisis*, claiming that stability and moderation are particularly important for judges and that “restraint in decision-making and respect for decisions once made are the keys to preservation of an independent judiciary and public respect for the judiciary’s role as a guardian of rights.”¹⁶⁹ *Stare decisis* also has many critics. Justice Scalia, for example, has stated, “I would think it a violation of my oath to adhere to what I consider a plainly unjustified intrusion upon the democratic process in order that the Court might save face.”¹⁷⁰

As with the separation-of-powers discussion, I do not intend to weigh the many positives and negatives of *stare decisis*. I believe, however, that both the advocates and detractors of *stare decisis* have overlooked a cost that is illuminated by the optimal search perspective. The cost of *stare decisis* is not simply that it allows incorrect decisions to linger. In addition, *stare decisis* changes the initial optimal decision for judges. Instead of choosing high-variance/low-expected-value policies and gaining the learning benefits of the optimal search approach, *stare decisis* pushes judges away from high-variance policies.¹⁷¹ When judicial policies are difficult or impossible to reverse because

167. For an illuminating discussion of *stare decisis*, see RICHARD A. POSNER, *THE FEDERAL COURTS: CHALLENGE AND REFORM* 371-82 (1996).

168. For one recent discussion, see Daniel A. Farber, *The Rule of Law and the Law of Precedents*, 90 MINN. L. REV. 1173 (2006).

169. Lewis F. Powell, Jr., *Stare Decisis and Judicial Restraint*, 47 WASH. & LEE L. REV. 281, 289-90 (1990).

170. *South Carolina v. Gathers*, 490 U.S. 805, 825 (1989) (Scalia, J., dissenting).

171. Judges can choose policies explicitly under the common law. See *Republican Party of Minn. v. White*, 536 U.S. 765, 784 (2002) (noting the power of state judges to “make” common law). When they are interpreting statutes, judges enjoy less discretion, as they are constrained by principles of statutory interpretation. Nevertheless, judges can choose

LEARNING THROUGH POLICY VARIATION

of stare decisis, judges will choose policies with high expected value that are unlikely to need changes. Extreme stare decisis leads to judicial policymaking under the real options approach, in which a low-expected-value precedent is kept while policymakers learn more about a potentially better option. As demonstrated above, the optimal search approach leads to better long-term outcomes, all else equal. Therefore, the foregone optimal search benefits represent an important and overlooked cost of stare decisis.

This Section applied the optimal search perspective described above to questions of institutional design. It is generally assumed that current separation of powers and stare decisis outcomes are viewed as a reasonable balancing of the checking benefits of separation of powers and stare decisis against their efficiency costs. This balance, however, excludes the optimal search perspective. The addition of the optimal search perspective impels institutional design toward less separation of powers and reduced emphasis on stare decisis.

The analysis of separation of powers and stare decisis introduced a consideration deemphasized in the earlier parts of this Article—the possibility (or inevitability) that public policymakers will not always choose policies that maximize the common good. The next Part further examines the interplay of public choice analysis and the optimal search approach.

IV. PUBLIC POLICYMAKING INCENTIVES AND THE OPTIMAL SEARCH APPROACH

This Article has explored the best policymaking choice for a public policymaker aiming to maximize long-run social welfare. The discussion has been normative, concerning how policymakers should choose policy to maximize long-term social welfare. When reversibility is low, high-variance policies are the best choice for this policymaker. But what if public policymakers have other goals, such as maintaining their positions or maximizing their own wealth? Under these conditions, will high-variance policies be chosen and will optimal search policies become a cover for policies that pursue other motives? In addition, are there any mechanisms for realizing the benefits of optimal search policies in the presence of self-interested public policymakers?

interpretive methods to implement a variety of different outcomes when interpreting statutes. See, e.g., Alexander Volokh, *Choosing Interpretive Methods: A Positive Theory of Judges and Everyone Else*, 83 N.Y.U. L. REV. 769 (2008). Thus, even when engaged in statutory interpretation, judges effectively “choose” policies to a certain degree, though they are constrained by principles of statutory interpretation.

This Part analyzes policymakers' incentives to enact high-variance policies. I will assume that policies are reversible without cost in order to focus on the role of incentives in optimal search. In addition, I assume that, instead of maximizing social welfare, public policymakers seek to jointly maximize some combination of social welfare and their own individual welfare.¹⁷² The individual welfare of policymakers may be increased through money, lack of criticism, reelection, and professional advancement. Some of these considerations will impede the production of high-variance policies via the optimal search approach while others will enhance such production.

A. High-Variance Policies and Reelection/Reappointment Incentives

First, consider a public policymaker seeking reelection to a position chosen by public vote. This policymaker chooses policies designed to convince the public to vote for her in upcoming elections. Both the public's risk preferences and the politician's status determine whether optimal search policies are pursued.

Assume that a successful policy increases reelection probabilities and that a failed policy decreases such probabilities. Under these conditions, high-variance/low-expected-value policies will not be chosen by policymakers. These policies are likely to fail and therefore hurt the policymaker's reelection chances, in spite of their long-run efficiency. The political process is unlikely to produce the optimal amount of policy variation for learning. The underproduction of policy variance will be exacerbated when voters are risk-averse, such that if given a choice between two portfolios of policies, they will choose the portfolio with a certain, albeit lower, expected value over a riskier portfolio with a higher expected value. Furthermore, if voters focus on the effects of new policies rather than old policies and most of the new policies are failures, then policymakers will avoid high-variance/low-expected-value innovation at all costs.

When donations from special interests raise reelection chances, then the optimal search approach may even provide "cover" for politicians seeking reelection by currying favor with the special interests. Suppose that policymakers seeking reelection must weigh the value of money provided by special interests against the electoral cost of instituting policies that favor the special interests but harm the average voter. In these conditions, policymakers may claim that a poor policy designed to favor a special interest is actually a high-variance optimal search policy. If the public cannot distinguish between

¹⁷² The discussion in this Part is influenced by Rose-Ackerman, *supra* note 108.

LEARNING THROUGH POLICY VARIATION

special-interest pandering and policies with genuinely high upsides, then the policymaker can reduce the cost of special-interest pandering. This cover will raise the amount of special-interest pandering.

This analysis, however, assumes that voters judge each policy simply by success or failure and are unable to distinguish policies with high upsides from policies that pander to special interests. If voters value extremely successful policies more than ordinary successes, and policy failures can be changed, then a package of high-variance policies may increase reelection chances. Policymakers can keep the successes, drop the failures, and achieve a net benefit for the public, increasing reelection chances. Relatedly, policymakers can educate voters about the benefits of variance. If voters understand the long-run value of policy change and social welfare is distributed evenly, then voters will reward policymakers choosing high-variance policies because of their long-run value. Finally, the cost of special-interest pandering is reduced with highly reversible policies. A policy may benefit a special interest, but if it is a failure and reversibility is easy, then such a policy may not last long. If failed policies benefiting special interests are quickly reversed, then special interests may have less incentive to pursue policies with no upside that simply benefit the special interest. Thus, depending upon the preferences of voters and the ease of reversibility, reelection incentives can either facilitate or impede the high-variance policies favored under the optimal search approach.

Political risk-taking incentives will also be affected by perceived reelection probabilities. Politicians expecting to lose future elections may be encouraged to choose high-variance/low-expected-value policies, because a very successful policy may enable future election victories while failure just makes already low approval ratings even lower.¹⁷³ While this appears to be a negative consequence of the “knife-edge” significance of receiving a majority of the vote,¹⁷⁴ the optimal search approach demonstrates that such policies may be socially desirable. Similarly, individuals with no chance of being reelected, such as politicians facing term limits, may be more likely to choose high-variance policies as they have “nothing to lose” by pursuing significant policy successes.

A similar analysis applies to appointed policymakers seeking reappointment. This category includes administrative policymakers whose job security depends upon the approval of elected officials or judges with fixed terms subject to reappointment.

173. See Susan Rose-Ackerman, *Risktaking and Electoral Competition*, 7 EUR. J. POL. ECON. 527, 527 (1991).

174. *Id.*

If reappointment depends upon the preferences of elected policymakers, then appointed policymakers may internalize the policymaking preferences of their appointers. In that case, policymaking by appointed policymakers seeking reappointment will be identical to the policymaking of elected policymakers discussed above.

In other cases, however, reappointment may depend upon simply avoiding substantial failures or any negative attention. Most policymakers may get reappointed, with only conspicuously poor policymakers getting denied reappointment. Under these conditions, the high-variance policies favored by the optimal search approach will not be chosen. The policymaker seeks to minimize the downside risk associated with the policies she chooses rather than the social welfare of the policies or the learning benefits of variance. High-variance policies have a higher probability of “failure” and therefore will not be chosen.

B. High-Variance Policies and Incentives for Political Advancement

While reelection or reappointment incentives will often reduce policymakers’ incentives to choose high-variance policies, the potential for policymaking advancement may raise the prevalence of high-variance policies. Because it is difficult to achieve higher office, candidates may be forced to take risks to attain such office. The risky choices policymakers take, however, will not always match the choices favored by the optimal search perspective.

Suppose, for simplicity, that all policymakers desire higher elected or appointed office. For example, there may be many congresspersons who strive to be elected to the Senate or to executive positions such as state governor or President. Similarly, there are many district court judges who have ambitions of a seat on an appeals court or the Supreme Court. Suppose further that the policymakers achieving higher office are the ones choosing the best policies. For example, one hundred senators are chosen from 435 representatives by choosing the one hundred representatives whose policies have produced the best outcomes.

Consider the incentives of a representative who is seeking advancement and must select from a set of policies. The representative aims to choose a policy that is among the one hundred best. To do this, the congressperson must take risks. If each representative chooses policies that maximize expected value, then a congressperson has less than a one-in-four chance of attaining the Senate. High-variance policies, by contrast, may provide a better chance of placing in the top hundred than expected-value maximization policies. These high-variance policies may have greater downsides than other policies, but congresspersons are less interested in the magnitude of the downsides. If a

LEARNING THROUGH POLICY VARIATION

policy falls below expectations, the representative is unlikely to advance, regardless of whether the failure is of small or great proportions. Thus, ambitions for higher positions may induce politicians to choose high-variance policies. Indeed, because the magnitude of downside risks plays almost no role in policy advancement, policymakers may give even less weight to poor outcomes than they would under the optimal search approach.

The scarcer the office, the greater the risks that will be taken to attain that office. No representative is likely to find the single most successful policy out of 435 without realizing the high upside of a high-variance policy. If there were 435 representatives for four hundred Senate seats, by contrast, congresspersons would be likely to become risk-averse, as they would be likely to attain the Senate so long as they were not associated with a failed policy with an undesirable downside.

For the most part, however, political hierarchies display a pyramidal shape. There are many policymakers on one rung competing for the rung above. High variance should therefore be induced by a government with multiple levels of hierarchy. Politicians may seek both reelection and advancement. They thus face conflicting incentives: advancement requires risk taking, while reelection may discourage risk. In total, the most political risks might be taken by a politician seeking advancement who has little fear of failing to gain reelection. The fewest risks will be taken by a politician in a competitive district who has no ambition for higher office. Because risky policies have high value under the optimal search perspective, “safe” districts may have greater political value than commonly believed.¹⁷⁵

When advancement requires appointment by one party and approval by another, such as appointment to the Supreme Court, risk-taking incentives are altered. On the one hand, judges aspiring to the Supreme Court must take enough risks to differentiate themselves from the mass of judges of similar ideology. If they do not, they are unlikely to get noticed or appointed. On the other hand, judges must take care to avoid offending the Senate and inducing a filibuster on their nomination. Given these conflicting incentives, risk-taking incentives for judges pursuing advancement are uncertain.

175. Cf. Michael S. Kang, *Race and Democratic Contestation*, 117 YALE L.J. 734 (2008) (arguing that competitive electoral districts are less important than creating conditions of “democratic contestation”). *But see* Samuel Issacharoff & Richard H. Pildes, *Politics as Markets: Partisan Lockups of the Democratic Process*, 50 STAN. L. REV. 643, 646 (1998) (“Only through an appropriately competitive partisan environment can one of the central goals of democratic politics be realized: that the policy outcomes of the political process be responsive to the interests and views of citizens.”).

C. Federalism and Incentives To Innovate

Subsection II.C.4 demonstrated that the benefits of experimentation increase in a multi-jurisdictional context, as learning through policy variation occurs across space as well as across time. Unfortunately, incentives to innovate in a federal system may be no greater than incentives to innovate in a uni-jurisdictional system, in spite of the added benefits of such innovation. Several factors explain why federalism may not generate sufficient useful experimentation.¹⁷⁶

First, federalism creates a free-rider problem for innovation. A policymaker in a state will strongly prefer to allow other states to attempt a high-variance/low-expected-value policy such as Policy B. The policymaker thereby gains the option value of Policy B without bearing its cost, free-riding on the efforts of the state that chooses Policy B. If every state hopes to free-ride, no state will innovate and the benefits of federalism as a laboratory will not be realized. Indeed, in some cases federalism leads to less innovation than does a uni-jurisdictional system.¹⁷⁷

176. For a thorough examination of the effects of federalism on innovation from a positive political theory perspective, see Rose-Ackerman, *supra* note 108.

177. To see this point, consider two jurisdictions that face the choice between Policy A and Policy B of Example 2. Section II.A demonstrated that in a uni-jurisdictional world, the optimal search approach dictates that Policy B is the efficient choice. When there are two jurisdictions, the choice of innovation (Policy B) versus no innovation (Policy A) can be modeled as the following game:

		JURISDICTION 2	
		POLICY A	POLICY B
JURISDICTION 1	POLICY A	\$500, \$500	\$628, \$586
	POLICY B	\$586, \$628	\$586, \$586

The value of \$628 in the game modeled above comes from retaining Policy A in the first period and observing the outcome of Policy B in the other jurisdiction. If Policy B fails in the other jurisdiction, then continue with Policy A in the second period. If Policy B succeeds in the other jurisdiction (and is observed), then switch to Policy B in the next period and enjoy the high upside of Policy B. In other words, one jurisdiction can free-ride off the other jurisdiction's experimentation with Policy B.

There is no pure strategy Nash Equilibrium to this game and therefore no unique solution. See, e.g., ANDREU MAS-COLELL, MICHAEL D. WHINSTON & JERRY R. GREEN, MICROECONOMIC THEORY 248-50 (1995). Choosing either policy is a rationalizable strategy

LEARNING THROUGH POLICY VARIATION

Second, if voters choose their elected officials based on outcomes in that jurisdiction exclusively, then each elected official will maximize his or her own jurisdiction's welfare. In these circumstances, the total amount of policy variance will be inadequate—even if no jurisdiction seeks to free-ride on the experimentation of another. Consider Policies *A* and *B* in Example 3 and assume that no jurisdiction seeks to free-ride on the experimentation of another jurisdiction and that no jurisdiction accounts for benefits of policy variance that accrue to other jurisdictions. Thus, both jurisdictions will treat the policy choice between Policy *A* and Policy *B* as if they were in a uni-jurisdictional setting. As shown in Subsection II.C.4.b, both jurisdictions will therefore choose Policy *A*, in spite of the fact that Policy *B* adds to social welfare.

Third, many laws are subject to network externalities.¹⁷⁸ In a network, the overall benefits of a good are proportional to the total number of individuals (the “network”) who use the good. A telephone, for example, is worthless if there is only one telephone, but it is an incredibly useful communication tool when there are many. Similarly, individuals benefit from familiar laws. Such laws save individuals the costs of having to learn about peculiarities in the law. This network effect further reduces the incentive to innovate. If someone invents a new communication system that is just as good (or even somewhat better) than the telephone but cannot be used in the current telephone system, the new system is unlikely to gain traction because the new system has no other users and no network. Similarly, an innovative law that is just as good as or slightly better than the existing law is unlikely to be tried because of the costs of building a new network around the innovative law. For example, a state may consider a new contract default better than the rule prescribed in the Uniform Commercial Code but decide that introducing a rule that is different from the rule of other states might raise costs for companies that do business in that state and others because they will have to deal with two sets of laws.

Fourth, when externalities or other market failures are a problem, federalism may not generate appropriate incentives. For example, if pollution in one state harms individuals in another state, then politicians in the first state will not have the appropriate incentives to regulate pollution. The individuals

for both jurisdictions, however. Therefore, we can make no definitive predictions about whether there will be innovation in this simple multi-jurisdictional world. The uni-jurisdictional world, by contrast, provides innovation using the optimal search approach. In this example, federalism can decrease innovation, but cannot increase it. When innovation is less attractive in the uni-jurisdictional world, however, there can be cases where federalism produces more innovation than does the optimal search approach.

178. See Michael Klausner, *Corporations, Corporate Law, and Networks of Contracts*, 81 VA. L. REV. 757, 764 (1995).

in the second state who are harmed by the pollution do not factor into the reelection prospects of a politician in the first state. Indeed, the presence of externalities may lead to a “race to the bottom.”¹⁷⁹ If pollution’s effects are largely out of state, then a state that has lax pollution restrictions will attract more businesses than a state with strict restrictions because the state with lax restrictions is a cheaper place to do business. To avoid the loss of businesses, the state with strict restrictions may loosen its pollution restrictions until all states have the level of pollution restrictions desired by the states with the most lax attitudes toward pollution.

In total, free-rider problems, failure to internalize benefits to other jurisdictions from innovation, network externalities, and spillovers may reduce policy variation in multi-jurisdictional systems to a level far below optimality.

V. IRREVERSIBILITY AND POLITICAL INCENTIVES: APPLICATIONS AND RECOMMENDATIONS

The previous three Parts dealt with several objections to the choice of high-variance policies. When policies are irreversible, high-variance/low-expected-value policies become undesirable. In addition, public choice incentives may blunt the formulation of high-variance/low-expected-value policies even when they are desirable. This Part therefore revisits the applications of the optimal search approach presented above in light of these complications.

A. Reversible Regulations

By assuming reversibility, Example 2 assumed the result. High-variance/low-expected-value policies are optimal searches when policies are reversible. Real-life regulations assume a sliding scale of reversibility. Regulations concerning pollutants with very short atmospheric lives, for example, will be more reversible than regulations concerning pollutants that linger. The desirability of high-variance/low-expected-value regulations depends critically upon where along this reversibility spectrum the regulation falls. Consequently, the informational value of a policy that would be analyzed under the OMB’s *Circular No. A-4* should include an examination of the potential persistence of a regulation’s effects.

179. Corporate law has a longstanding debate about whether state competition for corporate charters entails a race to the bottom or a race to the top. For important contributions to this literature, see *supra* note 73.

LEARNING THROUGH POLICY VARIATION

Regulations may also be subject to inertia caused by several factors. Parties that benefit from a failed regulation may fight harder to keep the regulation than they would to pass it in the first place due to status quo bias and loss aversion.¹⁸⁰ In addition, regulators seeking reappointment may avoid any policy change since such a change may reduce the chances of reappointment by “rocking the boat.”

It should be noted, however, that inertia in regulatory activities should be lower than in other contexts. Regulators face lower institutional barriers to changing policy than do other policymaking officials. For example, a new administration can change regulatory policy far more rapidly than it can change a statutory policy. In addition, inertia can be limited by enacting a sunset provision limiting the term of a high-variance/low-expected-value policy. This enables regulators to maximize the expected gains of learning through policy variation.

But what of regulators’ incentive to innovate? If each regulator controls one policy, and his or her chances of reappointment depend upon the success of that policy, he or she will have little incentive to choose high-variance/low-expected-value policies. This obstacle may prove to be a formidable barrier to innovation in many contexts.¹⁸¹ To increase the incentive to innovate, individual decisionmakers should have control over several policies. This will enable them to reap some notable successes from experimentation in addition to failures. Alternatively, regulatory policymakers might be granted extended terms to strengthen their incentive to maximize long-term social welfare rather than career chances.

B. Contract Default Rules and Increased Judicial Policymaking

The levels of irreversibility associated with contract default rule innovations are primarily associated with information costs. To properly assess the effects of a rule, contracting parties must know of the rule. Changing a rule therefore requires informing parties about the change. If the change proves to be a failure, then the new default rule should be reversed, incurring yet another

¹⁸⁰. See Kahneman et al., *supra* note 132.

¹⁸¹. Note, however, that there have been several examples of regulatory bodies running innovative policy experiments, such as the Moving to Opportunity experiment, see Jeffrey R. Kling, Jens Ludwig & Lawrence F. Katz, *Neighborhood Effects on Crime for Female and Male Youth: Evidence from a Randomized Housing Voucher Experiment*, 120 Q.J. ECON. 87 (2005), and the SEC’s pilot order restricting short sales of certain securities, see Order Extending Term of Short Sale Pilot, Exchange Act Release No. 53,684, 71 Fed. Reg. 24,765 (Apr. 20, 2006).

round of switching costs. The greater these switching costs relative to the long-run value of a policy, the more reluctant judges should be to choose high-variance optimal search policies. For example, mergers and acquisitions lawyers for publicly traded companies should be easier to educate about changes in contract law than the entire population of lawyers, who in turn are cheaper to inform than the public at large. As a result, enacting high-variance default rules on issues covered exclusively by public company transactional lawyers entails lower irreversible costs than do other contract default rule changes. All else equal, judges should develop innovative contract default rules in low cost contexts such as mergers and acquisitions of public companies.

Stare decisis poses another obstacle to the promulgation of new high-variance default rules. New default rules have higher variance in outcomes than existing rules, and stare decisis encourages judges to maintain the status quo. Stare decisis also poses an obstacle to reversing a failed high-variance default rule. Once the new rule is enacted, it gains precedential effect and is difficult to reverse. The irreversibility imposed by stare decisis should deter judges from instituting high-variance/low-expected-value rules, even if judges are willing to ignore stare decisis on their own account. If future judges adhere to stare decisis, then a high-variance/low-expected-value policy becomes relatively irreversible, negating the informational value of such a policy.

Even if stare decisis and learning costs are minor obstacles, judges may have inadequate incentives to choose high-variance innovative default rules for a number of reasons. Judges enacting new rules may be less likely to gain reelection because many of these rules will, on average, lead to bad outcomes, thereby encouraging voters to unseat the incumbent. Finally, in a multi-jurisdictional world, a jurisdiction with an innovative and idiosyncratic default rule loses out on network effects, raising costs for lawyers. Together, these factors may cause innovation in judicial policymaking to be far below its optimal level.

Other characteristics of judicial policymaking may counteract the aforementioned factors that diminish judicial risk taking. Many judges are appointed for life terms. As a result, they are free to prioritize social welfare over some other factor, such as reappointment or reelection probabilities. Because high-variance/low-expected-value policies have such high informational value, judges may be the best situated candidates to choose such optimal search policies. In addition, judges may choose high-variance/low-expected-value policies for selfish reasons, such as notoriety. A decision that adheres to precedent is unlikely to garner many citations and attention. Decisions overturning precedent, by contrast, will be instantly controversial. If

LEARNING THROUGH POLICY VARIATION

the new policy articulated in a decision proves to be a success, then the judge has likely established his or her reputation.¹⁸²

Sunset clauses for high-variance judicial policies may also have value. Such sunset clauses—such as the “twenty-five-year window” for affirmative action articulated by Justice O’Connor in *Grutter*¹⁸³—allow high-risk policies to be attempted while reducing the risk of irreversibility.

C. Shareholder Power

Shareholder voting policies currently exhibit significant inertia. Even seemingly minor changes that have been embraced without obvious effect in nations like Great Britain¹⁸⁴—such as allowing a nonbinding shareholder vote on executive compensation¹⁸⁵—raise significant outcries in the United States. Whatever the explanation, proponents of the status quo have successfully derailed significant changes to corporate voting rules, such as shareholder access to the corporate ballot. Shareholder power has witnessed very little learning through policy variation.

There are several explanations for the inertia in corporate voting policies in the face of serious disagreement about effective policy. The holdup is not structural in nature—the SEC could change shareholder access to the corporate

182. For example, Judge Baron Alderson gained considerable notoriety for instituting an information-forcing default rule in *Hadley v. Baxendale*. An interesting avenue for future research would be to compare life-appointed judges with elected judges to determine whether electoral incentives promote innovation through variation in private law rulemaking. One might hypothesize that life-appointed judges, removed from electoral pressures to avoid negative-expected-value common law rules, are more willing to enact new legal doctrines with high variance. See David E. Pozen, *The Irony of Judicial Elections*, 108 COLUM. L. REV. 265, 324-28 (2008) (describing scholarly support for elected judges’ more innovative or activist approaches and noting that state courts’ “common-law lawmaking powers are broadly respected and [their] decisions are relatively easily reversed through constitutional amendment or legislative action”). I thank William J. Rinner for raising this point.

183. See *Grutter v. Bollinger*, 539 U.S. 306 (2003). In *Grutter*, the Court noted that it “expect[s] that 25 years from now, the use of racial preferences will no longer be necessary to further the interest approved today.” *Id.* at 343.

184. Note how international differences can also serve to produce useful variation in laws. One concern, of course, is that differences between nations are so great that one cannot infer the effects of a variation in Country X will be similar to the effects of the same variation in Country Y.

185. See Sourafel Girma, Steve Thompson & Peter W. Wright, *Corporate Governance Reforms and Executive Compensation Determination: Evidence from the UK*, 75 MANCHESTER SCH. 65, 66-67 n.5 (2007).

ballot simply by issuing such a rule.¹⁸⁶ Instead, the SEC may prefer to avoid making any decision on a contentious issue such as shareholder access that has attracted so much attention.¹⁸⁷ Upsetting the status quo may lead to more trouble than would any other policy. This may explain the SEC's initial formulation and subsequent retraction of the shareholder access rule. Another cause of inertia may be that the lobbying power of those supporting the shareholder power status quo exceeds the lobbying power of those in favor of increasing shareholder power. Yet another explanation may simply be that the SEC views the status quo as having higher expected value than the proposed changes. Finally, the SEC may prefer to avoid the costs associated with the switchover to a new voting regime.

None of these explanations is perfectly satisfactory from a normative or even positive perspective. The SEC's attempt to pass a rule providing shareholder access to the corporate ballot stirred a maelstrom of controversy,¹⁸⁸ meaning that preserving the status quo failed to insure a quiet life for regulators. Moreover, avoidance of controversy deserves little to no normative weight. The asymmetric lobbying power explanation may have significant positive salience under a Republican administration, but lobbying power should be accorded little normative weight. The higher expected value for the status quo claim also fails from a normative perspective. If increasing shareholder power is easily reversible, then the optimal search approach demonstrates that the informational value of an innovative increased shareholder power policy probably trumps the high expected value but lower variance associated with the status quo. Finally, increasing (or subsequently decreasing) shareholder power is unlikely to entail prohibitive expenses.

186. See Roel C. Campos, Comm'r, SEC, Speech by SEC Commissioner: The SEC's Shareholder Access Proposal (Jan. 10, 2005) (transcript available at <http://www.sec.gov/news/speech/spcho11005rcc.htm>) ("Based on these considerations, I believe that the [shareholder access] rule is fully within our authority. I am not alone in my view. Several academics have expressed support for the SEC's authority to adopt the proxy access rule."). *But see* Bus. Roundtable v. SEC, 905 F.2d 406 (D.C. Cir. 1990) (suggesting limitations on the SEC's power to regulate substantive aspects of corporate voting).

187. The SEC received over thirty-four thousand comments on the issue. See Nazareth Speech, *supra* note 100.

188. See, e.g., Broc Romanek, *SEC Quickly Adopts Non-Shareholder Access Rule - Then Fireworks Erupt*, TheCorporateCounsel.net Blog, Nov. 29, 2007, <http://www.thecorporatecounsel.net/blog/archive/001638.html> ("The real fireworks began right after the meeting, when all sorts of investor groups, members of Congress, etc. issued statements disapproving the SEC's rulemaking (eg. CalPERS; CII; RiskMetrics; AFL-CIO; Rep. Frank; Sen. Dodd) - and some approved (eg. Marty Lipton). It's notable that the opposition is fairly organized on this issue; I can't imagine something like this happening even five years ago.").

LEARNING THROUGH POLICY VARIATION

Corporate voting primarily concerns two concentrated and experienced bodies—institutional shareholders and corporate management—who can cheaply be informed of any rule change. Compared to the purported benefits of “corporate democracy,” the switching costs appear quite low.

Indeed, learning through policy variation by instituting shareholder access to the corporate ballot offers the possibility of resolving this longstanding debate. Yet again, a “sunset provision” may offer a compromise that will satisfy both parties. While management advocates may refuse to accept a rule that guarantees shareholders access to the corporate ballot for an indefinite period, they may be willing to allow such access for a short period subject to a sunset provision. The sunset provision prevents inertia from allowing an inefficient shareholder access provision to linger. In addition, if shareholder access produces the parade of horrors that its opponents describe, then a brief period of shareholder access to the ballot will resolve the debate in favor of the status quo, conserving on lobbying energy. Conversely, advocates of shareholder access to the power would almost certainly prefer a limited trial of their desired policy than continued adherence to the status quo. If shareholder access proves to be the success its proponents foresee, then they will have a much stronger argument for reenacting the policy once it expires.

D. Federalism and Preemption

As discussed in Section IV.C, there exists a mismatch between the value of policy variance in a federalist context and the production of such variance. Free-riding and nonvaluation of informational benefits to other jurisdictions reduce policymakers’ incentives to enact high-variance policies relative to the socially optimal level. In addition, network externalities and other spillovers reduce the per-period value of policy differentiation between jurisdictions. In total, it is no surprise that federalism produces relatively few experiments.¹⁸⁹

Preemption doctrines exacerbate the problem. There is no question that by creating a legal network effect and eliminating spillovers from one jurisdiction’s policy to the next, preemption can improve outcomes in the current period. As discussed in Subsection II.C.4.c, however, this increase in per-period value comes at a significant cost—the elimination of information-producing policy variance that improves long-run policy outcomes. Because courts evaluating preemption have typically overlooked this dimension—learning through policy variation goes unmentioned in the Supreme Court’s preemption test—I advocate looser enforcement of preemption doctrine.

¹⁸⁹. See Rose-Ackerman, *supra* note 108, at 594.

Looser enforcement of preemption doctrine, however, does little or nothing to enhance policymakers' incentives to innovate. Innovative policies produce informational externalities—all jurisdictions learn from a single jurisdiction's experience. To encourage jurisdictions to internalize these externalities, Congress should adopt a system of grants to jurisdictions that enact innovative policies.¹⁹⁰ Grants will increase the incentives for innovation. Policymakers will weigh the expected benefits of the innovative policy and the value of the grant associated with the policy against the expected costs of policy innovation.

Grants are no panacea, however.¹⁹¹ Jurisdictions will have an incentive to claim that they are innovating to seek grants. Once they receive the grant, however, they will prefer to avoid making the potentially costly innovation. Without clear parameters regarding what does and does not constitute a policy innovation, grantmaking offers only a partial solution to the underproduction of policy variation.

E. Other Sources of Variation: Direct Experimentation

A federal system offers the possibility of learning through the experience of one jurisdiction without having to impose a high-variance policy on all jurisdictions. This is not the only way of achieving this goal, however. Policy variation can be produced at many different levels, such as the local, firm, or individual level. Drug trials, for example, induce variation in medication at the individual level; some participants in a drug trial receive one medication and some receive another. By experimenting on a small group of people and creating variance among them, drug experiments avoid the cost of experimenting on all individuals for some period.

Although they seldom do so today, policymakers can produce similar variation in public policies and learn from such variation. For example, instead of passing a rule granting shareholder access to the corporate ballot for a certain period to learn about the effects of such access, the SEC could randomly assign some companies to a shareholder access regime while allowing other companies to continue to prevent shareholder access.¹⁹² After observing the

^{190.} *See id.* at 615-16.

^{191.} *See id.* at 616.

^{192.} While random assignment may appear draconian, note that the SEC has randomly assigned corporations to different short-sale restriction regimes. *See* Order Extending Term of Short Sale Pilot, Exchange Act Release No. 53,684, 71 Fed. Reg. 24,765 (Apr. 20, 2006); SEC OFFICE OF ECON. ANALYSIS, ECONOMIC ANALYSIS OF THE SHORT SALE PRICE RESTRICTIONS UNDER THE REGULATION SHO PILOT (2007), <http://www.sec.gov/news/studies/2007/regshopiloto20607.pdf>.

LEARNING THROUGH POLICY VARIATION

effects of this policy variation for a certain period, the SEC would possess much greater information about the true effects of policy variance. In a related paper, Ian Ayres and I thoroughly analyze the promise and pitfalls of randomized public policy experiments, which are yet another means of learning through policy variation.¹⁹³ It should be emphasized, however, that current learning about the effects of policy occurs through policy change on all units within a jurisdiction rather than experimental assignment of policy variation to some subjects and not others.

CONCLUSION

This Article offers several contributions to theories of public policymaking. First, the Article emphasized that policies are not static and that this dynamism changes the policymaking calculus. In particular, policymakers should adopt the optimal search approach, which favors policies with higher variance (other things equal), whenever policies are reversible at low cost. This variance-preferring perspective has received little if any attention from scholars of policymaking, who generally debate whether policymaking should aim to maximize the average value of a policy or to take a risk-averse, Burkean approach.

Next, the Article expanded the dynamic policymaking context to include irreversible policies. When new policies cannot be changed and learning is possible, the policy calculus alters dramatically. Policymaking should become cautious to retain flexibility under the real options approach—a theoretical complement to the optimal search approach. When policies are partially reversible, then the flexibility benefits of the real options approach and the learning benefits offset each other, making a naïve expected-value maximization approach surprisingly attractive for previously unrecognized reasons. In sum, the best policy choice in the face of uncertain outcomes depends critically on the reversibility of the policy. Because some irreversibility is the deliberate outcome of institutional structures such as separation of powers, the optimal search approach demonstrates an underanalyzed cost of these designs—the prevention of policy improvement through learning from policy variation.

193. See Ian Ayres & Yair Listokin, *Randomizing Law* (Sept. 2008) (unpublished manuscript, on file with author).