INFORMATION ACQUISITION AND INSTITUTIONAL DESIGN

Matthew C. Stephenson

CONTENTS

INFORMATION ACQUISITION AND INSTITUTIONAL DESIGN

Matthew C. Stephenson[*]

Although good information is critical to effective decisionmaking, public agents' private incentives to invest in gathering information may not align with the social interest in their doing so. This Article considers how legal-institutional design choices affect government decisionmakers' incentive to invest in information, as well as how to manage the inevitable trade-off between promoting efficient use of information ex post and stimulating efficient acquisition of information ex ante. Using a simple theoretical framework, the Article considers a range of techniques for incentivizing information gathering, with particular attention to the structure of public institutions and public law.

INTRODUCTION

Good information is the lifeblood of effective governance.[1] In myriad contexts — from legislation to regulation to adjudication — public decisionmakers must operate in uncertain environments where the optimal choice depends, often substantially, on information about the likely consequences of different courses of action. This is not to say that these public decisions are or could be value-free exercises in technocratic neutrality. Nor is it to say that the relevant evidence must be of any particular type (for example, quantitative data) or that the analysis of such evidence must take any particular form (for example, cost-benefit analysis). It is rather to make the commonplace observation — so obvious that it ought to be uncontroversial — that many public decisions turn on some form of predictive judgment, such that a decisionmaker's choice does and should depend on the quality and content of the information available to her.

The importance of information is familiar to scholars who study how legal institutions structure and influence public decisionmaking.

Indeed, an important body of legal scholarship — which developed out of the Legal Process tradition — focuses on how to allocate authority among different potential decisionmakers (legislatures, chief executives, bureaucrats, judges, juries, voters, and others) in light of their relative "institutional competence," a phrase that connotes, among other things, expertise and access to information.[2] This literature also often considers how legal and political institutions might mitigate the principal-agent problem that can arise when the decisionmaker with the most expertise about some topic does not have the right preferences from the point of view of the relevant principal.[3] A related body

_____

[2] On the importance of institutional competence to the Legal Process school, see HENRY M. HART, JR. & ALBERT M. SACKS, THE LEGAL PROCESS 168–74, 696, 1009–10, 1111 (William N. Eskridge, Jr. & Philip P. Frickey eds., 1994); and William N. Eskridge, Jr. & Philip P. Frickey, *An Historical and Critical Introduction to* The Legal Process, *in* HART & SACKS, *supra*, at li, lx–lxi, xci–xcvi. *See also* Edward L. Rubin, Commentary, *The New Legal Process, the Synthesis of Discourse, and the Microanalysis of Institutions*, 109 HARV. L. REV. 1393, 1396 (1996); Jeffrey Rudd, *The Evolution of the Legal Process School's "Institutional Competence" Theme: Unintended Consequences for Environmental Law*, 33 ECOLOGY L.Q. 1045, 1052–54 (2006); Anthony J. Sebok, *Reading* The Legal Process, 94 MICH. L. REV. 1571, 1578–84 (1996) (book review).

Numerous modern scholars have built on the Legal Process tradition, producing a growing and increasingly sophisticated body of work that considers how public law institutions should allocate authority among various government decisionmakers, based on their relative institutional competence. *See, e.g.*, NEIL K. KOMESAR, IMPERFECT ALTERNATIVES 3–13 (1994); CASS R. SUNSTEIN, ONE CASE AT A TIME 259–63 (1999); ADRIAN VERMEULE, JUDGING UNDER UNCERTAINTY 63–85 (2006); Oren Bar-Gill & Elizabeth Warren, *Making Credit Safer*, 157 U. PA. L. REV. 1, 74–75 (2008); Mariano-Florentino Cuéllar, *Rethinking Regulatory Democracy*, 57 ADMIN. L. REV. 411, 490–97 (2005); Neal Devins, *Congressional Factfinding and the Scope of Judicial Review: A Preliminary Analysis*, 50 DUKE L.J. 1169, 1209–13 (2001); William N. Eskridge, Jr. & Lauren E. Baer, *The Continuum of Deference: Supreme Court Treatment of Agency Statutory Interpretations from* Chevron *to* Hamdan, 96 GEO. L.J. 1083, 1171–75 (2008); Richard H. Fallon, Jr., *Judicially Manageable Standards and Constitutional Meaning*, 119 HARV. L. REV. 1274, 1291 (2006); Philip P. Frickey & Steven S. Smith, *Judicial Review, the Congressional Process, and the Federalism Cases: An Interdisciplinary Critique*, 111 YALE L.J. 1707, 1713–16 (2002); Elizabeth Garrett & Adrian Vermeule, *Institutional Design of a Thayerian Congress*, 50 DUKE L.J. 1277, 1279–86 (2001); Clayton P. Gillette, *Expropriation and Institutional Design in State and Local Government Law*, 80 VA. L. REV. 625, 631–35 (1994); Jeff A. King, *Institutional Approaches to Judicial Restraint*, 28 O.J.L.S. 409, 432–40 (2008); M. Elizabeth Magill, *Beyond Powers and Branches in Separation of Powers Law*, 150 U. PA. L. REV. 603, 650–54 (2001); Nina A. Mendelson, *A Presumption Against Agency Preemption*, 102 NW. U. L. REV. 695, 717–22 (2008); Jide Nzelibe & John Yoo, *Rational War and Constitutional Design*, 115 YALE L.J. 2512, 2514–16 (2006); Jeffrey J. Rachlinski, *Bottom-Up Versus Top-Down Lawmaking*, 73 U. CHI. L. REV. 933, 934–37 (2006); Frederick Schauer, *The Dilemma of Ignorance:* PGA Tour, Inc. v. Casey Martin, 2001 SUP. CT. REV. 267, 288–95; Mark Seidenfeld, *Who Decides Who Decides: Federal Regulatory Preemption of State Tort Law*, 65 N.Y.U. ANN. SURV. AM. L. 611, 640–49 (2010); Catherine M. Sharkey, *Products Liability Preemption: An Institutional Approach*, 76 GEO. WASH. L. REV. 449, 484–90 (2008); Matthew C. Stephenson, *Public Regulation of Private Enforcement: The Case for Expanding the Role of Administrative Agencies*, 91 VA. L. REV. 93, 126–43 (2005); Cass R. Sunstein & Adrian Vermeule, *Interpretation and Institutions*, 101 MICH. L. REV. 885, 914–19 (2003); Ernest A. Young, *The Rehnquist Court's Two Federalisms*, 83 TEX. L. REV. 1, 65–69, 91–95 (2004).

[3] *See, e.g.*, Alícia Adserà et al., *Are You Being Served? Political Accountability and Quality of Government*, 19 J.L. ECON. & ORG. 445, 447–49 (2003); Scott Baker & Kimberly D. Krawiec, *The*

of scholarship investigates how legal and political institutions may encourage or inhibit the efficient aggregation of information from multiple sources.[4]

In focusing on how to encourage the most effective *use* of information (say, by allocating authority to the best-informed agents, aligning agents' policy preferences with the principal's, or ensuring efficient aggregation of information from multiple sources), the extant public law literature tends to neglect the related but distinct challenge of encouraging the efficient *acquisition* of information.[5] Many discussions

---

Penalty Default Canon, 72 GEO. WASH. L. REV. 663, 684–89 (2004); Christopher R. Berry & Jacob E. Gersen, *The Unbundled Executive*, 75 U. CHI. L. REV. 1385, 1391–98 (2008); Stephanos Bibas, *Prosecutorial Regulation Versus Prosecutorial Accountability*, 157 U. PA. L. REV. 959, 979–83, 996–97 (2009); Steven P. Croley, *Theories of Regulation: Incorporating the Administrative Process*, 98 COLUM. L. REV. 1, 22–25 (1998); Elizabeth Garrett, *The Promise and Perils of Hybrid Democracy*, 59 OKLA. L. REV. 227, 232–33 (2006); Jacob E. Gersen, *Unbundled Powers*, 96 VA. L. REV. 301, 309–15 (2010); Michael D. Gilbert, *Single Subject Rules and the Legislative Process*, 67 U. PITT. L. REV. 803, 840–41 (2006); Tom Ginsburg, *Economic Analysis and the Design of Constitutional Courts*, 3 THEORETICAL INQUIRIES L. 49, 50–55 (2002); Samuel Issacharoff & Daniel R. Ortiz, *Governing Through Intermediaries*, 85 VA. L. REV. 1627, 1635–38, 1649–52 (1999); Elena Kagan, *Presidential Administration*, 114 HARV. L. REV. 2245, 2331–39 (2001); Jonathan R. Macey, *Promoting Public-Regarding Legislation Through Statutory Interpretation: An Interest Group Model*, 86 COLUM. L. REV. 223, 244–46 (1986); David Markell, *"Slack" in the Administrative State and Its Implications for Governance: The Issue of Accountability*, 84 OR. L. REV. 1, 19–23 (2005); Jide O. Nzelibe & Matthew C. Stephenson, *Complementary Constraints: Separation of Powers, Rational Voting, and Constitutional Design*, 123 HARV. L. REV. 617, 639–43 (2010); Eric A. Posner & Adrian Vermeule, *The Credible Executive*, 74 U. CHI. L. REV. 865, 897–910 (2007); David S. Rubenstein, *"Relative Checks": Towards Optimal Control of Administrative Power*, 51 WM. & MARY L. REV. 2169, 2214–23 (2010); Matthew C. Stephenson, *Optimal Political Control of the Bureaucracy*, 107 MICH. L. REV. 53, 70 (2008).

  [4] *See, e.g.*, CASS R. SUNSTEIN, A CONSTITUTION OF MANY MINDS 7–16 (2009) [hereinafter SUNSTEIN, MANY MINDS]; CASS R. SUNSTEIN, INFOTOPIA 197–216 (2006); ADRIAN VERMEULE, LAW AND THE LIMITS OF REASON 79–95 (2009); Jacob E. Gersen, *Temporary Legislation*, 74 U. CHI. L. REV. 247, 266–78 (2007); Saul Levmore, *Conjunction and Aggregation*, 99 MICH. L. REV. 723, 734–46 (2001); Anne Joseph O'Connell, *The Architecture of Smart Intelligence: Structuring and Overseeing Agencies in the Post-9/11 World*, 94 CALIF. L. REV. 1655, 1676–78 (2006); Adrian Vermeule, *Many-Minds Arguments in Legal Theory*, 1 J. LEGAL ANALYSIS 1, 33–35 (2009) [hereinafter Vermeule, *Many-Minds Arguments*]; Adrian Vermeule, *The Parliament of the Experts*, 58 DUKE L.J. 2231, 2244–48 (2009) [hereinafter Vermeule, *Parliament*]; Adrian Vermeule, *Second Opinions* 11–19 (Harvard Law Sch. Pub. Law & Legal Theory Working Paper Series, Paper No. 10-38, 2010), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1646414.

  [5] There are, of course, important exceptions to this general statement. Indeed, a number of thoughtful contributions to the public law literature have noted ways in which legal and political institutions might influence government agents' incentives to acquire information. *See, e.g.*, WILLIAM J. STUNTZ, FIGHTING CRIME (forthcoming 2011); VERMEULE, *supra* note 4, at 57–60; Eric Biber & Berry Brosi, *Officious Intermeddlers or Citizen Experts? Petitions and Public Production of Information in Environmental Law*, 58 UCLA L. REV. 321, 325–26 (2010); Darryl K. Brown, *Cost-Benefit Analysis in Criminal Law*, 92 CALIF. L. REV. 323, 361 (2004); William W. Buzbee, *Adjudicatory Triggers of Enhanced Ambient Environment Information*, 83 IND. L.J. 583, 597–603 (2008); Kirsten Engel & Susan Rose-Ackerman, *Environmental Federalism in the United States: The Risks of Devolution*, *in* REGULATORY COMPETITION AND ECONOMIC INTEGRATION 135, 136 (Daniel C. Esty & Damien Geradin eds., 2001); Daniel C. Esty, *Environmental*

of optimal legal-institutional design take as given (implicitly or explicitly) the expertise of the various possible decisionmakers. Yet these agents often must exert costly efforts to obtain the evidence and to perform the analysis necessary to make better predictions about the consequences of different courses of action, and their incentives to do so may be shaped by the institutional environment. In the jargon of modern social science, public decisionmakers' expertise about policy decisions is often endogenous (produced by factors internal to the legal-institutional system) rather than exogenous (determined by factors external to, and therefore independent of, legal-institutional design choices).

The endogeneity of government agents' expertise may have profound consequences for a range of institutional design questions. As a general matter, agents' private incentives to invest in research may not align with the social interest in their doing so.[6] For this reason, institutional designers will often need to consider how to structure legal and institutional rules to promote not only the efficient aggregation and use of information, but also the efficient acquisition of information. Furthermore, there is often an unavoidable trade-off between inducing optimal use of information ex post and inducing optimal acquisition of information ex ante. This trade-off arises because designers often have to try to achieve both of these tasks simultaneously, with a relatively limited and crude set of mechanisms. Whenever one has to perform two tasks with one tool, it is likely that neither will be performed perfectly; if one of those tasks is ignored, it may be performed very badly indeed. Thus, failure to consider how legal institutions affect incentives to acquire information may lead to incomplete or misguided recommendations for institutional reform.

Protection in the Information Age, 79 N.Y.U. L. REV. 115, 120, 197–98 (2004); Elizabeth Garrett, *Framework Legislation and Federalism*, 83 NOTRE DAME L. REV. 1495, 1509–13 (2008); Jacob E. Gersen, *Overlapping and Underlapping Jurisdiction in Administrative Law*, 2006 SUP. CT. REV. 201, 212–14; Elizabeth Magill, Foreword, *Agency Self-Regulation*, 77 GEO. WASH. L. REV. 859, 890 (2009); O'Connell, *supra* note 4, at 1689–90; Peter H. Schuck, *Delegation and Democracy: Comments on David Schoenbrod*, 20 CARDOZO L. REV. 775, 782–83 (1999); Benjamin K. Sovacool, *The Best of Both Worlds: Environmental Federalism and the Need for Federal Action on Renewable Energy and Climate Change*, 27 STAN. ENVTL. L.J. 397, 424–25 (2008); Adrianne G. Threatt, *The Impact of Term Limits on the Congressional Committee System*, 6 GEO. MASON L. REV. 767, 781–82 (1998); Vermeule, *Many-Minds Arguments*, *supra* note 4, at 26–33; Vermeule, *Parliament*, *supra* note 4, at 2246–48; John Ferejohn, *The Lure of Large Numbers*, 123 HARV. L. REV. 1969, 1984–89 (2010) (reviewing CASS R. SUNSTEIN, A CONSTITUTION OF MANY MINDS (2009); and ADRIAN VERMEULE, LAW AND THE LIMITS OF REASON (2009)); Vermeule, *Second Opinions*, *supra* note 4, at 13–15. This Article builds on these prior contributions in the legal literature, as well as the more extensive analyses of this topic found in the political economy literature.

 [6] For a succinct and insightful early summary of this problem, which despairs of finding any real solution, see Gordon Tullock, *Public Decisions as Public Goods*, 79 J. POL. ECON. 913 (1971).

The goal of this Article is to move considerations of endogenous expertise from the periphery to the center of public law scholarship. In doing so, the Article sketches out how different institutional arrangements (arrangements that are often determined or shaped by law) might affect the production of useful information by government agents. In developing these points, the Article draws on an extensive body of literature in political economy and organization theory, a literature that has developed rapidly over the last two decades but that has not yet had much impact on legal scholarship. Rather than providing a technically rigorous survey of this literature, the Article synthesizes some of its most important concepts and findings and discusses the implications that are of greatest relevance to public law. The Article does not aspire to provide an in-depth analysis of the effects of any particular legal institution; the arguments and examples are deliberately general and abstract. Their purpose is to supply a set of principles that could be brought to bear, in conjunction with other context-specific considerations, in a range of applications.

Part I lays out a stylized theoretical framework for thinking about the challenge of inducing public decisionmakers to acquire policy-relevant information. Parts II and III then consider how various institutional arrangements might affect public decisionmakers' incentives to gather information. Part II analyzes a simple setting involving a single principal and a single agent; Part III focuses on more complex settings involving multiple agents. A brief conclusion follows.

I. THE PROBLEM OF PUBLIC UNDERINVESTMENT IN INFORMATION

Most government decisions must be made under conditions of substantial uncertainty, in which the optimal choice depends on information about consequences that can never be known with anything approaching certainty. Consider the following stylized examples, based on familiar public policy controversies:

- The President must decide whether to authorize a set of aggressive counterterrorism measures, including surveillance of electronic communications, coercive interrogation, and limits on the procedural rights of suspected terrorist detainees. In some possible states of the world, such policies are, on balance, in society's interests, given the gravity of the threat and the efficacy of the measures in question.[7] In other possible states of the world, the harm to civil liberties that these policies create outweighs their minimal security benefits; indeed, the policies might ac-

---

[7] *See, e.g.*, ERIC A. POSNER & ADRIAN VERMEULE, TERROR IN THE BALANCE 21–26 (2007); RICHARD A. POSNER, NOT A SUICIDE PACT 148–52 (2006).

tually undermine rather than aid the struggle against the terrorist threat.[8]

- Congress must vote on a health care reform bill. Whether the bill is a good idea or a bad idea for the country (or, more parochially, for each legislator's constituency) depends on the legislation's likely impact on health care costs, productivity, mortality and morbidity, inequality, and a host of other factors. There are plausible scenarios in which the bill, if enacted, would make the country better off, but also plausible scenarios in which the bill would make the country worse off.

- The Environmental Protection Agency (EPA) must decide how aggressively to respond to concerns about anthropogenic global climate change. The appropriate regulatory response depends on the severity of the risk, the efficacy of various regulatory responses, and the economic costs of these measures — all topics about which there is substantial uncertainty.[9] In addition to this scientific and economic uncertainty, there may also be political uncertainty: Will incremental administrative action prompt or inhibit a more comprehensive legislative response?[10] Will unilateral national action facilitate or undermine the prospects for negotiating a more comprehensive global treaty?[11]

- The Supreme Court must decide whether a statute that places limits on certain forms of campaign spending contravenes the constitutional guarantee of free speech. Although the question is framed as one of legal principle, the correct doctrinal answer depends on whether the government has a sufficiently compelling interest in maintaining the law.[12] This appraisal, in turn,

---

[8] *See, e.g.*, Louis Fisher, *Extraordinary Rendition: The Price of Secrecy*, 57 AM. U. L. REV. 1405, 1446 (2008); Nadine Strossen, *American Exceptionalism, the War on Terror, and the Rule of Law in the Islamic World*, 32 HARV. J.L. & PUB. POL'Y 495, 498 (2009).

[9] *See, e.g.*, Eric Biber, *Climate Change, Causation, and Delayed Harm*, 37 HOFSTRA L. REV. 975, 981–83 (2009); Daniel A. Farber, *Uncertainty* 31–41 (Univ. of Cal. Berkeley Pub. Law & Legal Theory Research Paper Series, Paper No. 1555343, 2010), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1555343.

[10] *See* Jason Scott Johnston, *Climate Change Confusion and the Supreme Court: The Misguided Regulation of Greenhouse Gas Emissions Under the Clean Air Act*, 84 NOTRE DAME L. REV. 1, 2 (2008); *cf.* J.R. DeShazo & Jody Freeman, *Timing and Form of Federal Regulation: The Case of Climate Change*, 155 U. PA. L. REV. 1499, 1500 (2007) (discussing how piecemeal adoption of state-level environmental laws may affect politics of adopting more comprehensive federal law).

[11] *See* Rachel Brewster, *Stepping Stone or Stumbling Block: Incrementalism and National Climate Change Legislation*, 28 YALE L. & POL'Y REV. 245, 247–49 (2010); John Holland, Note, *The United States and Its Climate Change Policy: Advocating an Alignment of National Interest and Ethical Obligations*, 23 NOTRE DAME J.L. ETHICS & PUB. POL'Y 623, 642–44 (2009).

[12] *See* Citizens United v. FEC, 130 S. Ct. 876, 898 (2010). For general discussions of how the legal question of the strength of the government's interest often has an important empirical and predictive dimension, see generally Richard H. Fallon, Jr., *Individual Rights and the Powers of*

> requires assessing the degree to which the restricted campaign activities increase the appearance or reality of political corruption, undermine confidence in democratic institutions, or cause other sufficiently severe social harms.[13]

All of the public decisionmakers in these examples — the President, Congress, the EPA, and the Supreme Court — are faced with decisions under uncertainty. While this sort of uncertainty is impossible to eliminate, all of these decisionmakers may have some capacity to acquire information about the likely consequences of different decisions. For example, these agents (or their subordinates) may collect or analyze additional data, study the scholarly literature, consult with knowledgeable outside parties, or perform a more systematic and comprehensive scenario analysis. Even when such activities are not feasible or appropriate, a decisionmaker might be able to achieve more accurate estimates of likely consequences by devoting more time, thought, and attention to the issues involved or by engaging in collective deliberation. In some cases, a policymaker might implement a pilot program that, while not itself cost justified, will generate useful empirical information that can then be used in the formulation of more general public policy.[14] As shorthand, we can refer to all of the foregoing ac-

---

Government, 27 GA. L. REV. 343 (1993); Stephen E. Gottlieb, *Compelling Governmental Interests: An Essential but Unanalyzed Term in Constitutional Adjudication*, 68 B.U. L. REV. 917 (1988); Kenneth L. Karst, *Legislative Facts in Constitutional Litigation*, 1960 SUP. CT. REV. 75; and Matthew C. Stephenson, *The Price of Public Action: Constitutional Doctrine and the Judicial Manipulation of Legislative Enactment Costs*, 118 YALE L.J. 2 (2008).

[13] *See, e.g.*, A USER'S GUIDE TO CAMPAIGN FINANCE REFORM (Gerald C. Lubenow ed., 2001); Stephen Ansolabehere et al., *Unrepresentative Information: The Case of Newspaper Reporting on Campaign Finance*, 69 PUB. OPINION Q. 213, 227–28 (2005); Stephen Ansolabehere et al., *Why Is There So Little Money in U.S. Politics?*, J. ECON. PERSP., Winter 2003, at 105, 125–26; Nathaniel Persily & Kelli Lammie, *Perceptions of Corruption and Campaign Finance: When Public Opinion Determines Constitutional Law*, 153 U. PA. L. REV. 119, 121–22 (2004); David M. Primo & Jeffrey Milyo, *Campaign Finance Laws and Political Efficacy: Evidence from the States*, 5 ELECTION L.J. 23, 34–36 (2006); Beth Ann Rosenson, *The Effect of Political Reform Measures on Perceptions of Corruption*, 8 ELECTION L.J. 31, 42–43 (2009).

[14] *See* Amihai Glazer & Refael Hassin, *Governmental Failures in Evaluating Programs*, 94 PUB. CHOICE 105, 107–10 (1998); Yair Listokin, *Learning Through Policy Variation*, 118 YALE L.J. 480, 491–95 (2008); *cf.* Sumon Majumdar & Sharun W. Mukand, *Policy Gambles*, 94 AM. ECON. REV. 1207 (2004) (modeling political incentives for government policy experimentation). Such information-gathering initiatives might also take the form of randomized experiments in which a range of policies (including some that are believed, ex ante, to be suboptimal) are implemented for different (randomly selected) subgroups, in the hopes of generating more and better information. *See, e.g.*, Michael Abramowicz et al., *Randomizing Law* (Yale Law Sch. Faculty Scholarship Series, Paper No. 161, 2010), *available at* http://digitalcommons.law.yale.edu/fss_papers/161; Esther Duflo et al., *Using Randomization in Development Economics Research: A Toolkit* (Ctr. for Int'l Dev. at Harvard Univ., Working Paper No. 138, 2006), *available at* http://www.hks.harvard.edu/centers/cid/publications/faculty-working-papers/cid-working-paper-no.-138; D. James Greiner & Cassandra Wolos Pattanayak, What Difference Representation? (Jan. 12, 2011) (unpublished manuscript), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1708664. Experimental initiatives along these lines would qualify as "costly

tivities collectively as "research," even though some of them might not resemble research in a conventional sense.[15]

The main benefit of such research activities is that additional information may lead to better decisions. Research, however, is costly: It requires a decisionmaker (or her staff) to devote time, resources, and mental effort to studying a particular issue rather than to something else. Research may also entail delay — "paralysis by analysis" — that is costly as soon as it appears there is at least one alternative policy better than the status quo. For these reasons, there is a limit both to how much research is socially desirable and to how much research one can expect a public decisionmaker to undertake. We hope and expect that legislators will study the likely consequences of health care reform before voting on a bill; we do not hope, and should not expect, that legislators will spend all of their time, and the entirety of the gross domestic product, researching this issue. We would like the EPA to devote substantial effort to studying the global climate change problem before deciding whether to adopt aggressive new restrictions on carbon dioxide emissions, but once we are sufficiently confident that *something* needs to be done, then after a certain point it is no longer worthwhile for the EPA to continue to study the issue in order to develop an even better regulatory response. As a general matter, we would like our public decisionmakers to invest in research up to the point where the marginal social benefit of additional research (in the form of improved policy decisions) is equal to the marginal social cost (typically the opportunity costs associated with the diversion of resources and delay).[16]

If government decisionmakers are (approximately) rational, we can expect that they will choose a level of research effort that (approximately) equates their marginal benefits with their marginal costs.[17] In many cases, however, the marginal *social* costs and benefits of research investment will not align perfectly with the relevant government agent's *private* marginal costs and benefits. Such misalignment leads

research" as this Article uses that term, as proceeding in this fashion entails costs (both the expected cost of implementing what appears to be a suboptimal policy with respect to some populations and the expected costs of implementing and monitoring the experiment), but might yield valuable information that could lead to more informed policy decisions in the future.

[15] These examples, and most of the discussion in this Article, involve situations where research might reduce *empirical* uncertainty. A similar logic would apply to moral or legal reasoning, if one postulates that the soundness of a decisionmaker's moral or legal conclusion is positively correlated with the amount of effort the decisionmaker invests in considering the issue.

[16] Of course, taking action does not preclude additional research; indeed, taking action will often be useful in generating further useful information. *See supra* note 14. The statement in the text is therefore a bit of an oversimplification, but the same basic logic would apply to more complex decisions that involve ongoing revision and learning-by-doing.

[17] *See, e.g.*, Jonathan Baert Wiener, *Managing the Iatrogenic Risks of Risk Management*, 9 RISK: HEALTH, SAFETY & ENV'T 39, 73 (1998).

to socially suboptimal investment in information. Although it is theoretically possible for a public decisionmaker's research effort to be too high or too low, there are reasons to think that in most cases it will be too low. Most importantly, a typical government agent internalizes only a fraction of the aggregate social benefit associated with making a better public policy decision, but she internalizes the lion's share of the research costs. In contrast, although some of the costs associated with additional research are social costs, much of the research cost to the agent derives from things like forgone leisure, greater mental effort, and other private costs that are socially trivial.[18] In his classic formulation of this problem, Professor Gordon Tullock used the example of a judge deciding a case. Tullock reasoned that, if the judge is confronted with a difficult legal question,

> [h]e can produce a quick solution to the problem without much thought. If, however, he wants to be sure that he makes the "correct" decision, he must devote a great deal of time and thought to it. This is a private cost, and the decision will primarily produce public goods. Ordinary public-goods reasoning would imply that he would underinvest in this private expenditure to obtain the public good of a superior decision.[19]

This problem is likely to exist even when the relevant agents are dedicated public servants who care deeply about making good decisions. As a *relative* matter, the cumulative social utility from making even a slightly better decision on a matter like health care, climate change, counterterrorism, or campaign finance likely dwarfs the private utility that accrues to the responsible decisionmaker.[20] An additional reason why pre-decision research investment is likely to be too low (explored in greater detail in Part III) is that when multiple agents are responsible for researching a policy decision, the acquisition of policy-relevant information may be subject to a collective action problem.

---

[18] Furthermore, research activities compete for time and attention with an agent's other tasks, and many forms of research may be less rewarding to the agent, as a relative matter, than these other tasks, even if the social value of research is large. *See* Eric Biber, *Environmental Law's Monitoring Problem* (Univ. of Cal. Berkeley Pub. Law & Legal Theory Research Paper Series, Paper No. 1680000, 2010), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1680000.

[19] Tullock, *supra* note 6, at 915 (footnote omitted).

[20] *See id.* In this issue, Professor Michael Abramowicz proffers the intriguing suggestion that it might be possible to allow innovative policymakers to internalize more of the benefits of their innovations by, for example, granting a state government that implements a successful new program "exclusive rights" to that policy innovation, and perhaps requiring other states wishing to follow the model to pay a "licensing fee." *See* Michael Abramowicz, *Orphan Business Models: Toward a New Form of Intellectual Property*, 124 HARV. L. REV. 1362 (2011). Abramowicz, however, does not flesh out this possibility (which is peripheral to the main point of his article), and it would seem to face formidable practical difficulties. The more general idea of leveraging competition among different policymaking entities (for example, states) is explored in section III.C, pp. 1479–82.

Thus, a government agent's private marginal benefit from additional research may often be systematically lower than the social marginal benefit of such research, or the agent's private marginal research cost may be systematically higher than the social marginal cost of that research.[21] If this principal-agent problem is a real and serious one — one that leads to systematically worse decisions of law and policy than we would prefer — what might we do about it? The natural answer is to find ways to reduce the relevant government agents' marginal research costs, to increase their marginal research benefits, or both.

The most straightforward way a social planner could try to align an agent's private marginal research costs and benefits with social marginal costs and benefits would be to offer the agent some form of compensation (such as a higher wage) for performing additional research (or for making a higher-quality decision). Alternatively, one could threaten the agent with some kind of direct punishment (a fine or termination of employment) if the agent invests too little in research or if the final outcome is adjudged a failure. Indeed, much of the existing economics literature on the analogous problem in the private firm context focuses on this sort of contingent compensation mechan-

---

[21] Although this Article focuses on the problem of systematic *under*investment in information, which is likely to be more prevalent for the reasons sketched in the text, *see supra* pp. 1430–31, a structurally similar problem may lead to systematic *over*investment in information. This latter problem may take the form of the "paralysis by analysis" problem noted earlier, *see supra* p. 1430, in which extra study is socially costly because it preserves a socially inefficient status quo for a longer time. Even if there is no delay problem, there may be situations in which a government agent diverts too many resources from other socially useful tasks in order to research a particular problem. One would expect to observe this sort of overinvestment when the government agent's private marginal benefit of doing additional research is higher than the marginal social benefit (as when the agent has an excessive stake in the success of one particular project) or when the agent's private marginal cost of research is lower than the marginal social cost (as when the agent pays for additional research by diverting resources from activities that are more valuable to society than to the agent). *See* Matthew C. Stephenson, *Bureaucratic Decision Costs and Endogenous Agency Expertise* 18–19 (Harvard Univ. John M. Olin Ctr. for Law, Econ. & Bus., Discussion Paper No. 553, 2006), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=921439.

Furthermore, if decisionmakers suffer from certain forms of cognitive bias, acquisition of more information may actually *worsen* decisionmaking by, for example, fostering overconfidence in potentially false beliefs. *See* Charles G. Lord et al., *Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence*, 37 J. PERSONALITY & SOC. PSYCHOL. 2098, 2105–08 (1979); Matthew Rabin & Joel L. Schrag, *First Impressions Matter: A Model of Confirmatory Bias*, 114 Q.J. ECON. 37, 38–39 (1999); *see also* Edward L. Glaeser & Cass R. Sunstein, *Extremism and Social Learning*, 1 J. LEGAL ANALYSIS 263, 263–64 (2009) (arguing that deliberation can sometimes cause group polarization, which can lead to more rather than fewer erroneous decisions); Cass R. Sunstein, *Group Judgments: Statistical Means, Deliberation, and Information Markets*, 80 N.Y.U. L. REV. 962, 1004–06 (2005) (same).

In settings where overinvestment in information is the principal concern, the analysis in the remainder of this Article would still be relevant, but the normative conclusions would be inverted, as the social objective in this mirror-image case would be to reduce, rather than increase, research incentives.

ism.[22]   Even in contexts where literal contracts are not available —
such as the relationship between politicians and voters — scholars
have suggested that something like this mechanism (in particular, the
threat of ex post punishment for policy failure) may help redress the
inherent agency problem in political representation.[23] This same basic
logic might induce political agents to acquire more and better informa-
tion before taking risky action.

   There are, however, several difficulties with this solution in the
context of public decisionmaking. First, it is often impossible for the
principal to observe (or to infer with sufficient accuracy) how much a
given agent has invested in research, and it is often similarly difficult
to observe the quality of a policy decision after the fact, especially for
policies that have long-term consequences.[24] Second, it may be diffi-
cult for the principal to credibly commit to follow through on its prom-
ise to pay a bonus for acquiring information or to impose a penalty for
failing to do so. This problem is exacerbated by the difficulty of ob-
serving research effort and the quality of policy outcomes, which
makes it more difficult to assess whether the principal has reneged.
Third, if the principal is society in general, it may be infeasible for so-
ciety, acting collectively, to agree to an employment contract with gov-
ernment officials that implements this sort of refined incentive scheme.
Fourth, established features of our current public institutions — in-
cluding the rules for compensating judges, civil servants, and senior
legislative and executive officials — may constrain our ability to offer
these sorts of contingent rewards.[25] Although these institutional rules

---

[22] *See generally, e.g.*, PATRICK BOLTON & MATHIAS DEWATRIPONT, CONTRACT THEORY
(2005); BERNARD SALANIÉ, THE ECONOMICS OF CONTRACTS (1997). For specific applica-
tions of this optimal contracting approach to investment in information (as opposed to other forms
of effort), see, for example, Tim S. Campbell et al., *Incentives for Information Production and
Optimal Job Assignment with Human Capital Considerations*, 60 ECONOMICA 13 (1993); Joel S.
Demski & David E.M. Sappington, *Delegated Expertise*, 25 J. ACCT. RES. 68 (1987); Joel S.
Demski & David E.M. Sappington, *Hierarchical Regulatory Control*, 18 RAND J. ECON. 369
(1987); Richard A. Lambert, *Executive Effort and Selection of Risky Projects*, 17 RAND J. ECON.
77 (1986); Tracy R. Lewis & David E.M. Sappington, *Information Management in Incentive Prob-
lems*, 105 J. POL. ECON. 796 (1997); Todd T. Milbourn et al., *Managerial Career Concerns and
Investments in Information*, 32 RAND J. ECON. 334 (2001); and Dezsö Szalay, *Contracts with
Endogenous Information*, 65 GAMES & ECON. BEHAV. 586 (2009).

[23] *See* David Austen-Smith & Jeffrey Banks, *Electoral Accountability and Incumbency*, *in*
MODELS OF STRATEGIC CHOICE IN POLITICS 121, 123–24 (Peter C. Ordeshook ed., 1989);
John Ferejohn, *Incumbent Performance and Electoral Control*, 50 PUB. CHOICE 5, 7 (1986);
Adam Meirowitz, *Probabilistic Voting and Accountability in Elections with Uncertain Policy
Constraints*, 9 J. PUB. ECON. THEORY 41, 42–43 (2007); Nzelibe & Stephenson, *supra* note 3, at
627–36; Matthew C. Stephenson & Jide O. Nzelibe, *Political Accountability Under Alternative
Institutional Regimes*, 22 J. THEORETICAL POL. 139, 143 (2010).

[24] *See* Canice Prendergast, *A Theory of "Yes Men,"* 83 AM. ECON. REV. 757, 757–58 (1993).

[25] *See* Dezsö Szalay, *The Economics of Clear Advice and Extreme Options*, 72 REV. ECON.
STUD. 1173, 1173 (2005).

are themselves changeable, they may be costly to change and may serve other desirable functions, and so they might have to be treated as exogenous constraints. For these and other reasons, a contingent compensation contract may not work all that well for a range of important public decisionmaking contexts, and this Article will put this sort of mechanism to one side.

If contingent compensation contracts are not available, a principal might nonetheless decrease an agent's marginal research cost simply by supplying the agent with more resources (for example, a higher budget or more staff). Doing so will tend to increase the resources that the agent is willing to allocate across a range of tasks, including research.[26] However, simply increasing resources is unlikely to be a complete solution to the problem. Providing any given agent with more resources is both socially costly and constrained by other political or institutional factors. In addition, resources alone do not address the problem of incentive misalignment; indeed, while providing more resources will increase investment in socially valuable activities, such as research, some of these resources will be diverted to activities with lower social value but greater private value to the agent.[27] It is therefore unlikely that increasing resources, even if socially desirable, would entirely solve the underinvestment problem.

Another way that an institutional designer might increase a government agent's marginal research benefit, or reduce her marginal research cost, is through the selection of agents. Indeed, perhaps the most obvious way to foster a more expert government is to staff the government with smarter people.[28] Additionally, a principal can try to select "public-spirited" agents who care intensely about policy outcomes, as such agents' private marginal benefit from additional research will tend to be higher.

The above observations are familiar (though they are not always framed in the context of incentivizing the acquisition of information) and so this Article will not pursue them further, except to note that attempting to use selection to redress the problem of underinvestment in

---

[26] *See, e.g.*, Michael M. Ting, *The "Power of the Purse" and Its Implications for Bureaucratic Policy-Making*, 106 PUB. CHOICE 243, 245 (2001). For example, some evidence indicates that sizeable increases in law enforcement budgets in the early 1970s and again in the early 1990s led to greater increases in prosecutions than in arrests, which is at least suggestive evidence that larger budgets led to more information gathering, which in turn led to higher-quality law enforcement. *See* STUNTZ, *supra* note 5.

[27] *See* Ting, *supra* note 26, at 243–44.

[28] "Smart" here is a casual way to characterize agents with low marginal research costs. Of course, an agent's marginal research cost is only one aspect of an agent's qualifications and must be balanced against others. *See* Anne Joseph O'Connell, Qualifications: Law and Practice of Selecting Agency Leaders 3, 20–22 (Jan. 3, 2011) (unpublished manuscript) (on file with the Harvard Law School Library).

research may present some difficult trade-offs. First, the most public-spirited agents may not always be the smartest, and vice versa. Even if one is focused narrowly on incentivizing research, it is important to keep in mind that the most important consideration is the *ratio* of marginal research costs to marginal research benefits; considering either in isolation may be misleading. Second, there is the well-known trade-off between competence and loyalty — a trade-off the Chinese Communists colorfully described as the "red-versus-expert" problem.[29] Often the agent with the lowest marginal research costs also has preferences that differ substantially from the principal's, while the agent with policy preferences most similar to the principal's faces relatively high marginal research costs.

The analysis in Parts II and III focuses on how legal and institutional rules can affect government agents' incentives to invest in expertise through mechanisms other than contingent performance contracts and selection mechanisms. The Article also brackets the potentially important role of institutional reforms that improve the overall efficiency of government, thereby lowering the marginal research costs.[30] The remainder of the analysis will treat each government agent's marginal research costs as fixed, focusing instead on how different legal and institutional design choices might affect a government agent's research incentives through their effects on the agent's marginal research benefit.

It is analytically useful to decompose the agent's marginal research benefit into two components, each of which corresponds to a question that a rational agent would ask herself at the moment she must decide whether to invest an additional unit of effort in attempting to acquire decision-relevant information.

First, the agent will assess her expected utility if she does not acquire additional information. That is, she will ask herself: "If I don't invest in trying to learn anything else, what will the final decision be, and what is my expected utility from that decision?" Call this the agent's (expected) "default payoff." The default payoff may be the known or expected payoff from a particular decision, or it may reflect the agent's subjective beliefs about the probability of different outcomes if the agent does not do any additional research.

---

[29] *See* Martin K. Whyte, *"Red vs. Expert": Peking's Changing Policy*, PROBS. COMMUNISM, Nov.–Dec. 1972, at 18, 18.

[30] For example, this Article does not deal with mechanisms like petition systems, private rights of action, prediction markets, or other tools that might lower the costs to government agencies of gathering widely dispersed information. *See* MICHAEL ABRAMOWICZ, PREDICTOCRACY 162–93 (2007) (prediction markets); Biber & Brosi, *supra* note 5, at 327–31 (petition systems); Stephenson, *supra* note 2, at 98–121 (private rights of action). The Article's analysis would, however, apply to the government's decision to create such systems in the first place, as such a decision might itself turn on costly information about its likely effects on the quality of future public decisions.

Second, the agent will estimate her expected utility if she does invest the additional unit of effort in research. In other words, the agent will ask herself: "If I invest in research, what will my expected utility be from the (possibly) better-informed decision that will result?" Call this the agent's (expected) "research payoff." An agent's calculation of her research payoff entails her making subjective probability assessments of what she might learn from additional research, as well as what policy consequences would flow from different possible pieces of new information.[31]

---

[31] This formulation posits that government decisionmakers maximize their subjective expected utility (SEU): given their personal subjective (and perhaps implicit or subconscious) beliefs about the probabilities of different outcomes (which may not be correct), agents will act to maximize their expected utility, given their values. *See* LEONARD J. SAVAGE, THE FOUNDATIONS OF STATISTICS 6–9 (1954). This is a standard axiom in most modern political economy. That said, there are a number of objections to the SEU framework. While a full treatment of these issues is well beyond the scope of this Article, two are particularly pertinent to the application of SEU theory to investment in information.

First, some economists, following the work of Professor Frank Knight, have suggested a distinction between "risk" (situations in which different possible outcomes can be assigned numerical probabilities a priori) and "uncertainty" (situations in which even the probabilities of different outcomes are themselves unknown). *See generally* FRANK H. KNIGHT, RISK, UNCERTAINTY AND PROFIT 197–203 (1921). In one interpretation of Knight's distinction, agents cannot act as SEU maximizers under conditions of uncertainty, because it is impossible for them to calculate the probabilities of different outcomes (or even to identify what those outcomes are); they must use some other decision procedure. One might assert that government agents trying to decide how much to invest in research face a situation of uncertainty rather than risk. This is especially so if the range of possible policy choices is unknown ex ante, such that research might uncover entirely new options that the decisionmaker does not even know about when she makes her research decision.

However, while it is often true that one cannot confidently assign precise probability estimates to particular outcomes on the basis of objective a priori or statistical considerations, SEU theory does not require that sort of precision; it requires only that agents behave as if they assigned probability distributions over outcomes. *See, e.g.*, MILTON FRIEDMAN, PRICE THEORY 282 (1976). (Indeed, Professors Stephen LeRoy and Larry Singell have persuasively argued that Knight himself never intended his concept of uncertainty to entail a rejection of subjective probability, but rather meant to characterize settings in which insurance markets cannot function due to adverse selection and moral hazard problems. *See* Stephen F. LeRoy & Larry D. Singell, Jr., *Knight on Risk and Uncertainty*, 95 J. POL. ECON. 394, 394–96 (1987).) Admittedly, an important body of empirical evidence indicates that the SEU assumption is sometimes violated — in particular, that individuals exhibit "uncertainty" or "ambiguity" aversion, in addition to risk aversion. *See, e.g.*, Daniel Ellsberg, *Risk, Ambiguity, and the Savage Axioms*, 75 Q.J. ECON. 643, 643 (1961). But recent work has suggested that ambiguity avoidance may in fact be consistent with SEU theory if agents can take unobservable actions that affect the probabilities of different outcomes. *See* David Kelsey & Frank Milne, *Induced Preferences, Nonadditive Beliefs, and Multiple Priors*, 40 INT'L ECON. REV. 455, 457 (1999). Moreover, although genuine ambiguity aversion would imply a higher level of research investment than SEU theory would predict (all else equal), it would likely not affect the comparative statics predicted by SEU theory. That said, for readers who find the risk-uncertainty distinction compelling and who doubt that agents can act as SEU maximizers under conditions of pure uncertainty, it is worth emphasizing that many government decisions of interest do involve conditions of risk (or a mix of risk and uncertainty). Thus, the Knightian objection would at most limit the scope of the analysis developed in this Article.

The agent's marginal research benefit from investing an additional unit of effort in research is simply the difference between the research payoff and the default payoff. Thus, the strength of the agent's research incentive is a decreasing function of her default payoff and an increasing function of her research payoff. That conclusion is both quite straightforward — perhaps obvious — and quite abstract. The balance of this Article seeks to flesh out this simple observation by exploring how a range of institutional design choices are likely to affect an agent's default payoff and research payoff. Although these institutional structures are presented in stylized form, they correspond to familiar features of real-world decisionmaking structures used by legislatures, courts, and executive branch agencies and officials.

Before proceeding, however, it is worth pausing to consider a potential objection to the whole enterprise. The analysis in this Article presumes that both the principal and the agent care about information, in the sense that it is possible that new information, acquired via costly research effort, might alter their (induced) policy preferences. Yet there may be many issues — possibly including a number of the high-profile political issues discussed above — for which additional information is unlikely to change actors' minds (except perhaps by intensifying their prior beliefs).[32] This observation is unlikely to be true for

---

Second, a related concern, more specific to the context of information gathering, is a kind of infinite regress problem: deciding how much to invest in research requires assessing the costs and benefits of an additional unit of research, but these costs and benefits are themselves uncertain; the agent may learn more about the potential value of research (the range of possible results, the likely costs, and so forth), but learning this information is also costly; this fact implies that the agent must make a decision regarding how much to invest in researching how much to research; this second-order research problem raises the same issue; and so on. *See* John Conlisk, *Why Bounded Rationality?*, 34 J. ECON. LITERATURE 669, 686–88 (1996); Barton L. Lipman, *How to Decide How to Decide How to . . .: Modeling Limited Rationality*, 59 ECONOMETRICA 1105, 1106 (1991); Wiener, *supra* note 17, at 73–74. While this infinite regress problem has important philosophical and technical ramifications, it is not a serious difficulty with the analysis in this Article. The suggestion that the infinite regress problem means that it is impossible for an agent to optimize, *see, e.g.*, *id.* at 74, is not necessarily true; under some reasonable assumptions, one can still characterize an agent's behavior as the solution to an optimization problem, despite the regress, *see generally* Lipman, *supra*.

Moreover, despite the conceptual challenge posed by the infinite regress problem, from a practical modeling standpoint it makes sense to bracket the infinite regression and to consider only one research decision and the final decision. This approach would be equivalent to assuming that the agent's information at the decision stage is endogenous (determined at the research stage), but her information at the prior research stage is exogenous. That is a plausible simplifying assumption to make for an applied analysis, and one that corresponds to much observed behavior. *See* Conlisk, *supra*, at 688.

[32] *See, e.g.*, Glaeser & Sunstein, *supra* note 21, at 272–73; Dan M. Kahan & Donald Braman, *Cultural Cognition and Public Policy*, 24 YALE L. & POL'Y REV. 149, 156 (2006); Dan M. Kahan & Donald Braman, *More Statistics, Less Persuasion: A Cultural Theory of Gun-Risk Perceptions*, 151 U. PA. L. REV. 1291, 1317–18 (2003); Rabin & Schrag, *supra* note 21, at 41–47; James Andreoni & Tymofiy Mylovanov, Diverging Opinions 26–27 (May 7, 2010) (unpublished manuscript) (on file with the Harvard Law School Library).

all issues, however. Moreover, this Article's analysis would still be pertinent so long as the relevant decisionmakers cared *somewhat* about additional information; the analysis assumes neither that decisionmakers process information optimally nor that they lack strong prior policy views. There may be a subset of cases in which actors' preferences truly are insensitive to additional information. Such cases are analytically equivalent to the setting in which all decision-relevant information is exogenous rather than endogenous. In such settings the considerations in this Article would not apply.[33] But as long as it is possible, at least under some circumstances, for better information to lead to different (and better) policy decisions, then the analysis laid out in the balance of this Article will be relevant.

II. RESEARCH INCENTIVES OF A SINGLE AGENT

In exploring the myriad ways that legal and institutional structures may affect public decisionmakers' research incentives, it is helpful to begin with a simple setting involving a single principal and a single agent. The principal might be a particular institution, such as a legislative assembly, a chief executive, or an appellate court, with the agent a delegate or subordinate of that institution, such as a legislative committee, an administrative agency, or a trial court. Alternatively, one could think of the principal as society in general and of the agent as a government body that is supposed to serve the public interest.[34] The following discussion explores how the principal's institutional design choices can have a powerful effect on the agent's incentives to acquire information and how recognition of this fact may alter conventional conclusions about optimal institutional design. Part II focuses on the

---

[33] Perhaps, as noted earlier, there are also cases where an actor's information-processing ability is not only imperfect, but so distorted that additional information actually worsens the quality of her decisions (from her own perspective), even though the actor herself does not realize it. *See supra* note 21. Again, in such cases this Article's analysis would still be relevant, but the normative conclusions would be reversed, since the social objective would be to reduce rather than increase the agent's research investment.

[34] Although the positive analysis developed in this Article applies even if the "principal" is a real-world government agent or institution, rather than "society," it is important to emphasize that a normative evaluation of different "solutions" to the principal-agent problem described below may depend critically on an assessment of whether the principal's inability to control the agent is in fact a problem. Under some circumstances — in particular, those where the principal has (or may have) the wrong preferences from a social point of view — one may not want to make the agent more responsive to the principal's interest. *Cf.* Stephenson, *supra* note 3, at 55 (developing a model in which constraining the ability of an elected political leader to control subordinates is socially optimal). Normative interpretations of the positive results discussed below must be sensitive to this consideration. Although this Article is motivated by underlying normative concerns about government performance, the principal objective is to elucidate how various institutional mechanisms might influence an agent's incentive to invest in research — a positive, rather than normative, inquiry.

incentive effects of three familiar ways in which institutional designers use legal and other mechanisms to shape decisionmaking by public agents: ex ante substantive and procedural constraints on the agent's discretionary authority, evidentiary standards and requirements, and ex post review of the agent's policy choices.[35]

A. *Ex Ante Substantive and Procedural Constraints*

Often, government authorities and institutional designers delegate authority to an agent but at the same time set up a structure that constrains and channels how the agent exercises her discretion. This control strategy — which seeks to make use of the agent's policy expertise while mitigating concerns about the agent's policy bias — is particularly relevant to public law, which is sometimes thought to be one of the principal mechanisms that implement these ex ante control structures.[36] For example, constitutional rules are thought to both establish and constrain the discretion of the various branches of government. Likewise, statutes that delegate to other agents — such as administra-

---

[35] While the discussion within each section provisionally assumes that a particular institutional design choice is feasible, in some contexts some of those choices might *not* be feasible. For this reason, the discussions in each of the sections are, for the most part, self-contained. Each section notes, where appropriate, the assumptions being made about the sorts of mechanisms that the principal is and is not able to employ. Sometimes the discussion explicitly notes which of two possible control strategies is better for the principal, if both are available, but outside these specific instances the analysis makes no general claims about which mechanisms are superior.

[36] *See* Kathleen Bawn, *Choosing Strategies to Control the Bureaucracy: Statutory Constraints, Oversight, and the Committee System*, 13 J.L. ECON. & ORG. 101, 103–05 (1997); Mathew D. McCubbins et al., *Administrative Procedures as Instruments of Political Control*, 3 J.L. ECON. & ORG. 243, 246–48 (1987) [hereinafter McCubbins et al., *Administrative Procedures*]; Mathew D. McCubbins et al., *Structure and Process, Politics and Policy: Administrative Arrangements and the Political Control of Agencies*, 75 VA. L. REV. 431, 440–44 (1989) [hereinafter McCubbins et al., *Structure and Process*]; Richard J. Pierce, Jr., *The Role of the Judiciary in Implementing an Agency Theory of Government*, 64 N.Y.U. L. REV. 1239, 1248–50 (1989); Mark Seidenfeld, *Bending the Rules: Flexible Regulation and Constraints on Agency Discretion*, 51 ADMIN. L. REV. 429, 479–87 (1999); Richard B. Stewart, *The Reformation of American Administrative Law*, 88 HARV. L. REV. 1667, 1699–1701 (1975).

It is worth observing that a system that imposes ex ante constraints on an agent's discretion often requires the principal to make a credible commitment to these institutions. Such precommitment, however, is problematic: if the principal has an ex post incentive to allow an agent to make an ostensibly prohibited decision or to block a supposedly permitted choice, then it is unclear how the system of ex ante constraints is enforceable. Nonetheless, commitment appears to be possible at least some of the time, even in settings without an outside enforcer. *See* Nzelibe & Stephenson, *supra* note 3, at 648–49; Daniel F. Spulber & David Besanko, *Delegation, Commitment, and the Regulatory Mandate*, 8 J.L. ECON. & ORG. 126, 146–50 (1992); *cf.* Daniel Diermeier, *Commitment, Deference, and Legislative Institutions*, 89 AM. POL. SCI. REV. 344, 344–45 (1995) (discussing the control of committee specialization exercised by the legislative floor majority despite a commitment problem). The assumption that precommitment is possible would, of course, eliminate one of the reasons why the principal could not use the more efficient mechanism of paying the agent directly for doing research, but other factors might continue to make such direct contingent compensation arrangements infeasible. *See supra* pp. 1433–34.

tive agencies or courts — often contain substantive or procedural restrictions that are designed to limit or condition the agents' autonomy.

The simplest form that this sort of ex ante control strategy may take is the advance specification, by the principal, of the set of actions that an agent is allowed to select and the corresponding set of prohibited actions. The principal, in other words, may use legal rules to establish the agent's "discretionary window."[37] When deciding how much discretion to delegate (that is, the size and location of the discretionary window), the principal must weigh the potential informational gains of delegation against the costs associated with potential agency bias. The extensive literature on this basic delegation problem has generated a number of important (if straightforward) hypotheses. The first is that a principal is willing to delegate more discretion to an agent with expected policy preferences similar to the principal's own — a hypothesis sometimes referred to as the "ally principle."[38] A natural corollary is that if the principal has the power to influence the agent's policy preferences, the principal will select an agent with preferences closer to the principal's own. A related corollary is that if the principal is not able to select a perfect ally as its agent, it will locate the discretionary window to (partially) offset the agent's predicted bias.

A second prominent hypothesis — sometimes dubbed the "uncertainty principle" — is that a principal will delegate more discretion to the agent (that is, it will expand the size of the discretionary window) when the link between policies and outcomes is less certain ex ante.[39] This uncertainty principle may be thought of as a special case of a more general "expertise principle," according to which the principal's willingness to delegate increases as the agent's expected informational advantage increases.[40] This hypothesis fits comfortably with one of

---

[37] *See* DAVID EPSTEIN & SHARYN O'HALLORAN, DELEGATING POWERS 60 (1999); Sean Gailmard, *Discretion Rather than Rules: Choice of Instruments to Control Bureaucratic Policy Making*, 17 POL. ANALYSIS 25, 25–42 (2009); Sean Gailmard & John W. Patty, *Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise*, 51 AM. J. POL. SCI. 873, 882–84 (2007). This research builds on the seminal model developed by Bengt Holmström in the general context of organizational design. *See* Bengt Holmström, *On the Theory of Delegation, in* BAYESIAN MODELS IN ECONOMIC THEORY 115 (Marcel Boyer & Richard E. Kihlstrom eds., 1984).

[38] *See* Jonathan Bendor et al., *Theories of Delegation*, 4 ANN. REV. POL. SCI. 235, 236 (2001); Jonathan Bendor & Adam Meirowitz, *Spatial Models of Delegation*, 98 AM. POL. SCI. REV. 293, 293 (2004); Frederick J. Boehmke et al., *Whose Ear to Bend? Information Sources and Venue Choice in Policy-Making*, 1 Q.J. POL. SCI. 139, 140–41 (2005); John D. Huber & Charles R. Shipan, *Politics, Delegation, and Bureaucracy, in* THE OXFORD HANDBOOK OF POLITICAL ECONOMY 256, 260 (Barry R. Weingast & Donald A. Wittman eds., 2006).

[39] *See* Bendor et al., *supra* note 38, at 242, 248; Bendor & Meirowitz, *supra* note 38, at 293; Huber & Shipan, *supra* note 38, at 260.

[40] *See* Matthew C. Stephenson, *Statutory Interpretation by Agencies, in* RESEARCH HANDBOOK ON PUBLIC CHOICE AND PUBLIC LAW 285, 288 (Daniel A. Farber & Anne Joseph O'Connell eds., 2010). The underlying logic of the expertise principle is that delegation is more attractive when the agent has better information about the connection between policy and out-

the classic explanations (and justifications) for the growth of the administrative state: the bureaucracy's superior expertise, especially on complex technical matters, is a key factor that leads Congress to delegate broad authority to agencies.[41]

The specification of a discretionary window is a relatively crude control strategy, as it entails delegating unconstrained discretion within a range and totally prohibiting anything outside that range. A principal might employ a more nuanced version of essentially the same control strategy by establishing a system of variable rewards and penalties that makes some policy choices relatively more or less attractive.[42] For example, the legislature might empower an agency to set a regulatory standard, but impose procedural requirements that make increasing the stringency of the regulatory standard progressively more difficult (and therefore more costly) for the agency.[43] The legislature might achieve a similar result by subjecting the agency's decision to review by a third-party overseer who demands increasingly onerous and costly justifications as the agency's proposed policy deviates further from the status quo (or some favored policy specified by the principal).[44] Or the legislature might make additional resources available

---

comes. The agent's expertise advantage will increase as the principal becomes more uncertain about the likely effects of different policies, holding the agency's (un)certainty fixed; this is the traditional uncertainty principle modeled in much of the early literature on the political economy of delegation. *See, e.g.*, EPSTEIN & O'HALLORAN, *supra* note 37, at 17–18. Likewise, the attractiveness of delegation increases if one holds the principal's uncertainty about the connection between policies and outcomes fixed but increases the agent's certainty about this connection. Because what matters is the *relative* certainty of the principal and the agent about the policy-outcome connection (which one can think of as their relative "expertise"), rather than some general level of uncertainty as such, the expertise principle seems a more accurate characterization of the basic argument.

[41] *See generally* JAMES M. LANDIS, THE ADMINISTRATIVE PROCESS (1938); David B. Spence & Frank Cross, *A Public Choice Case for the Administrative State*, 89 GEO. L.J. 97 (2000).

[42] *See* Gailmard, *supra* note 37, at 26–28; Matthew C. Stephenson, *Bureaucratic Decision Costs and Endogenous Agency Expertise*, 23 J.L. ECON. & ORG. 469, 469–73 (2007); Stephenson, *supra* note 12, at 4.

[43] *See* McCubbins et al., *Administrative Procedures*, *supra* note 36, at 248–53; McCubbins et al., *Structure and Process*, *supra* note 36, at 431; David B. Spence, *Managing Delegation Ex Ante: Using Law to Steer Administrative Agencies*, 28 J. LEGAL STUD. 413, 417–18 (1999); Emerson H. Tiller, *Controlling Policy by Controlling Process: Judicial Influence on Regulatory Decision Making*, 14 J.L. ECON. & ORG. 114, 119–24 (1998); *cf.* Matthew C. Stephenson, *The Strategic Substitution Effect: Textual Plausibility, Procedural Formality, and Judicial Review of Agency Statutory Interpretations*, 120 HARV. L. REV. 528, 553–57 (2006) (arguing that courts are likely to be more deferential to administrative agencies' statutory interpretations when the agencies have employed costly formal procedures).

[44] *See* Matthew C. Stephenson, *A Costly Signaling Theory of "Hard Look" Judicial Review*, 58 ADMIN. L. REV. 753 (2006); Stephenson, *supra* note 12; Stephenson, *supra* note 43; *cf.* Ethan Bueno de Mesquita & Matthew C. Stephenson, *Regulatory Quality Under Imperfect Oversight*, 101 AM. POL. SCI. REV. 605 (2007) (discussing oversight of agencies by courts, the Office of Information and Regulatory Affairs, and other institutions).

to the agency only if the agency chooses certain actions.[45] One might classify these approaches collectively as the manipulation, by the principal, of the agent's relative "enactment costs" for different policies.[46]

This more nuanced "enactment cost manipulation" strategy (also sometimes referred to as the use of a "menu law"[47]) is always at least as good for the principal as is fixing a discretionary window, and usually better.[48] Of course, this approach may not always be possible: it may be more difficult to establish and enforce variable rewards and penalties than it is simply to set a range of permissible choices. But if the variable enactment cost approach is feasible, then the principal could conceivably eliminate (or at least substantially mitigate) the agency problem by using these rewards and punishments to align the agent's policy preferences with the principal's own.

While both of these control strategies have attracted a great deal of scholarly attention (both in general and in particular applications), most of the extant scholarship, particularly in the legal literature, operates under the implicit assumption that the relevant agents' informational endowments are exogenous. Yet both of these forms of ex ante control may affect an agent's incentive to gather costly information. Taking the endogeneity of information into account both complicates the standard hypotheses concerning the optimal design of ex ante control systems and suggests alternative strategies for inducing better-informed public policy decisions.

Imposing substantive or procedural constraints on the scope of an agent's policy discretion may affect the agent's research incentives in two ways:

First, these constraints may affect the agent's default payoff, thereby strengthening the agent's research incentives. The simplest way to illustrate this phenomenon is by considering a simple version of the discretionary window strategy in which the principal and the agent have identical policy preferences. In such a case, the standard ally principle would imply that the principal would prefer to delegate the agent unconstrained discretion. But suppose the principal were to

---

[45] *See* Gailmard, *supra* note 37; *cf.* David P. Baron, *Legislative Organization with Informational Committees*, 44 AM. J. POL. SCI. 485, 487 (2000) (discussing a similar control strategy in the context of a legislative chamber overseeing a committee).

[46] *See* Stephenson, *supra* note 42, at 469; Stephenson, *supra* note 12, at 2; *see also* Gersen, *supra* note 4, at 248; Gersen, *supra* note 3, at 302; John O. McGinnis, Medellín *and the Future of International Delegation*, 118 YALE L.J. 1712, 1715 (2009); Ernest A. Young, *Executive Preemption*, 102 NW. U. L. REV. 869, 877 (2008).

[47] *See* Gailmard, *supra* note 37, at 26.

[48] *See id.* Indeed, the discretionary window approach can be thought of as simply a special case of the enactment cost approach in which the enactment cost of any policy within the window is zero and the enactment cost of any policy outside the window is so large that the agent would never select it. *See id.*

prohibit the agent from choosing the policy that the agent (and the principal) would prefer ex ante. Doing so would lower the agent's default payoff, which tends to increase the agent's incentive to do research.[49] For instance, if a jury has to choose between convicting a defendant of a serious felony and letting him go, the jury may deliberate much more intensively than would be the case if the jury had the option to convict the defendant of a lesser offense that carried a lighter sentence.[50] Likewise, if a regulator must choose between banning a potentially hazardous product and leaving it unregulated, the regulator has a stronger incentive to accurately estimate the true dangers of the product than would be the case if the regulator could adopt some intermediate regulatory standard, such as a licensing regime. As the principal further restricts the agent's choice set by eliminating, say, the agent's second-, third-, and fourth-most-preferred policies ex ante, the agent's default payoff worsens further, and the agent's incentive to do research grows stronger.[51] A similar logic applies in contexts where the principal can manipulate the agent's enactment costs: raising the relative costs of the agent's ex ante preferred option lowers the agent's default payoff relative to her informed payoff, and this effect in turn stimulates greater research effort.[52]

Thus, one important — and perhaps surprising — general principle regarding the optimal structure of delegated authority is that when information is endogenous the principal may prefer to prohibit, or make relatively less desirable, those policy choices that the agent would be most inclined to favor ex ante, even if the agent's preferences are close-

---

[49] *See* Szalay, *supra* note 25, at 1174. Of course, eliminating the ex ante preferred option also reduces the agent's research payoff, in that research might show that the best course of action is indeed the one that the agent preferred ex ante. But in many plausible decision settings, the elimination of the ex ante preferred option will reduce the agent's default payoff by more than it reduces her research payoff. This outcome is especially likely when the ex ante preferred option is in fact unlikely to be the true optimum, but is instead a "compromise" option located between more extreme choices that might be very good or very bad. *See id.* It is, however, possible to construct cases in which elimination of the agent's ex ante preferred option actually weakens research incentives by reducing the agent's research payoff more than it reduces the agent's default payoff.

[50] The Supreme Court, however, has held that putting a jury to such a choice is unconstitutional, at least in capital contexts. *See* Beck v. Alabama, 447 U.S. 625, 627 (1980). The example in the text is not meant in any way as a critique of *Beck* or as an endorsement of the strategy of putting a jury to this sort of choice. It is merely a way to illustrate the intuition that when one is forced into an unhappy choice between extremes, one may think harder about the choice than would be the case if an attractive "moderate" or "compromise" option were readily available.

[51] *See* Szalay, *supra* note 25, at 1174. Of course, as the principal rules out additional possible policy choices beyond the agent's ex ante most-preferred choice, the agent's research payoff also decreases, which tends to weaken the agent's research incentives. *See supra* note 49. Even if the effect on the agent's default payoff is larger initially, past a certain point the net effect will reverse.

[52] *See* Stephenson, *supra* note 42, at 470.

ly aligned with those of the principal. Of course, an institutional design strategy that rules out the ex ante preferred option may have serious disadvantages: the ex ante preferred option (for example, conviction on the lesser offense, enactment of the licensing regime) may turn out to be the most desirable choice even after the agency has invested optimally in information (either because the agency's research efforts are unsuccessful or because the agency learns that the ex ante preferred option, or something close to it, is indeed the best policy to adopt). Nonetheless, it is possible that these ex post losses will be outweighed by the informational gain induced by restricting the agency's discretion.

Second, ex ante substantive and procedural constraints on an agent's discretion may also alter the agent's research payoff. Imagine, for example, that the principal narrows the scope of the agent's discretion, such that certain policy options are not available even if the agent learns information that would cause the agent to prefer those options. This change reduces the agent's research payoff, because narrowing the agent's discretion means that she may learn information that she cannot use (because she is legally barred from taking the action implied by that information). All else equal, this reduction in the agent's research payoff tends to decrease her incentive to do costly research.[53] Likewise, if the principal raises the relative enactment costs of policies other than the agent's ex ante preferred policy, then the agent's research payoff decreases and the agent's research incentive weakens.[54] In contrast, if the principal expands the agent's freedom to choose policies that seem ex ante undesirable — but that the agent might prefer given sufficient evidence — then the agent's research incentives become stronger. Similarly, if the principal lowers the relative enactment costs of policies other than the agent's ex ante preferred policy, the agent has a greater incentive to do research, as such research is more likely to affect her policy choice.

This analysis suggests that a principal might prefer to delegate substantial discretion even to a non-ally in order to induce greater research investment.[55] Likewise, even when the principal has the capac-

---

[53] *See* Philippe Aghion & Jean Tirole, *Formal and Real Authority in Organizations*, 105 J. POL. ECON. 1, 2–3 (1997); *see also* Kathleen Bawn, *Political Control Versus Expertise: Congressional Choices About Administrative Procedures*, 89 AM. POL. SCI. REV. 62, 63 (1995) (assuming, though not deriving, a positive correlation between the scope of an agency's authority and its expertise); *cf.* Jody Freeman, *Collaborative Governance in the Administrative State*, 45 UCLA L. REV. 1, 67 (1997) (suggesting, in the context of negotiated rulemaking, that the impulse to limit the scope of the negotiation by restricting the number of possible alternatives "conflicts with the reality that problem solving requires an effort not to foreclose creative ideas or new conceptions of the issues to be negotiated").

[54] *See* Stephenson, *supra* note 42, at 470–71.

[55] *See* Aghion & Tirole, *supra* note 53, at 2–3.

ity to manipulate the agent's enactment costs in order to perfectly align the policy preferences of the principal and the agent, the principal may not wish to do so.[56] In addition to these qualifications to the ally principle, the above discussion suggests an inversion of the expertise principle, which assumes that agencies' superior expertise (partially) explains the prevalence of broad delegations. While this assumption may well be true, it may also be the case that the existence of broad delegations partially explains why agencies have superior expertise.[57]

The main implication of the preceding discussion is that modifying the agent's discretion — either by altering her discretionary window or by manipulating her enactment costs — may simultaneously affect both the agent's incentives to gather information ex ante and the agent's incentives regarding use of that information ex post. This double effect may create a difficult trade-off for the principal. Even when the principal has the ability to precommit to a schedule of enactment costs that could align the agent's policy preferences perfectly with the principal's, the endogeneity of agency research effort may lead the principal to prefer a somewhat different schedule of enactment costs. The principal can offset the agent's policy bias by raising the relative enactment costs of policies that the agent tends to prefer more strongly than does the principal. But the principal can give the agent stronger incentives to acquire information by raising the relative enactment costs of the policies the uninformed agent would tend to favor. These considerations may often go together, but they need not. Suppose, for example, that the agent is an environmental regulator deciding whether to implement a new pollution control rule, and the aggregate environmental benefit of the rule takes some value, B, between 0 and 100, with any number in that range equally likely. Suppose further that the agency views the cost of the regulation as 40, but the legislature (the principal) views the cost as 60. If the agency's information about B is exogenous, the legislature will prefer to make the agency pay an enactment cost of 20 if it adopts the new regulation. Doing so will perfectly align the agency's policy preferences with the legislature's, leading the agency to adopt the regulation if, but only if, B is greater than 60. But if the agency's information about B is endogenous, the agency's incentive to invest in such information will be

---

[56] *See* Stephenson, *supra* note 42, at 470–71. The same change in institutional structure may often affect *both* the agent's default payoff *and* her research payoff. In such cases, assessing the net effect on the agent's research incentives requires sorting out the relative strength of these effects. This inquiry in turn requires more specific assumptions, based on more detailed information about the specific setting.

[57] *Cf.* Aghion & Tirole, *supra* note 53, at 2–3 (suggesting that granting more authority to an agent increases the agent's initiative and incentive to gain expertise); Gailmard & Patty, *supra* note 37, at 875 (suggesting that "allowing agents to bend policy to their liking" incentivizes "expertise development and a career in public service").

strongest if the enactment cost is 10 (as this enactment cost would make the agency ex ante indifferent between regulation and nonregulation). As the legislature increases the agency's enactment cost from 0 to 10, the legislature is both strengthening the agency's research incentive and aligning the agency's (induced) policy preferences more closely with the legislature's. But for enactment costs between 10 and 20, these interests trade off: increasing enactment costs from 10 to 20 moves the agency's (induced) policy preferences closer to the legislature's, but at the cost of weakening the agency's research incentives.

Thus, the principal might sometimes prefer to forgo the use of incentive schemes that would achieve a closer alignment between the principal's and agent's policy preferences. Indeed, a rational principal might even prefer to use an enactment cost scheme that moves the preferences of the agent *further* from those of the principal, if doing so sufficiently reduces the agent's expected default payoff and/or increases the agent's expected research payoff.[58] The problem for the principal is that she must use one tool — the schedule of enactment costs — to pursue these two objectives simultaneously. Thus, the delegation scheme that is optimal when the agent's information is exogenous may not be optimal when the agent's information is endogenous.

B. Evidentiary Rules and Burdens

Often the principal (or the overseers charged with ensuring that the agent faithfully implements the scheme enacted by the principal) will not be able to observe, comprehend, or verify the evidence in the agent's possession, or even observe accurate proxies for the agent's research effort. In those cases, the principal's ex ante control mechanisms will be limited to those discussed in section A. In other cases, however, the principal might be able to precommit to a system in which the constraints on the agent's decision are a function of the agent's information or research effort. Such constraints might take three forms. The first involves the imposition of evidentiary standards — burdens of production and proof — that an agent must meet before she is permitted to take certain actions. Second, the agent might be prohibited from considering certain kinds of evidence (or, equivalently, disabled from engaging in certain kinds of research). Third, the

---

[58] *See* Stephenson, *supra* note 42. To see this possibility, suppose in the example in the preceding paragraph that the agency viewed regulatory costs as equal to 60, and the legislature viewed these costs as equal to 70. If the legislature wanted to align the agency's ex post policy preferences with its own, it would impose an enactment cost equal to 10. But if the legislature wanted to maximize the agency's research incentives, it would provide an enactment *subsidy* equal to 10 (that is, a negative enactment cost). If the benefit to the legislature of increasing agency expertise is sufficiently large, then even a legislature that is more concerned about regulatory costs than is the agency might want to make regulatory interventions relatively more attractive to the agency.

agent's discretion might be conditioned on observable evidence of her research effort; this approach can be thought of as granting deference to an agent's perceived "expertise." All three of these mechanisms are familiar to public law scholars interested in the design of public decisionmaking systems. And, as the following discussion illustrates, analysis of all three may be quite different when one replaces the implicit assumption of exogenous expertise with the more realistic assumption of endogenous expertise.

　　1. Standards of Proof. — A familiar and intuitive way that a principal might use ex ante legal or procedural rules to mitigate a potential agency problem is for the principal to establish a system (perhaps enforced by a third-party adjudicator) in which an agent may take certain actions if, but only if, the agent produces adequate supporting evidence. For example, a regulatory agency might be permitted to adopt a stringent regulation only if it provides to some oversight body (for example, a court or review board) detailed scientific and economic data establishing that the regulation's likely benefits outweigh its costs. Similarly, a law enforcement agent (for example, a prosecutor or administrative enforcement division) may have the power to impose civil or criminal penalties on a defendant only if the agent can present to an adjudicator (for example, a judicial or administrative tribunal) sufficient evidence of the defendant's culpability. These sorts of constraints can be thought of as the establishment, by the principal, of the standard of proof that the agent must satisfy to take different possible actions.[59] The imposition of a standard of proof implies that the overseer can observe (perhaps imperfectly) the content of the evidence in the agent's possession.

　　If the agent's research effort is exogenous, then the principal will prefer a standard of proof that permits (indeed, requires) the agent to take whatever action the principal prefers in light of the agent's evidence at the time of decision. Of course, the strength of the evidence and the precise meaning of the linguistic formulations for various standards of proof are often open to debate. But, feasibility concerns aside, as a conceptual matter the institutional design problem of selecting a standard of proof is trivial when the agent's research effort is exogenous and the product of that research effort is observable.

　　The problem is more complicated once one considers the effect that the standard of proof may have on the agent's research effort. Perhaps most importantly, imposing a more demanding standard of proof can have a similar effect to prohibiting or penalizing the choice that the

---

　　[59] This section assumes that the principal commits to a standard of proof ex ante. Sometimes, of course, the "standard of proof" is just the amount of evidence the principal — or some other overseer — will demand ex post before deciding that a particular course of action is correct. This possibility is taken up in section II.C, pp. 1453–61.

agent is inclined to make ex ante — worsening the agent's default payoff and thereby strengthening the agent's incentive to do research.[60] This phenomenon may sometimes lead a principal to prefer a standard of proof that appears too stringent when viewed from an ex post perspective.[61] In effect, a more demanding evidentiary standard increases research effort by mandating a default option that the agent views as undesirable ex ante; this penalty default strengthens the agent's incentive to find enough evidence to enable the agency to avoid its default payoff.[62] A prosecutor who believes a defendant to be guilty is more likely to invest heavily in evidence gathering if she knows she must prove her case beyond a reasonable doubt than if she knows she must prove her case only by a preponderance of the evidence. A regulatory agency that wants to require a new auto safety technology will be more inclined to conduct rigorous, in-depth studies if the agency is required by statute to provide "clear and convincing evidence" that the regulation is justified than would be the case if the agency could implement the regulation so long as it provided "substantial evidence" of regulatory benefits.

The main drawback of this strategy, though, is that if the agent's greater research effort does not guarantee that she will acquire additional hard information, then committing to a stringent standard of proof may lead to suboptimal decisions ex post. This cost would need to be balanced against the informational gains that a more demanding standard of proof might spur. Furthermore, increasing the stringency of the standard of proof will not necessarily enhance the agent's research incentives. Indeed, under some circumstances, raising the standard of proof can *weaken* the agent's research incentives. There are two ways this weakening may occur. First, if the principal cannot affect the agent's default payoff (for example, if the agent both prefers the status quo ex ante and can guarantee herself the status quo outcome by failing to take action), then raising the standard of proof for *other* policy choices may reduce the agent's research payoff without affecting her default payoff. Second, even if the principal imposes a default payoff that the agency views as undesirable, if the standard of proof becomes too stringent the agent may simply decide not to act be-

---

[60] *See* Matthew C. Stephenson, *Evidentiary Standards and Information Acquisition in Public Law*, 10 AM. L. & ECON. REV. 351 (2008).

[61] *See* Bueno de Mesquita & Stephenson, *supra* note 44; Hao Li, *A Theory of Conservatism*, 109 J. POL. ECON. 617 (2001); Stephenson, *supra* note 60.

[62] This situation is roughly analogous to the use of penalty defaults to stimulate costly action in other contexts. *Cf.* EINER ELHAUGE, STATUTORY DEFAULT RULES (2008); Ian Ayres & Robert Gertner, *Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules*, 99 YALE L.J. 87 (1989); Baker & Krawiec, *supra* note 3; 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).

cause the agency anticipates that she will probably end up with her default payoff no matter how much research she does.[63] This outcome is one manifestation of the familiar concern about how overly demanding evidentiary requirements can "ossify" public decisionmaking.[64] So increasing the stringency of the standard of proof may strengthen the agent's research incentives, but only up to the point where the standard becomes so hard to satisfy that the agent is better off not even trying.[65]

In addition, the principal may be able to manipulate evidentiary standards in order to improve the agent's research payoff. Consider a case in which the principal and the agent both view a certain policy choice as relatively undesirable ex ante, but the amount of evidence it would take to convince the agent to adopt that policy is less than the amount of evidence that the principal would demand. If information were exogenous, the principal would simply insist that the agent take the action only if the agent had enough evidence to convince the principal that it was a good idea. But if research were endogenous, the principal's ex post optimal standard of proof might discourage the agent from researching at all. If the principal lowered the standard of proof somewhat, the agent would have more of an incentive to at least look into the option in question, and this benefit to the principal from the increase in research might outweigh the expected ex post cost.

2. Exclusion of Probative Evidence. — Many legal and institutional rules attempt to prohibit a decisionmaker from considering certain types of apparently useful information. Most obviously, the law of evidence excludes numerous categories of probative evidence.[66] Something similar also occurs in the regulatory context. For example, some statutes prohibit agencies from taking certain seemingly relevant factors — such as regulatory costs — into account.[67] Requirements that

---

[63] *See* Stephenson, *supra* note 60.

[64] *See* Bueno de Mesquita & Stephenson, *supra* note 44; Thomas O. McGarity, *Some Thoughts on "Deossifying" the Rulemaking Process*, 41 DUKE L.J. 1385 (1992); Richard J. Pierce, Jr., *Seven Ways to Deossify Agency Rulemaking*, 47 ADMIN. L. REV. 59 (1995).

[65] *See, e.g.*, JERRY L. MASHAW & DAVID L. HARFST, THE STRUGGLE FOR AUTO SAFETY (1990) (suggesting that something like this outcome occurred in the auto safety area, where escalating judicial demands for rigorous evidence in support of design standards caused the Department of Transportation to abandon that regulatory approach, in favor of the less efficient ex post recall of models shown to be unsafe).

[66] *See, e.g.*, MIRJAN R. DAMAŠKA, EVIDENCE LAW ADRIFT 84–86 (1997); *see also* ALEX STEIN, FOUNDATIONS OF EVIDENCE LAW (2005); Richard D. Friedman, *Toward a Partial Economic, Game-Theoretic Analysis of Hearsay*, 76 MINN. L. REV. 723 (1992); Dale A. Nance, *The Best Evidence Principle*, 73 IOWA L. REV. 227 (1988); Dale A. Nance, *Conditional Relevance Reinterpreted*, 70 B.U. L. REV. 447 (1990).

[67] There is a voluminous literature debating the merits of permitting and prohibiting regulatory agencies to take costs into account when setting regulatory policies. *See, e.g.*, FRANK ACKERMAN & LISA HEINZERLING, PRICELESS 35–40 (2004); RICHARD L. REVESZ & MICHAEL A. LIVERMORE, RETAKING RATIONALITY 151–90 (2008); CASS R. SUNSTEIN, THE COST-

force a decisionmaker to ignore relevant evidence have widely noted disadvantages.[68]   They may, however, have the advantage of strengthening an agent's incentives to invest effort in acquiring or analyzing *other* sorts of evidence that may be more costly for the agent to acquire, but also more probative.[69]

Consider the following example.   Suppose that a legislature delegates to an administrative agency the authority to adopt a pollution control regulation.   Both the environmental benefits (B) and the economic costs (C) of the regulation are uncertain; ex ante, there is a 25% chance that the regulation will produce high environmental benefits ($B = 100$), a 50% chance that it will produce moderate benefits ($B = 50$), and a 25% chance that it will produce no benefits ($B = 0$).   On the cost side, there is a 50% chance that the economic costs of the regulation are low ($C = 40$) and a 50% chance that they are high ($C = 60$).   Adopting the regulation gives both the agency and the legislature a payoff of $B - C$.   The agency has the ability to do additional research on both costs and benefits.   At a private cost of 12, the agency can conduct an additional scientific analysis that reveals the true environmental benefit of the regulation, B.   For a private cost of 4, the agency can perform an additional economic analysis that reveals the regulation's true economic burden, C.

Ideally, the legislature would like the agency to do additional research with respect to both benefits and costs, as this would improve the legislature's expected utility by 15.[70]   However, although the agen-

---

BENEFIT STATE (2002); Bruce A. Ackerman & Richard B. Stewart, *Reforming Environmental Law: The Democratic Case for Market Incentives*, 13 COLUM. J. ENVTL. L. 171 (1988); John P. Dwyer, *The Pathology of Symbolic Legislation*, 17 ECOLOGY L.Q. 233, 255–57 (1990); Christopher T. Giovinazzo, *Defending Overstatement: The Symbolic Clean Air Act and Carbon Dioxide*, 30 HARV. ENVTL. L. REV. 99 (2006); John D. Graham, *Saving Lives Through Administrative Law and Economics*, 157 U. PA. L. REV. 395 (2008); Shi-Ling Hsu, *Fairness Versus Efficiency in Environmental Law*, 31 ECOLOGY L.Q. 303 (2004).

[68] *See, e.g.*, Susan Haack, *Epistemology Legalized: Or, Truth, Justice, and the American Way*, 49 AM. J. JURIS. 43, 55–61 (2004); Richard A. Posner, *An Economic Approach to the Law of Evidence*, 51 STAN. L. REV. 1477 (1999); Frederick Schauer, *On the Supposed Jury-Dependence of Evidence Law*, 155 U. PA. L. REV. 165, 168–70 (2006).

[69] *See* Benjamin Lester et al., *Information Acquisition and the Exclusion of Evidence in Trials*, J.L. ECON. & ORG. (forthcoming 2011), *available at* http://jleo.oxfordjournals.org/content/early/2009/12/30/jleo.ewp040.full.pdf; *see also* Friedman, *supra* note 66; Nance, *The Best Evidence Principle*, *supra* note 66; Nance, *Conditional Relevance Reinterpreted*, *supra* note 66.

[70] The legislature's expected utility if the agency remains uninformed is 0. If the agency learns both B and C before regulating, then there is a 50% chance that the agency will not regulate (a 25% chance that B is low, plus a 25% chance that B is moderate and C is high), a 12.5% chance that both B and C are high (in which case the agency will regulate, giving the legislature a net payoff of 40), a 12.5% chance that B is high and C is low (in which case the agency regulates and the legislature's net payoff is 60), and a 25% chance that B is moderate and C is low (in which case the agency regulates and the legislature gets a payoff of 10). Thus, if the agency learns both B and C, the legislature's expected utility is $(0.5 \times 0) + (0.125 \times 40) + (0.125 \times 60) + (0.25 \times 10) = 15$.

cy would also realize this utility gain, the agency's total private cost of doing research on both benefits and costs is 16. Left to its own devices, the agency would do additional research on the regulation's economic costs but not its environmental benefits. Research on costs improves the agency's expected utility by 5 units but costs the agency only 4, for a net gain of 1.[71] If the agency conducts research only on environmental benefits, the expected utility from the final policy choice will be 12.5.[72] This expected payoff is larger than the cost of researching environmental benefits (12), but the *net* expected utility gain to the agency, 0.5, is smaller than the net gain from conducting research solely on the regulation's economic costs.

Now suppose that the legislature prohibits the agency from considering costs when deciding whether to regulate. Such a rule might seem irrational ex post: if the agency has hard evidence that C is high, then forcing the agency to regulate even when B is moderate or uncertain leads to a net expected utility loss for both the agency and the legislature. But viewed from an ex ante perspective, such an evidentiary rule can improve the legislature's welfare. Left to its own devices, the agency would research C rather than B, as this approach improves the agency's expected utility by 1, whereas investigating B rather than C improves the agency's utility by 0.5. The legislature, however, would prefer that the agency research B rather than C if the choice is one or the other: if the agency learns B, the legislature realizes an expected utility gain of 12.5, whereas if the agency researches only C, the legislature's expected utility gain is 5. Thus, prohibiting the agency from researching C induces the agency to choose its second-best option — research on B — which is better for the legislature.

Observe that the exclusionary rule illustrated by the preceding example is equivalent to the imposition of a more demanding standard of proof of the sort discussed in section 1. By excluding consideration of C, the legislature reduces the agency's default payoff if it does not do research on B. The only twist here is that in order to achieve this re-

---

[71] If the agency learns C but not B, then the agency will regulate if C is low but not if C is high. In the 50% of cases where the agency learns that C is high, the agency does not regulate, and the agency's payoff is 0. In the 50% of cases where the agency regulates (after learning that C is low), the payoff is $(0.5 \times 10) + (0.25 \times 60) - (0.25 \times 40) = 10$. This means that, if the agency learns C but not B, the expected payoff to the legislature and to the agency (prior to netting out research costs) is $0.5 \times 10 = 5$.

[72] If the agency learns B but not C, there is a 75% chance that the legislature's expected payoff is 0. (There is a 25% chance that the agency learns that B is low, in which case the agency will not regulate, and a 50% chance that the agency learns that B is moderate, in which case the agency is indifferent between regulation and nonregulation, as both yield a 0 payoff.) There is, however, a 25% chance that the agency will learn that B is high, in which case the agency will regulate, giving the legislature an expected payoff of $(0.5 \times 40) + (0.5 \times 60) = 50$. Thus, if the agency learns B but not C, the expected payoff to the legislature (and the payoff to the agency prior to netting out research costs) is $0.25 \times 50 = 12.5$.

sult, the legislature excludes consideration of *other* probative evidence, in this case, *C*. But that is more a difference of technique than substance. More generally, this example illustrates that when an agent may engage in a variety of research activities that are partial substitutes, and the principal and the agent differ on the optimal allocation of agency effort across different research activities (for example, because the forms of research that are most informative are also the costliest for the agent), the principal may prefer to commit to evidentiary rules that discount the significance of the evidence that the agent would be most inclined to acquire and magnify the importance of the sorts of evidence that the principal views as most useful.[73] Doing so can lead to ex post inefficiencies but may nonetheless be optimal if the informational gains induced by this evidentiary standard outweigh those ex post losses.[74]

3. Standards of Deference. — The preceding sections have considered situations in which the principal can commit to a set of constraints on the agent's discretion that are conditional on the quality of the agent's evidence but are not contingent on the agent's research effort. If, however, the principal can commit to a system in which the agent's freedom of action is conditional on the agent's research effort (or a sufficiently good proxy for research effort), then the principal can induce both optimal ex ante research and optimal ex post decisionmaking by threatening the agent with a very bad outcome — mandating a decision very far from the agent's ex ante preferred policy — if the agent fails to invest what the principal views as the optimal amount of research effort. However, this approach may not always be feasible. Suppose, for instance, that the agent always has the ability to ensure that the final outcome is the status quo, for example, by simply failing to take action. This possibility prevents the principal from lowering the agent's default payoff below her status quo payoff.

If there is a lower bound on how much the principal can decrease the agent's default payoff, then the principal can strengthen the agent's research incentives further only by improving the agent's research payoff. The preceding sections have illustrated a variety of ways the principal might do this, such as giving the agent more discretion to choose policies other than her ex ante preferred option and low-

---

[73] *See* Lester et al., *supra* note 69, at 2 (arguing that judges exclude evidence that jurors would find easy to understand to focus the jurors on more difficult, but more probative, evidence).

[74] This conclusion does not mean that the exclusion of particular kinds of evidence is the *best* way to induce an agent to invest in gathering relevant information. The principal would likely be better off if it could impose the optimal ex post evidentiary standard and then offer the agent some kind of separate transfer payment, or threaten some separate punishment, based on the quantity or quality of the agent's evidence. But sometimes these more efficient mechanisms are not feasible. *See supra* pp. 1433–34.

ering the standard of proof associated with such policies. If the principal can make the agent's discretion contingent on the agent's research effort, the principal can use these tools more effectively. For example, the principal could guarantee the agent a broader range of policy discretion, or a lower standard of evidentiary proof, if the agent invested more effort in research. Doing so is a way for the principal to "pay" the agent for doing costly research when more direct subsidies are not feasible. A principal that employs such a strategy establishes, in essence, a variable standard of deference in which agents more likely to be "expert" — by virtue of their observable investment in research activities — have greater freedom of action.

The idea that there might be a positive correlation between an agent's expertise and her policy discretion is a familiar one — this is exactly what the well-known uncertainty principle (or expertise principle) would predict.[75] But the standard argument assumes that the principal gives better-informed agents more discretion solely because they have more of an informational advantage vis-à-vis the principal. The argument sketched above suggests a quite different (though not mutually exclusive) reason why a principal might confer more discretion on an agent that invests substantially in acquiring information: to create stronger incentives for the agent to invest in costly research. Thus, as suggested earlier, policy discretion may be both a response to the fact that the agent has better information and a cause of the agent's acquisition of such information.[76] This possibility, in turn, suggests that when the latter effect is taken into account, the positive correlation between expertise and discretion may be substantially stronger than the conventional analysis would predict. "Paying" the agent for research by offering more discretion or a lower standard of proof is costly to the principal, because doing so may allow the agent to make a policy choice ex post that differs from the principal's optimal choice, but that payment may be worth it if it stimulates a sufficient increase in the agent's research effort.

C. Ex Post Oversight

Sections A and B considered the effect of legal and institutional "rules of the game" that the principal establishes ex ante. Another familiar and important institutional mechanism for addressing the agency problem involves ex post oversight by the principal (or some other

---

[75] *See supra* notes 39–40; *see also* Jody Freeman & Adrian Vermeule, Massachusetts v. EPA*: From Politics to Expertise*, 2007 SUP. CT. REV. 51, 52 (arguing that recent case law suggests that courts are endeavoring "to ensure that agencies exercise expert judgment free from outside political pressures").

[76] *See* section II.A, pp. 1439–46.

overseer acting on the principal's behalf).[77] Of course, many of the institutional mechanisms discussed in the earlier sections also presumed some form of oversight, but in those cases the overseer was enforcing a set of rules established in advance. Indeed, a critical presumption of the preceding sections was that the principal could credibly precommit to those rules and requirements. Such broad precommitment, however, may not always be feasible.[78] Instead, in many settings the principal can only establish forms of review in which the overseer makes whatever decision is optimal ex post, rather than enforcing a set of rules that would be optimal ex ante.

This kind of after-the-fact review is an important feature of many legal and political institutions. Most obviously, appellate courts review the decisions of lower courts.[79] Courts may also review legislative and administrative decisions for reasonableness under various constitutional and statutory standards.[80] Congress oversees, and may overturn by statute (or indirectly, through other forms of coercion), the decisions of executive branch departments.[81] There are also a variety of oversight mechanisms within the executive branch, such as review of proposed regulations by the Office of Management and Budget.[82] This section

---

[77] This sort of oversight has often been discussed in the more general context of political audits of bureaucratic performance. *See, e.g.*, Jeffrey S. Banks, *Agency Budgets, Cost Information, and Auditing*, 33 AM. J. POL. SCI. 670 (1989); Jeffrey S. Banks & Barry R. Weingast, *The Political Control of Bureaucracies Under Asymmetric Information*, 36 AM. J. POL. SCI. 509 (1992); Jonathan Bendor et al., *Politicians, Bureaucrats, and Asymmetric Information*, 31 AM. J. POL. SCI. 796 (1987); Bueno de Mesquita & Stephenson, *supra* note 44; Anne Joseph O'Connell, *Auditing Politics or Political Auditing?* (Univ. of Cal. Berkeley Pub. Law & Legal Theory Research Paper Series, Paper No. 964656, 2007), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id= 964656; Anne Joseph O'Connell, Who Walks the Watchdog? Bureaucratic Oversight and the Government Accountability Office (July 5, 2007) (unpublished manuscript) (on file with the Harvard Law School Library). There is a closely related literature on the design of hierarchical court systems. *See* Charles M. Cameron & Lewis A. Kornhauser, *Appeals Mechanisms, Litigant Selection, and the Structure of Judicial Hierarchies*, *in* INSTITUTIONAL GAMES AND THE U.S. SUPREME COURT 173 (James R. Rogers et al. eds., 2006); Andrew F. Daughety & Jennifer F. Reinganum, *Appealing Judgments*, 31 RAND J. ECON. 502 (2000); Steven Shavell, *The Appeals Process as a Means of Error Correction*, 24 J. LEGAL STUD. 379 (1995); Matt Spitzer & Eric Talley, *Judicial Auditing*, 29 J. LEGAL STUD. 649 (2000).

[78] *See supra* note 36 (noting the difficulties with the precommitment assumption).

[79] *See, e.g.*, Shavell, *supra* note 77; Spitzer & Talley, *supra* note 77.

[80] *See, e.g.*, Stephenson, *supra* note 44, at 757–67 (discussing generally controversies surrounding judicial review of agency policy choices); Stephenson, *supra* note 12, at 7–8 (discussing general phenomenon of judiciary engaging in interest balancing when conducting constitutional review).

[81] *See, e.g.*, JOEL D. ABERBACH, KEEPING A WATCHFUL EYE (1990); Mathew D. McCubbins & Thomas Schwartz, *Congressional Oversight Overlooked: Police Patrols Versus Fire Alarms*, 28 AM. J. POL. SCI. 165 (1984); Barry R. Weingast & Mark J. Moran, *Bureaucratic Discretion or Congressional Control? Regulatory Policymaking by the Federal Trade Commission*, 91 J. POL. ECON. 765 (1983).

[82] *See, e.g.*, Christopher C. DeMuth & Douglas H. Ginsburg, Commentary, *White House Review of Agency Rulemaking*, 99 HARV. L. REV. 1075, 1075 (1986).

considers the implications of ex post oversight when the agent's research effort is endogenous.

The discussion that follows considers separately three cases, each of which makes different assumptions about the information that the overseer has at the time it reviews the agent's decision. In the first case, the agent's information, once acquired by the agent, is fully transparent, such that the overseer can verify the content of that information. The second case entails a complete absence of transparency, in which the overseer can observe neither the agent's research effort nor the content of her evidence. In the third case, the overseer cannot observe the agent's information but can observe the agent's research effort (or some other proxy for the probability that the agent has good information).[83] In each case, the analysis focuses on how standard positive and normative evaluations of ex post oversight might change when the agent's research effort is endogenous.

1. Oversight Under Full Transparency. — In the full transparency case, the principal can independently assess and verify the evidence that the agent had at the time she made her decision. In this case, if there are no constraints on the principal's ability to revise the agent's decision, then the agent is relegated to the role of a de facto research assistant: although the agent may make a recommendation and provide supporting evidence, only the evidence (not the recommendation) really matters, because the principal makes the actual decision. If the agent's research effort is exogenous, then the principal will clearly prefer ex post oversight both to a setting in which the agent has unfettered discretion and to a setting in which the principal makes the policy choice without consulting the agent's evidence. Furthermore, the agent's own policy preferences are irrelevant in this case, so the principal does not care about influencing those preferences.

If, however, the agent's research effort is endogenous, the principal's calculations may be quite different. The effect of ex post oversight on the agent's incentive to do costly research turns on the nature of the preference divergence between the principal and the agent. If the expected divergence between the preferred policies of the principal and the agent is approximately the same in low-uncertainty and high-uncertainty environments (or if the principal and the agent are more likely to agree on policy in high-uncertainty environments), then ex post oversight with full transparency tends to weaken the agent's re-

---

[83] All three of these cases presume that any additional information the principal acquires must come from the agent. Another possibility is that the principal might also be able to invest costly effort in acquiring information, and might make this choice endogenously, in response to the agent's decision. This possibility is taken up in Part III, pp. 1461–82, which considers information gathering in contexts with multiple agents (one of whom might be an overseer acting on behalf of the principal).

search incentives relative to the case in which the agent has unconstrained discretion. The reason is that although ex post oversight reduces both the agent's default payoff and her research payoff, the latter effect is stronger.[84] This result suggests that ex post oversight with full transparency may exacerbate the problem of underinvestment in information and may therefore be suboptimal for the principal even compared to an alternative that gives a somewhat biased agent unconstrained policy discretion.[85]

There is another possibility, however. In some plausible scenarios, the expected divergence between the preferred policies of the principal and the agent will narrow as uncertainty decreases. In other words, the principal and the agent may find themselves in agreement regarding the policy prescriptions that follow from particular forms of hard evidence, even though they have dramatically different views of the best decision to make when such evidence is unavailable. For example, a bureaucratic agency and a legislative overseer might agree that a certain chemical should be banned if it is carcinogenic and left unregulated if it is not, but nonetheless disagree about what to do if the scientific evidence is inconclusive. The agency, for instance, might prefer to err on the side of safety, while the legislature prefers erring on the side of nonintervention.

If the policy preferences of the agent and the principal tend to converge in information-rich environments (or, equivalently, if they tend to diverge more sharply in information-poor environments), then the adverse impact of oversight on the agent's default payoff becomes

---

[84] To illustrate, suppose that an administrative agency must select a degree of regulatory stringency, S, and that the agency's ideal level of stringency depends on the level of risk from the targeted activity, R. Suppose further that R is initially uncertain; it is somewhere between o and 100, with any value in that range equally likely. Suppose that the agency's utility from the policy choice is $100 - |S - R|$. That is, the agency would prefer to set $S = R$, and its utility decreases linearly in the distance between these values. If there is no oversight, then if the agency remains uninformed, it would select $S = 50$ for a default payoff of 75, whereas if the agency learns R, it will select $S = R$ for a research payoff of 100. Now suppose the agency's decision is reviewed by a legislature that prefers to set $S = R - 10$. In this case, then under full transparency the final outcome if the agency remains uninformed is $S = 40$, which gives the agency a default payoff of 74. This result is lower than the agency's default payoff without oversight, but only slightly. If, however, the agency learns R, then the final outcome will be $S = R - 10$, which gives the agency a research payoff of 90 — much lower than its research payoff in the absence of oversight. The intuition that this abstract example illustrates is that the overseer can take advantage of the agent's information to shift the policy outcome even further from the agent's ideal. The more information the agent acquires, the greater the overseer's ability to take advantage of it, to the agent's detriment.

[85] *See* Aghion & Tirole, *supra* note 53, at 6–7; Sean Gailmard, *Expertise, Subversion, and Bureaucratic Discretion*, 18 J.L. ECON. & ORG. 536 (2002); Gailmard & Patty, *supra* note 37, at 875–86; *cf.* Thomas W. Gilligan & Keith Krehbiel, *Collective Decisionmaking and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures*, 3 J.L. ECON. & ORG. 287, 288–89 (1987); Thomas W. Gilligan & Keith Krehbiel, *Organization of Informative Committees by a Rational Legislature*, 34 AM. J. POL. SCI. 531, 533 (1990).

larger relative to the adverse impact on the agent's informed payoff. This relationship reduces the adverse impact of oversight on the agent's research incentives and, past a certain point, can actually flip the effect, making ex post oversight strengthen rather than weaken an agent's research incentives. When this is so, then even if establishing ex post oversight would be undesirable if agency research were exogenous (if, for example, setting up the institutions to facilitate meaningful legislative oversight would involve substantial costs), the positive effect on agency research incentives may make such oversight desirable on balance.

This conclusion has another potentially important implication: if the principal has some ability to select (or at least to influence) the agent's preferences, the principal may not necessarily want to select a perfect ally. When the agent's decision is subject to ex post review, the principal can stimulate greater research effort by selecting an agent with policy preferences that tend to converge with the principal's as more information becomes available but differ from the principal's in information-poor environments. In other words, the principal's ideal agent may be one who agrees about ultimate policy objectives in general terms but has a quite different approach to decisionmaking under uncertainty. For example, the principal may prefer an agent who places different relative weights on the error costs associated with false positives and false negatives than does the principal. For similar reasons, the principal might sometimes prefer an agent with similar underlying policy preferences but quite a different level of risk aversion.

2. Oversight Without Transparency. — In many settings involving ex post oversight, the substantive content of the agent's evidence is not transparent to the principal (either because the principal literally never sees the evidence or because the principal lacks the capacity to understand or verify it). In such cases, the agent may try to persuade the principal to choose a particular policy by announcing the content of the evidence it has uncovered. However, because the principal cannot verify the agent's announcement, that announcement is, in economics jargon, "cheap talk." If the agent's preferences diverge from those of the principal, the agent has an incentive to misrepresent the content of her evidence. The principal, knowing this, has an incentive to discount the agent's representations. This discounting of the agent's report is stronger when the preference divergence between the principal and the agent is larger: when the agent's preferences and the principal's preferences are far apart, the agent has strong incentives to give

a misleading report, and the principal consequently pays little attention to the agent's recommendation.[86]

For example, suppose a scientist with (known) proregulatory policy preferences must deliver to a more conservative policymaker a report on the health hazards of a certain product. If the scientist issues an honest report, the policymaker will regulate the product much less stringently than the scientist believes is optimal. So the scientist has an incentive to issue a report suggesting that the health hazards of the product are even worse than they actually are, in the hopes of inducing the policymaker to adopt a more stringent policy. The policymaker, however, is aware of this incentive and discounts the report accordingly. Thus, in equilibrium, either the scientist will issue an extreme report (which the policymaker will interpret in light of the scientist's known bias)[87] or, if the scientist determines that the true health hazard is more moderate, the scientist will issue a moderate report that the policymaker will simply ignore, choosing instead her ex ante preferred policy. This latter result is inefficient, as valuable information (contained in the scientist's report) is lost because the policymaker does not trust the scientist sufficiently. This inefficiency in communication is exacerbated by the fact that, if the principal is also uncertain about the agent's preferences or ability, the agent has an incentive to issue a re-

---

[86] The seminal work on cheap talk signaling is Vincent P. Crawford & Joel Sobel, *Strategic Information Transmission*, 50 ECONOMETRICA 1431 (1982). *See also* ARTHUR LUPIA & MATHEW D. MCCUBBINS, THE DEMOCRATIC DILEMMA (1998); Wouter Dessein, *Authority and Communication in Organizations*, 69 REV. ECON. STUD. 811 (2002); Joseph Farrell & Robert Gibbons, *Cheap Talk with Two Audiences*, 79 AM. ECON. REV. 1214 (1989); Joseph Farrell & Matthew Rabin, *Cheap Talk*, 10 J. ECON. PERSP. 103 (1996); Vijay Krishna & John Morgan, *A Model of Expertise*, 116 Q.J. ECON. 747 (2001). For applications of cheap talk models to political and legal institutions, see David Austen-Smith, *Strategic Models of Talk in Political Decision Making*, 13 INT'L POL. SCI. REV. 45 (1992); Jacob E. Gersen & Eric A. Posner, *Soft Law: Lessons from Congressional Practice*, 61 STAN. L. REV. 573, 589–90, 596–97 (2008); Leslie Johns, *A Servant of Two Masters: Communication and the Selection of International Bureaucrats*, 61 INT'L ORG. 245 (2007); Timothy R. Johnson et al., *Oral Advocacy Before the United States Supreme Court: Does It Affect the Justices' Decisions?*, 85 WASH. U. L. REV. 457, 468, 503–12 (2007); John W. Patty, *The Politics of Biased Information*, 71 J. POL. 385 (2009); Otto H. Swank et al., *A Theory of Policy Advice*, 15 J.L. ECON. & ORG. 602, 603–04 (1999); and Sean Gailmard & John W. Patty, Congressional Development of the Institutional Presidency (Oct. 2010) (unpublished manuscript) (on file with the Harvard Law School Library).

[87] If the scientist learns that the product is actually not hazardous at all, and issues an "extreme" report to that effect, the policymaker will view the report as more credible because it runs counter to the scientist's expected preferences. If the scientist learns that the product is extremely hazardous, and issues a report to that effect, the policymaker may listen because the scientist is recommending something that goes well beyond what even someone with her preferences would ordinarily endorse, but will still discount the report somewhat, and will issue a regulatory standard that is not as stringent as what the scientist recommends.

Readers familiar with the economic analysis of cheap talk games cited in note 86 will recognize that the explanation in the text is not entirely accurate, and omits important subtleties and complexities. The informal discussion here is intended to give nontechnical readers a basic intuition for why preference divergence leads to informational losses in these settings.

port that makes the agent appear in the best light, further devaluing the informative content of her recommendation.[88] In the above example, if the scientist fears dismissal from her post if she is perceived as "too left wing," she may always (or almost always) issue a report to the effect that the health hazards of the product under investigation are minimal.

The most important finding of the extant literature on cheap talk settings is that both the principal and the agent are better off when their preferences are more closely aligned.[89] The reason for this result, suggested by the example above with the scientist and the policymaker, is that an agent can never use a cheap talk report to bias the expected policy outcome in a favorable direction; the most she can do is to make the policy decision more accurate by providing better information to the principal. Preference divergence leads to informational losses for both the principal and the agent but does not alter the fact that the policy selected will be whichever one seems to be in the principal's best interest. For this reason, it may sometimes actually be better for the principal to eliminate or constrain its ability to review the agent's decision: doing so allows the agent to bias the policy outcome but also makes the ultimate choice more informed.[90]

These basic results, developed primarily in models that presume that the agent's information is exogenous, may be strengthened when the agent's research decision is endogenous. After all, the agent knows that she cannot do anything to shift the expected policy outcome in a favorable direction; the most she can do is to provide the principal with a report that will enable the principal to make a more informed decision. However, the greater the divergence between the preferences of the principal and the agent, the less likely the agent is to acquire information that will lead to *any* change in the principal's policy decision. This tendency is easiest to see in extreme cases: if the principal and the agent are perfectly aligned, then any information the agent acquires will be useful, because any report by the agent can induce the principal to alter the final policy decision; if the principal and agent have such different preferences that the principal will disregard *any* report by the agent, then the agent has no incentive whatsoever to acquire information, since this information will always be ignored. (In

---

[88] *See* Timur Kuran & Cass R. Sunstein, *Availability Cascades and Risk Regulation*, 51 STAN. L. REV. 683 (1999); Stephen Morris, *Political Correctness*, 109 J. POL. ECON. 231 (2001); *cf.* Prendergast, *supra* note 24, at 757–59 (explaining the incentive of workers who are subjectively evaluated to conform their opinions to those of their supervisors).

[89] *See* Crawford & Sobel, *supra* note 86, at 1442.

[90] *See* Bendor & Meirowitz, *supra* note 38; Boehmke et al., *supra* note 38; Gailmard & Patty, *supra* note 86; *cf.* Sven E. Feldmann, Bureaucratic Expertise and Learning from Interest Groups (June 1, 2005) (unpublished manuscript) (on file with the Harvard Law School Library).

the above example, if the scientist must do costly research to assess the health hazards of a particular product, she is unlikely to do so if she expects the policymaker always to choose nonregulation, no matter how dire the scientist's report.) Preference alignment, then, has two desirable effects in this setting: it not only improves information transmission (holding the agent's research effort constant), but also may strengthen the agent's incentive to invest in research.

3. Oversight with Observable Research Effort. — In some settings, the principal might be able to observe the agent's research *effort* (or some other proxy for the expected quality of the agency's information), even if the substantive *content* of the agent's information remains hidden. For example, even when a legislature, court, or other overseer is unable to assess the quality or accuracy of the technical evidence that an administrative agency proffers in support of a proposed regulation, the overseer might be able to assess how much effort the agency has invested in acquiring this evidence.[91]

Recall that in settings where the principal could condition a set of ex ante rules on the agent's observable research investment, the principal had substantial ability to induce the agent to invest optimally in research. This conclusion no longer holds in ex post oversight settings, because the principal often cannot credibly threaten to punish the agent (say, by reversing the agent's decision) if the agent fails to exert enough research effort. Even when the agent's research effort is observable, when the substantive content of the agent's evidence is not transparent to the principal, the agent will find herself in a situation where her report is cheap talk.

Even in this cheap talk setting, however, the observability of the agent's research effort may have important effects on the agent's research incentives. As noted in the discussion of ex post oversight with full transparency,[92] there may be cases in which the preferences of the principal and agent tend to converge in high-information environments but diverge in low-information environments, as well as cases in which the opposite is true. When either of those conditions holds, the agent's research investment is an indirect proxy for the degree of preference alignment between the principal and the agent. And, as discussed earlier, in cheap talk settings, the agent's report is more likely to be influential when the preferences of the agent are closer to those of the principal.

These relationships suggest that when the preferences of the principal and the agent tend to converge in information-poor environments but diverge in information-rich environments, then an agent has even

---

[91] *Cf.* Stephenson, *supra* note 44, at 773–75; Stephenson, *supra* note 43, at 551.

[92] *See* section II.C.1, pp. 1455–57.

weaker incentives to invest in research if the principal can observe the agent's research effort. Perhaps more interestingly, if preferences tend to converge as information improves, then the principal will rationally treat the agent's research effort as a proxy for preference alignment. Thus, the agent's report is more influential when the agent has done a great deal of research, which in turn strengthens the agent's research incentives. This result implies a potentially important deviation from the ally principle: although the principal prefers an agent who acts as a close ally in information-rich environments, the principal may have an incentive to select an agent with substantially different preferences in information-poor environments. Selecting such an agent is costly to the principal, in that the principal will not be able to extract as much useful information from the agent if the agent remains relatively uninformed. But this preference divergence intensifies the agent's incentive to invest in research, and it is possible that this informational gain outweighs the losses associated with less efficient information transmission.

III. RESEARCH INCENTIVES WITH MULTIPLE AGENTS

Analyzing the research incentives of a single agent is useful in building intuition and establishing some basic concepts. However, most public decision settings entail at least the potential for multiple agents to contribute to information gathering. Most of the relevant government "agents" — including legislatures, bureaucratic agencies, judiciaries, and executives — are in fact composed of many individuals,[93] and many public decisions involve contributions from more than one of these entities. Indeed, much recent work in institutional design theory has moved beyond the simple question of which among several possible agents should have principal responsibility for a particular policy decision to the more complex question of how to arrange decisionmaking systems that entail input from many different agents.[94]

---

[93] *See, e.g.*, Neal Kumar Katyal, *Internal Separation of Powers: Checking Today's Most Dangerous Branch from Within*, 115 YALE L.J. 2314, 2316–18 (2006); Kenneth A. Shepsle, *Congress Is a "They," Not an "It": Legislative Intent as Oxymoron*, 12 INT'L REV. L. & ECON. 239, 240–41 (1992); Adrian Vermeule, *The Judiciary Is a They, Not an It: Interpretive Theory and the Fallacy of Division*, 14 J. CONTEMP. LEGAL ISSUES 549, 549, 554–55 (2005).

[94] *See, e.g.*, Justin Fox & Matthew C. Stephenson, *Judicial Review as a Response to Political Posturing*, 105 AM. POL. SCI. REV. (forthcoming 2011); Justin Fox & Richard Van Weelden, *Partisanship and the Effectiveness of Oversight*, 94 J. PUB. ECON. 674 (2010); Saul Levmore, *Bicameralism: When Are Two Decisions Better than One?*, 12 INT'L REV. L. & ECON. 145, 145–46, 162 (1992); Nzelibe & Stephenson, *supra* note 3, at 631–37; Torsten Persson et al., *Separation of Powers and Political Accountability*, 112 Q.J. ECON. 1163 (1997); James R. Rogers, *Legislative Incentives and Two-Tiered Judicial Review: A Game Theoretic Reading of* Carolene Products *Footnote Four*, 43 AM. J. POL. SCI. 1096 (1999); Stephenson & Nzelibe, *supra* note 23.

In addition, an important strain of the political economy literature considers settings involving multiple *principals*, which compete to influence the behavior of one or more common agents. *See, e.g.*, B. Douglas Bernheim & Michael D. Whinston, *Common Agency*, 54 ECONOMETRICA

This literature, however, like the more traditional institutional competence literature, tends to neglect the related question of how the design of multi-agent systems affects the efficient production (or discovery) of information.

Many of the considerations discussed in Part II may also be present in settings involving multiple agents. The focus of Part III, however, is on additional considerations peculiar to multi-agent settings. The following three sections focus on broad categories of decisionmaking systems that involve contributions from several different agents. The first, and most straightforward, considers settings in which multiple agents are independently responsible for providing informational inputs into some decision, whether in the form of votes or nonbinding advice. The second section turns to settings in which two or more agents investigate a similar policy question sequentially rather than simultaneously. The third section turns to systems that seek to induce competition among government agents with respect to the provision of useful information.

A. Simultaneous Information Aggregation

A recent trend in both the scholarly and popular literatures on public decisionmaking emphasizes the advantages of aggregating dispersed information from a large number of parties, as an alternative (or supplement) to reliance on a smaller number of expert decisionmakers.[95] The basic insight underlying what has been popularly dubbed the "wisdom of crowds"[96] dates back at least to the Marquis de Condorcet's famous Jury Theorem, which demonstrates that a group of lay jurors who decide by majority rule can arrive at the correct answer more often than can a single expert.[97] "Many minds," according to one

---

923 (1986); David Martimort & Lars Stole, *The Revelation and Delegation Principles in Common Agency Games*, 70 ECONOMETRICA 1659 (2002); Takuro Yamashita, *Mechanism Games with Multiple Principals and Three or More Agents*, 78 ECONOMETRICA 791 (2010). For reasons of space and tractability, this Article does not consider endogenous information acquisition in multi-principal settings. In the real world, however, many important government decision problems may involve multiple principals, and a natural extension of this line of research would consider endogenous information acquisition in such settings.

[95] *See, e.g.*, SCOTT E. PAGE, THE DIFFERENCE (2007); SUNSTEIN, MANY MINDS, *supra* note 4; JAMES SUROWIECKI, THE WISDOM OF CROWDS (2004); JEREMY WALDRON, THE DIGNITY OF LEGISLATION 124–66 (1999).

[96] *See* SUROWIECKI, *supra* note 95, at xiv.

[97] *See* MARQUIS DE CONDORCET, AN ESSAY ON THE APPLICATION OF MATHEMATICS TO THE THEORY OF DECISION-MAKING (1785), *reprinted in* CONDORCET: SELECTED WRITINGS 33 (Keith M. Baker ed., 1976). For a succinct modern formalization (and extension) of the Condorcet Jury Theorem, see Krishna K. Ladha, *The Condorcet Jury Theorem, Free Speech, and Correlated Votes*, 36 AM. J. POL. SCI. 617 (1992). An excellent informal summary of the theorem and its most important extensions appears in Vermeule, *Many-Minds Arguments*, *supra* note 4, at 4–9.

influential modern formulation of Condorcet's insight, are often better than few.[98] The Condorcet Jury Theorem has inspired a range of recommendations for legal and institutional reform (as well as further justifications for existing institutions).[99] The common thread linking these analyses is the proposition that, at least for the types of decisions where the assumptions of the Condorcet Jury Theorem are likely to be satisfied, decisionmaking systems should strive to aggregate the dispersed information held by large numbers of parties.

The Condorcet Jury Theorem focuses on the aggregation of information through voting, and many of the modern applications of the theorem likewise focus on voting systems or settings involving aggregation of nonbinding advice. Institutional designers sometimes also seek contributions from multiple agents by designing systems with built-in institutional overlap or redundancy, both horizontal and vertical. ("Horizontal" redundancy refers to settings in which two or more separate agents are responsible for a similar task, such as producing information about a public policy problem. "Vertical" redundancy refers to systems of appeal and review, in which a hierarchically superior agent may scrutinize, and possibly revise, a decision made by a hierarchically inferior agent.) Redundant systems are thought to act as a form of insurance: if one agent fails in her task, another agent's contributions may compensate. Furthermore, if agents' contributions are partial rather than perfect substitutes (that is, if the agents' functions overlap but are not fully redundant), then the contributions from multiple agents may add value to the final outcome even if none of them shirk.[100]

This sort of institutional redundancy or overlap may be valuable in the production of policy-relevant information. Instead of charging one scientific panel with ascertaining the causes and likely consequences of global climate change, the government can assign this responsibility to multiple panels. Instead of relying on a single court to scrutinize the

---

[98] *See, e.g.*, SUNSTEIN, MANY MINDS, *supra* note 4.

[99] *See* SUNSTEIN, INFOTOPIA, *supra* note 4, at 197–216; SUNSTEIN, MANY MINDS, *supra* note 4 (explaining how many-minds arguments do and should play a role in constitutional interpretation); VERMEULE, *supra* note 4 (arguing that the Condorcet Jury Theorem and similar many-minds arguments suggest reasons to favor lawmaking by nonjudicial institutions); Michael Abramowicz, *En Banc Revisited*, 100 COLUM. L. REV. 1600, 1632–36 (2000); Lewis A. Kornhauser & Lawrence G. Sager, *Unpacking the Court*, 96 YALE L.J. 82, 96–100, 115–17 (1986); Levmore, *supra* note 4, at 734–39.

[100] *See, e.g.*, JONATHAN B. BENDOR, PARALLEL SYSTEMS 209–45 (1985); JAMES Q. WILSON, BUREAUCRACY 274 (1989); Robert M. Cover, *The Uses of Jurisdictional Redundancy: Interest, Ideology, and Innovation*, 22 WM. & MARY L. REV. 639 (1981); Katyal, *supra* note 93, at 2314–27; Martin Landau, *Redundancy, Rationality, and the Problem of Duplication and Overlap*, 29 PUB. ADMIN. REV. 346 (1969); O'Connell, *supra* note 4; Nancy Staudt, *Redundant Tax and Spending Programs*, 100 NW. U. L. REV. 1197, 1222–39 (2006); Michael M. Ting, *A Strategic Theory of Bureaucratic Redundancy*, 47 AM. J. POL. SCI. 274 (2003).

facts of a contested case, the judicial system might provide for a second look at the facts by an appellate panel. The President may consult the head of a cabinet department for policy advice but also create a "stovepipe" that allows the President to acquire information from other agents elsewhere in the bureaucracy, or may seek a second opinion from another government department with partly overlapping responsibilities.[101]

While many analysts have stressed the advantages of drawing on information from a larger number of agents, the enthusiasm for many-minds arguments has not gone unchallenged. Many scholars have pointed out the limits of the Condorcet Jury Theorem, particularly in cases where decisionmakers have strongly correlated biases,[102] where they vote strategically,[103] or where there is no underlying consensus on the "correct" answer.[104] Likewise, institutional redundancy has well-known costs, including wasteful duplication of effort and socially unproductive turf battles.[105] There is now a vigorous debate in the public law literature about the implications, if any, of many-minds arguments for the design of public institutions. Much (though certainly not all) of this debate focuses on how best to make use of information that various agents already possess, without considering the implications of different arrangements for the incentive to gather information. Yet this latter consideration may have important implications for questions regarding the optimal number of agents involved in providing informational inputs into a particular public decision, questions about the optimal collective decision rule for translating these separate inputs into a final policy choice, and questions regarding the optimal composition of multimember bodies. Let us consider each in turn.

1. The Number of Agents. — By far the most important implication that endogenous expertise may have for the many-minds debate —

---

[101] *See* Sean Gailmard & John W. Patty, Stovepiping (Aug. 27, 2010) (unpublished manuscript) (on file with the Harvard Law School Library); Vermeule, *Second Opinions*, *supra* note 4.

[102] *See* Ladha, *supra* note 97, at 625–30; Christian List & Philip Pettit, *An Epistemic Free-Riding Problem?*, *in* KARL POPPER: CRITICAL APPRAISALS 128 (Philip Catton & Graham Macdonald eds., 2004).

[103] *See* David Austen-Smith & Jeffrey S. Banks, *Information Aggregation, Rationality, and the Condorcet Jury Theorem*, 90 AM. POL. SCI. REV. 34 (1996).

[104] *See* Paul H. Edelman, *On Legal Interpretations of the Condorcet Jury Theorem*, 31 J. LEGAL STUD. 327, 337–39 (2002); N.W. Barber, *Two Meditations on the Thoughts of Many Minds*, 88 TEX. L. REV. 807, 817–18 (2010) (book review).

[105] *See generally* WILSON, *supra* note 100; Stuart Minor Benjamin & Arti K. Rai, *Fixing Innovation Policy: A Structural Perspective*, 77 GEO. WASH. L. REV. 1 (2008); Neal Devins & Michael Herz, *The Uneasy Case for Department of Justice Control of Federal Litigation*, 5 U. PA. J. CONST. L. 558 (2003); William V. Roth, Jr., *The "Malmanagement" Problem: Finding the Roots of Government Waste, Fraud, and Abuse*, 58 NOTRE DAME L. REV. 961 (1983); Nathan Alexander Sales, *Share and Share Alike: Intelligence Agencies and Information Sharing*, 78 GEO. WASH. L. REV. 279 (2010). The role of competition among agents is explored in greater detail in section III.C, pp. 1479–82.

and the one that has received the most attention in the extant literature — is that increasing the number of agents involved in researching a public decision problem tends to reduce each individual agent's incentive to invest heavily in doing that research. Information acquisition may often suffer from a collective action problem, in which each agent is tempted to free ride, shirking her own obligations to invest in costly research in the hopes that others will do enough of the work to produce a sufficiently high-quality final decision.[106] As is generally true with collective action problems, the more agents that are involved in the production of the public good (in this case, information), the stronger is the incentive to free ride.[107]

This collective action problem may provide a reason to be skeptical of wisdom-of-crowds arguments. The traditional Condorcet Jury Theorem framework, as well as most of its more modern formulations and permutations, assumes that each agent gets some "signal" of the correct answer to the question at issue and that the probability that the agent gets the right signal is exogenous (that is, unaffected by the agent's own strategic choices). While these assumptions may be reasonable in some settings, as when the relevant population already possesses information on the question at hand,[108] it is often more plausible that the quality of each agent's signal depends on the amount of effort she devotes to investigating the issue. As the number of agents involved in making a decision increases, the collective action problem outlined above intensifies, and the quality of each individual agent's information is likely to worsen. (This effect — known as the "rational ignorance" problem — has long been recognized in other contexts,

---

[106] The seminal modern work on the collective action problem is MANCUR OLSON, JR., THE LOGIC OF COLLECTIVE ACTION (1965). For applications of the collective action problem specifically to the production of information for public policymaking, see Brian Galle & Joseph Leahy, *Laboratories of Democracy? Policy Innovation in Decentralized Governments*, 58 EMORY L.J. 1333, 1346–71 (2009); Garrett, *supra* note 5, at 1509–13; Li, *supra* note 61; Nicola Persico, *Committee Design with Endogenous Information*, 71 REV. ECON. STUD. 165, 169–83 (2004); and Ting, *supra* note 100, at 277–88. *Cf.* Wolfgang Pesendorfer & Asher Wolinski, *Second Opinions and Price Competition: Inefficiency in the Market for Expert Advice*, 70 REV. ECON. STUD. 417 (2003).

[107] There may, however, be countervailing factors when agents have incentives to acquire information beyond their interest in making the correct collective decision. For example, if agents receive a private benefit from appearing well informed (as might occur when the appearance of expertise may improve the agent's reputation), then expanding the size of the group might sometimes *strengthen* each agent's incentive to acquire information, because a larger group means both more people within the group to impress and more rivals in the competition to impress outside observers. While it is worth acknowledging this possibility, it nonetheless seems plausible that in a broad range of decision contexts, the free rider problem will predominate, such that the incentive to invest in information decreases as the size of the decisionmaking body increases.

[108] *See* VERMEULE, *supra* note 4, at 88–89; César Martinelli, *Would Rational Voters Acquire Costly Information?*, 129 J. ECON. THEORY 225, 226–27 (2006).

such as voting by citizens and shareholders, and in private markets.[109]) So increasing the number of agents involved in making a decision presents a trade-off: doing so increases the *quantity* of signals received (which tends to increase the probability of making the correct decision) but also decreases the average *quality* of those signals (which tends to decrease the probability of making the correct decision).[110] Often the former effect will outweigh the latter effect, but under plausible circumstances — particularly those in which each agent must pay some fixed cost to get any sort of informative signal about a policy problem — increasing the number of agents involved may actually decrease the overall quality of information.[111]

The collective action problem also has implications for the use of horizontal institutional redundancy in the production of information. As noted above, the main advantage of this sort of redundancy is that it increases the odds that even if one agent fails to complete her assigned research task adequately, another agent will pick up the slack. The problem is that when agents must invest costly effort in order to produce useful research, and two or more agents are responsible for producing the same sort of information (or when the types of information each one produces are sufficiently similar to function as partial substitutes), the collective action problem tends to induce suboptimal research effort.[112] Institutional redundancy thus gives rise to a trade-off that parallels the trade-off present when information aggregation takes the form of voting: greater systemic redundancy increases the probability of effective information production when the research effort of the individual agents is taken as given, but for precisely that reason redundancy tends to reduce the research effort of each individual agent and may worsen the aggregate production of information.

The above analysis suggests, as a general rule of thumb, that legal or institutional systems that draw on the information of a larger number of agents, whatever their other advantages, run the risk of weaken-

---

[109] *See* ANTHONY DOWNS, AN ECONOMIC THEORY OF DEMOCRACY (1957); FRANK H. EASTERBROOK & DANIEL R. FISCHEL, THE ECONOMIC STRUCTURE OF CORPORATE LAW 83 (1991); George J. Stigler, *The Economics of Information*, 69 J. POL. ECON. 213 (1961).

[110] *See* Drora Karotkin & Jacob Paroush, *Optimum Committee Size: Quality-Versus-Quantity Dilemma*, 20 SOC. CHOICE & WELFARE 429, 429–34 (2003); List & Pettit, *supra* note 102, at 128; Kaushik Mukhopadhaya, *Jury Size and the Free Rider Problem*, 19 J.L. ECON. & ORG. 24 (2003); Persico, *supra* note 106, at 167–82; Vermeule, *Many-Minds Arguments*, *supra* note 4, at 28–29. This insight dates back at least to Jeremy Bentham. *See* JON ELSTER, EXPLAINING SOCIAL BEHAVIOR 413 (2007) (quoting Bentham as asserting that "[t]he greater the number of voters the less the weight and the value of each vote, the less its price in the eyes of the voter, and the less of an incentive he has in assuring that it conforms to the true end and even in casting it at all" (internal quotation marks omitted)).

[111] *See* Minoru Kitahara & Yohei Sekiguchi, *Condorcet Jury Theorem or Rational Ignorance*, 10 J. PUB. ECON. THEORY 281, 292 (2008); Martinelli, *supra* note 108, at 226.

[112] *See* Ting, *supra* note 100, at 275–77.

ing agents' research incentives. To put the problem in the terminology developed earlier in this Article, when multiple agents are responsible for providing decision-relevant information (whether via voting, advice giving, or some other mechanism), the difference between each agent's research payoff and default payoff tends to shrink. Whether the net effect on aggregate information acquisition is positive or negative will depend on context, but there are at least some situations in which the exacerbation of the rational ignorance problem outweighs the informational benefits of making use of many minds.

This bleak assessment may not apply, however, when the types of information provided by different agents to a single decisionmaker are complements rather than substitutes. (Pieces of information are substitutes if the possession of one piece of information *decreases* the marginal value of acquiring the second piece. Pieces of information are complements when the possession of one piece of information *increases* the marginal value of acquiring the second piece.) Indeed, when effective research entails generating complementary forms of information (that is, where the informational whole is greater than the sum of its parts), dividing complementary research tasks among different agents may enhance rather than undermine each agent's research incentives.[113] Consider a setting in which a government decision turns on two complementary pieces of information: the expected value (for both the principal and all agents) of acquiring, and basing the decision on, both pieces of information is 10 units of utility, while the expected value of acquiring one piece of information, but not both, is 3 units.[114] Suppose further that the private cost to an agent of acquiring one of these pieces of information is 6 units, so that the total private cost of acquiring both pieces is 12 units. If both research tasks are assigned to one agent, that agent will not research either type of evidence. However, assigning the two research tasks to two separate agents can induce research effort: it is in each agent's interest to invest in research if she expects the other to do so; thus, there is an equilibrium in which the government can induce optimal research investment by dividing

---

[113] *See* Persico, *supra* note 106, at 167; Ting, *supra* note 100, at 284–87.

[114] For an example of pieces of information that might be complementary, consider a hypothetical municipal planning board deciding whether to implement a given land use regulation or local public works project. The board might be uncertain about whether the regulation will increase or decrease housing density and whether social welfare overall is positively or negatively correlated with housing density. If the board has good information on both issues, it is well equipped to make the right decision. But having good information on only one of the two questions is not very useful: even if the board knows that the regulation will increase (or decrease) housing density, the board cannot make a good urban planning decision if it remains unsure whether it should be trying to increase or decrease density; likewise, even if the board knows that it should be trying to increase (or decrease) density, this information is of little use if the board remains ignorant regarding whether the regulation at issue will increase or decrease density.

the investigation of complementary questions among different agents, even though it is impossible to achieve this result if these tasks are all a single agent's responsibility.[115]

Although the above example is abstract, it illustrates a significant, if subtle, substantive point: the impact of dividing research tasks among multiple agents may depend critically on whether the types of information they are charged with producing are substitutes or complements.[116]　When different types of information are substitutes — when the marginal value of one type of information declines as more information of other types becomes available — then dividing the responsibility for acquiring information, whatever its other advantages, will tend to dampen research incentives. In contrast, dividing responsibility for researching complementary types of information may strengthen research incentives. It seems reasonable to conjecture that in most cases the former scenario is more likely, which would imply that it makes sense to begin with the presumption, prevalent in the extant literature, that increasing the number of agents tends to undermine individual agents' research incentives. Nonetheless, this presumption must be qualified by the recognition that division of tasks may in fact enhance research incentives in those cases where different types of information are complementary.

2. The Collective Decision Rule. — For reasons similar to those discussed in the preceding section, the endogeneity of agents' information may complicate the analysis of the voting rule that collective bodies ought to employ. For instance, an important strain of scholarship has identified significant advantages to supermajority rules under certain circumstances, in particular, circumstances in which the institutional designer would like to build in a bias in favor of the status quo (or some other default outcome).[117]　A supermajority rule has the at-

---

[115] When tasks are divided, there is also an equilibrium in which neither agent invests in research effort; if one agent does not invest in research, the optimal response of the other agent is not to do research either. In order to determine which equilibrium the agents would select, one would need to appeal to other criteria — perhaps including the fact that the equilibrium in which both invest in research yields a higher payoff to both of them than does the equilibrium in which neither invests in research. Likewise, if the agents make their research decisions sequentially rather than simultaneously, then the *only* equilibrium is the one in which both agents invest in research. The second agent to move will mimic the first agent's choice, which means that the first agent will choose to invest in research, knowing that the second will rationally follow suit.

[116] *See* Ting, *supra* note 100, at 287–88.

[117] *See, e.g.*, Jacob E. Gersen & Adrian Vermeule, Chevron *as a Voting Rule*, 116 YALE L.J. 676, 679–88, 708–12 (2007); John O. McGinnis & Michael Rappaport, *The Condorcet Case for Supermajority Rules*, 16 SUP. CT. ECON. REV. 67 (2008); John O. McGinnis & Michael B. Rappaport, *Majority and Supermajority Rules: Three Views of the Capitol*, 85 TEX. L. REV. 1115 (2007); Jed Handelsman Shugerman, *A Six–Three Rule: Reviving Consensus and Deference on the Supreme Court*, 37 GA. L. REV. 893, 954–62 (2003); *see also* Ruth C. Ben-Yashar & Shmuel I. Nitzan, *The Optimal Decision Rule for Fixed-Size Committees in Dichotomous Choice Situations:*

tractive feature of allowing a change from the default policy only if a sufficiently large number of agents have independently concluded the change is a good idea, which would imply a particularly high degree of confidence that the change is justified.

This argument for supermajority rules typically presumes that each voter votes independently and sincerely — assumptions that may not always hold true.[118] Even if these assumptions do hold, however, the benefits of supermajoritarianism can be undermined when the quality of each voter's information is endogenous: holding group size constant, as the size of the supermajority required to effect a policy change becomes larger, the probability that any individual will cast the decisive vote will change and may often decrease. This possibility, in turn, implies that if each agent's information (or a substantial portion of it) remains private, the difference between the research payoff and default payoff shrinks for each agent as the supermajoritarian threshold for policy change approaches unanimity.[119] Thus, changing the collective decision rule — say, by switching from a simple majority rule to a supermajority rule — may reduce the accuracy of the decision by reducing aggregate investment in information.

It is easiest to illustrate this effect with an extreme case. Suppose that a large committee of homogeneous agents may vote (simultaneously and by secret ballot) to adopt a new policy. In keeping with the traditional arguments in favor of supermajority rules, let us suppose (despite the problematic nature of this supposition) that each agent casts her vote independently and sincerely. Let us assume further (and in contrast with the more conventional analyses of supermajoritarian voting rules) that information is endogenous: each agent is initially uninformed but can improve her probability of casting a correct vote by investing costly effort in research. Assume that the results of each individual agent's research remain her private information. Thus, each agent's research is valuable to her only if she ends up casting the pivotal vote. The probability of casting the deciding vote under a simple majority rule is not very high, but at least if the group is not too large, it is not trivial. This probability gives each agent an incentive to do at least some research. Under a unanimity rule, however, the probability that any given agent will cast the deciding vote is likely

The General Result, 38 INT'L ECON. REV. 175, 181–83 (1997) (providing a general model that establishes conditions under which a supermajority rule is optimal).

[118] *See, e.g.*, Timothy Feddersen & Wolfgang Pesendorfer, *Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts Under Strategic Voting*, 92 AM. POL. SCI. REV. 23, 23, 31–32 (1998).

[119] *See* Persico, *supra* note 106, at 165–68; *see also* David Austen-Smith & Timothy Feddersen, *Deliberation and Voting Rules*, *in* SOCIAL CHOICE AND STRATEGIC DECISIONS 269 (David Austen-Smith & John Duggan eds., 2005).

to be much lower. Unless the quality of each agent's signal is very high (such that the signals exhibit an extremely strong positive correlation), the likelihood (from the standpoint of any given agent) that all other agents will vote for the policy change is quite low, which makes her vote likely irrelevant to the outcome.

For example, consider a committee of three voters, *A*, *B*, and *C*. Suppose that the ex ante probability that a proposed new policy is better than the status quo is 50%. To keep the example simple, let us focus on the research investment decision of agent *A*. We will assume that *B*'s and *C*'s research investments are exogenous; both *B* and *C* have a 75% chance of correctly assessing whether the new policy is superior to the status quo; they make these assessments independently; and they vote in accordance with their individual assessments. Thus, the probability that both *B* and *C* vote for the change (making *A*'s vote pivotal under a unanimity rule but meaningless under a simple majority rule) is 31.25%. The probability that *B* and *C* split their votes (making *A*'s vote pivotal under a simple majority rule but meaningless under a unanimity rule) is 37.5%.[120] Thus, in this example *A* would invest more in research under a simple majority rule than she would under a unanimity rule.[121]

Of course, *B* and *C* are likely also to make endogenous decisions regarding research effort. To fully solve for the equilibrium of even this simple model, one would need to derive a research investment level for all three agents such that none would have an incentive to deviate. Moreover, the preceding analysis assumed that agents vote sincerely and independently. While these assumptions are useful simplifications and are common in the extant legal literature on supermajority rules, they are likely wrong in many cases. Relaxing the assumptions of sincere and independent voting may substantially alter predictions about equilibrium behavior.[122] A full analysis of the impact of supermajority rules would also need to take these considerations into account.

---

[120] Call a "Y" vote a vote for the new policy and an "N" vote a vote for the status quo, and let p be the probability that agents *B* and *C* correctly assess whether it is a good idea to adopt the new policy. Ex ante, agent *A* recognizes that there are four possible ways that the other two agents might vote (YY, YN, NY, and NN), and the respective probabilities of each of these scenarios, given the assumptions of the example, are $0.5 \times (p^2 + (1 - p)^2)$, $p \times (1 - p)$, $p \times (1 - p)$, and $0.5 \times (p^2 + (1 - p)^2)$. Under simple majority rule, agent *A*'s vote is pivotal in the second and third scenarios, but not in the first or fourth. The values in the text are generated by setting $p = 0.75$.

[121] This result, however, would not hold for very high levels of accuracy. To illustrate, consider an extreme case in which *B* and *C* get a signal that is 100% accurate ($p = 1$). In that case, *A* can never cast a deciding vote under a simple majority rule (because the other two will always vote either YY or NN), but may well cast a deciding vote under a unanimity rule (if *B* and *C* vote YY).

[122] *See, e.g.*, Austen-Smith & Banks, *supra* note 103; Feddersen & Pesendorfer, *supra* note 118.

We need not pursue that full analysis here, however, as the preceding discussion merely illustrates the more general point that a change in the collective decision rule, by altering each agent's probability of being pivotal, may alter each agent's research incentive. The most interesting potential consequence of this insight for the existing legal literature on supermajority rules is that, whatever their other virtues and vices, supermajority rules may induce a substantial reduction in both individual and aggregate information production.[123]

A crucial feature of the preceding analysis, however, is the assumption that the information each agent acquires remains her private information. Sometimes, though, costly information gathering improves the accuracy of every agent in the collective decisionmaking body. In such situations, a supermajority rule may actually strengthen research incentives, because the pivotal voter in the group will be more skeptical of the proposed policy than the median voter will be. Thus, those agents who favor a particular change will not be able to stop their research efforts and proceed to a vote once they have enough information to convince a majority of the group; rather, they will have to continue to invest in information gathering until they have enough information to convince the requisite supermajority.[124] In other words, "veto power or unanimity represents a constraint that induces deliberation: when parties can block outcomes, actors have an incentive to find reasons which are convincing to all, not just to the majority."[125] (Note that in this case, the supermajority requirement functions as a more stringent standard of proof and induces information acquisition in a similar fashion.[126]) Thus, supermajority rules not only may affect agents' incentives to invest in information, but also may have differential effects on their incentives to acquire different *kinds* of information. More demanding supermajority requirements may reduce agents' incentives to invest in the production of *private* information, as supermajority rules tend to decrease each agent's probability of casting the pivotal vote, but more demanding supermajority requirements may simultaneously increase agents' incentives to invest in *public* or *shared* information, because agents sympathetic to a proposed change will need to work harder to win over the (more skeptical) pivotal voter.

3. Preference Heterogeneity. — The preceding discussion implicitly assumed that the different agents involved had similar preferences.

---

[123] *See* Persico, *supra* note 106, at 165–66.

[124] *See* Alessandro Lizzeri & Leeat Yariv, Sequential Deliberation 21 (Nov. 4, 2010) (unpublished manuscript) (on file with the Harvard Law School Library).

[125] Erik Oddvar Eriksen, *Democratic or Technocratic Governance?* 15 (NYU Sch. of Law Jean Monnet Working Papers, Paper No. 6/01, 2001), *available at* http://centers.law.nyu.edu/jeanmonnet/papers/01/011201.html.

[126] *See supra* pp. 1447–49.

Yet one of the issues that has attracted the most attention from legal scholars interested in collective decisionmaking is the role of diversity, particularly preference heterogeneity. Indeed, much of the institutional design literature hypothesizes that increasing diversity on multimember decisionmaking bodies tends to lead both to more moderation and to more accuracy.[127] The literature has suggested a variety of reasons for this pattern, including the hypotheses that diverse groups are less likely to exhibit correlated biases or groupthink,[128] that diverse groups tend to seek consensus and avoid dissent,[129] and that minority members can act as whistleblowers, threatening to disclose malfeasance or sloppiness to the principal (or some other monitor).[130]

The fact that information acquisition is often endogenous may supply another perspective on the differences between the behaviors of ideologically diverse and ideologically homogeneous groups, as well as another tentative reason to favor diversity — particularly in the form of including the representation of minority views — in collective decisionmaking. Certain forms of ideological diversity may create stronger incentives for agents to acquire information, and this effect may imply that ideologically diverse bodies are sometimes superior not only to bodies whose members exhibit strongly correlated biases, but also to homogeneous bodies whose members are entirely unbiased.[131] The reason is that a minority faction's default payoff tends to be systematically worse than the majority faction's, and this payoff structure gives the minority faction a stronger incentive to invest in information, in the hopes of changing the majority's mind.[132]

---

[127] *See* PAGE, *supra* note 95, at 153–58; CASS R. SUNSTEIN ET AL., ARE JUDGES POLITICAL? AN EMPIRICAL ANALYSIS OF THE FEDERAL JUDICIARY 135–40 (2006); Cass R. Sunstein, *Deliberative Trouble? Why Groups Go to Extremes*, 110 YALE L.J. 71, 92–93, 108–09 (2000).

[128] *See* PAGE, *supra* note 95, at 48–50; CASS R. SUNSTEIN, WHY SOCIETIES NEED DISSENT 140–44 (2003); Ladha, *supra* note 97, at 627.

[129] *See* Harry T. Edwards, *The Effects of Collegiality on Judicial Decision Making*, 151 U. PA. L. REV. 1639, 1651 (2003); Joshua B. Fischman, *Understanding Voting Behavior in Circuit Court Panels* 38–42 (Univ. of Va. Sch. of Law Pub. Law & Legal Theory Research Paper Series, Paper No. 2010-27, 2010), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1636002; Joshua B. Fischman, Decision-Making Under a Norm of Consensus: A Structural Analysis of Three-Judge Panels 15 (Jan. 4, 2008) (unpublished manuscript), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=912299; *cf.* RICHARD A. POSNER, HOW JUDGES THINK 32–34 (2008) (describing how appellate judges' reluctance to dissent can produce consensus even on ideologically diverse panels).

[130] *See* Frank B. Cross & Emerson H. Tiller, *Judicial Partisanship and Obedience to Legal Doctrine: Whistleblowing on the Federal Courts of Appeals*, 107 YALE L.J. 2155 (1998); Jonathan P. Kastellec, *Panel Composition and Judicial Compliance on the U.S. Courts of Appeals*, 23 J.L. ECON. & ORG. 421 (2007).

[131] *See* Matthew Spitzer & Eric Talley, Left, Right, and Center: Strategic Information Acquisition and Diversity in Judicial Panels (Aug. 23, 2010) (unpublished manuscript), *available at* http://www.law.northwestern.edu/colloquium/law_economics/documents/Talley2010.pdf.

[132] *See id.* at 19.

To illustrate, consider a hypothetical three-judge appellate panel. Assume that each of the three judges may be one of two types, which for simplicity we will call Democrat and Republican. The judges have to make some collective decision, by majority rule, whether to overturn a lower court decision. Assume that, on a cursory analysis of the decision below (focusing on, for example, information about the identities of the parties, the reputation of the lower court judge, and the nature of the issue), a Democratic judge would think there is a 75% chance that the lower court decision was correct, while a Republican judge would put the probability of a correct decision below at 25%. Each judge receives a payoff of 100 units of utility if the panel decides the appeal correctly (upholding a correct lower court decision or reversing an incorrect decision) and receives a payoff of 0 otherwise. Each judge can also invest additional time and effort into a more sophisticated analysis of the case that will identify the legally correct answer (and do so in a way that is verifiable by her colleagues). Such effort would cost the judge 40 units of utility.

In a homogeneous panel, the judges will tend to agree on what appears to be the best decision. The incentive of any one of them to expend additional effort to ascertain the correct answer is relatively low, for two reasons. First, it is unlikely that additional information will lead to a different decision. (Put another way, each judge's default payoff is high relative to her research payoff.) Second, the collective action problem means that each of the three judges has an incentive to free ride on the others' efforts. So, for example, on a panel consisting of three Democrats, each judge's expected payoff from affirming the decision below is 75. If any one judge invests extra effort to ascertain the right answer, the judges who free ride get a payoff of 100, while the judge who does the work gets a payoff of 100 − 40 = 60. Thus, no judge is willing to do the extra work, and the work will not get done.[133]

Now consider a panel with two Democratic judges and one Republican judge (or vice versa). The minority judge has considerably stronger incentives to invest in additional research to attempt to persuade her colleagues to change their minds because the minority judge's default payoff (25) is low relative to her research payoff (100). This phenomenon is similar to the setting discussed in section II.C.1, in which the principal reviewed an agent's decision and the principal and agent had preferences that diverged substantially in low-information environments but converged in high-information environments. In that setting, the agent's incentive to invest in research was higher than in a case where the agent was a closer ally of the

---

[133] The assumption here, of course, is that the judges cannot solve the collective action problem by cooperating, for example, by sharing the workload.

principal, with preferences that were similar in both low- and high-information environments. The reason was that an agent that disagreed with the principal about what to do in a low-information environment had a greater interest in finding more information. In the case of multi-judge panels (or multi-agent panels more generally), the majority faction is analogous to the principal in the earlier example, while the minority faction is analogous to the agent.

The analogy to the earlier case also highlights that this positive effect of ideological diversity may not obtain when the preference divergence between the majority and minority factions on the panel remains constant (or increases) as the amount of information increases. In these settings, ideological diversity might actually *dampen* research incentives. After all, if the minority judge knows that any information she provides might be used by the majority faction to craft a ruling that is even more distant from her ideal ruling, then she has considerably weaker incentives to do research. A panel composed of homogeneous judges, while still subject to a collective action problem, would not suffer from this difficulty, and the individual judges might invest more in research than would a minority judge who would view the case as a lost cause from the beginning and refuse to add anything to the discussion.[134]

B. Sequential Information Aggregation

When crafting a system in which multiple agents may supply or use information in making a policy decision, institutional designers must often take into account the timing of the agents' inputs. While this inquiry may entail a variety of considerations, a common general issue is whether agents' contributions are simultaneous or sequential. This timing decision, as relevant here, determines whether some agents can observe other agents' decisions and/or information prior to making their own decisions.[135] (In that sense, a system in which agents make their decisions in chronological sequence, but later-moving agents remain ignorant of prior agents' decisions, would be considered simultaneous rather than sequential for purposes of the present discussion.)

While many decisionmaking systems rely principally on simultaneous contributions (for example, many voting systems and some forms

---

[134] This characterization of the minority judge's incentives would be less apt in contexts where a dissent — especially a well-researched and well-reasoned dissent — could have substantial influence with a broader audience or on future decisions, even if it did not alter the outcome of the present case. *See* Cross & Tiller, *supra* note 130, at 2159 & n.17. Thus, the incentive effects described in this section are likely to be relatively weaker in such settings, though they may still exist to some degree.

[135] *See* Aghion & Tirole, *supra* note 53, at 6–7; Vermeule, *Many-Minds Arguments*, *supra* note 4, at 31–33.

of horizontal redundancy), others employ sequential aggregation of information. For example, appeals systems (and other forms of hierarchical oversight) often involve sequential investigation of the same issue, first by a front-line agent (such as a lower court or programmatic agency), then by an oversight agent (such as an appellate court or a centralized regulatory review body). Likewise, systems that develop bodies of law or policy over time, through something like a precedent-based common law approach, can be thought of as entailing a kind of sequential aggregation of information. These systems differ in that hierarchical oversight involves subsequent review of the *same* decision, whereas the development of common law involves a series of *different* (but related) decisions. Nonetheless, these systems share the common feature that, in contrast to secret-ballot voting and other forms of simultaneous information aggregation, appeals systems and common law systems allow later decisionmakers to incorporate the information contained in earlier decisions.

This additional information can be an advantage. Indeed, systems that operate according to common law principles, or something similar, are often celebrated for their ability to draw on the accumulated wisdom of past experience.[136] And much of the literature on hierarchical review similarly notes the fact that the reviewing agent can use the decision below as an important source of information when making its ruling.[137] Sequential decisionmaking systems, however, are also vulnerable to certain well-known pathologies, chief among them "information cascades" (sometimes also known as "herding"), in which subsequent decisionmakers ignore their own information in favor of following prior judgments, on the logic that there is more information contained in the earlier decision (or decisions) than in the present decisionmaker's own independent assessment of the evidence.[138] The

---

[136] *See, e.g.*, SUNSTEIN, MANY MINDS, *supra* note 4, at 36–55, 107–11; David A. Strauss, *Common Law Constitutional Interpretation*, 63 U. CHI. L. REV. 877, 891–92 (1996).

[137] *See* Fox & Stephenson, *supra* note 94 (manuscript at 22–23); Fox & Van Weelden, *supra* note 94, at 679–81, 683; Vermeule, *Second Opinions*, *supra* note 4, at 8; *cf.* Charles M. Cameron et al., *Strategic Auditing in a Political Hierarchy: An Informational Model of the Supreme Court's Certiorari Decisions*, 94 AM. POL. SCI. REV. 101, 104 (2000) (analyzing a setting in which the overseer may draw inferences from the first agent's action when deciding whether to review that action, but the overseer's decision on the merits is independent of the first agent's choice); Spitzer & Talley, *supra* note 77, at 658 (same).

[138] *See* Abhijit V. Banerjee, *A Simple Model of Herd Behavior*, 107 Q.J. ECON. 797, 798 (1992); Sushil Bikhchandani et al., *A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades*, 100 J. POL. ECON. 992, 1001 (1992); Robert J. Shiller, *Conversation, Information, and Herd Behavior*, 85 AM. ECON. REV. 181, 181 (1995). For discussions of the information cascade problem specifically in the context of public law and institutional design, see, for example, SUNSTEIN, MANY MINDS, *supra* note 4, at 171–73; Andrew F. Daughety & Jennifer F. Reinganum, *Stampede to Judgment: Persuasive Influence and Herding Behavior by Courts*, 1 AM. L. & ECON. REV. 158, 180–82 (1999); Kuran & Sunstein, *supra* note 88, at 746–60; Eric Talley, *Precedential Cascades: An Appraisal*, 73 S. CAL. L. REV. 87, 121–24 (1999); Vermeule, *Many-Minds*

problem with an information cascade is that it can stifle the further aggregation of useful information. There is an extensive literature, in both law and social science, on the information cascade problem, as well as other difficulties that may arise from strategic interactions in sequential decision settings.[139] The limited aim of this section is to suggest some additional considerations that may be relevant when each agent's research effort is endogenous.

Consider a setting involving the simplest form of sequential decisionmaking by multiple agents: a two-period hierarchical oversight system in which the first agent (the "agency") makes some decision (say, whether to regulate a potentially hazardous substance), and a second agent (the "overseer") decides whether to uphold, reverse, or modify that decision. Note that this setup is quite similar to the setting considered in section II.C, in which the overseer could review, and potentially reject or revise, the decision made by the agent. The difference is that in the earlier example the overseer could not make its own independent research investment; the overseer observed only the agent's information (II.C.1), only her research effort (II.C.3), or nothing at all (II.C.2). In the present setting, the overseer may also acquire information, but doing so is costly.

The first important observation about this setting is that the endogeneity of research effort may intensify the information cascade problem. In the basic version of the problem, a later-moving agent may ignore her own assessment of the evidence if it contradicts the information implicit in the earlier decision (or decisions). But if the later-moving agent's information is good enough, she will still use it. In the two-period oversight context sketched above, this possibility means that the overseer will "defer" to the agency's judgment if, but only if, the agency's signal is believed to be of sufficiently higher quality than the overseer's. If, however, the overseer must invest costly effort in order to learn useful information, the overseer might rationally prefer not to do her own research and to defer instead to the agency's decision. Put another way, if the earlier decision is sufficiently informative, the overseer's default payoff is relatively high, which reduces her research incentive.

The same consideration applies even when the overseer would always make at least some use of her own information, for example, by changing policy incrementally. When incremental policy change is

---

Arguments, *supra* note 4, at 31–33; Vermeule, *Parliament*, *supra* note 4, at 2253–56; and Vermeule, *Second Opinions*, *supra* note 4, at 15. A similar phenomenon — that of "yes men" — may arise when the first-moving agent can anticipate what the later-moving agent will believe and has a material incentive to report the result that the later-moving agent is most likely to view as correct. *See* Prendergast, *supra* note 24, at 757–58.

[139] *See* sources cited *supra* notes 137–138.

possible, the standard version of the information cascade problem is mitigated, since later-moving agents can continue to make rational use of their information; that their own information may not lead to very large changes in policy is immaterial, as that information is still efficiently aggregated with the information contained in prior decisions.[140] If each agent's signal is exogenous, one might conclude that so long as agents are not constrained to make crude categorical decisions, information cascades are not a serious problem. That optimistic conclusion is drawn into question when each agent's research effort is endogenous: if later-moving agents recognize that their information will lead only to very small incremental policy changes, they may not have sufficiently strong incentives to invest in high-quality information.[141] This effect gives rise to a vicious cycle, as low-quality information is even less likely to result in meaningful policy changes.

The above discussion suggests (tentatively) that endogenous information acquisition may exacerbate the information cascade problem by reducing later-moving agents' incentives to invest in research. Endogenous information acquisition may, however, have quite different implications for the behavior of early-moving agents. One possibility is that if the agency can anticipate that the overseer will defer to its decision, then the agency may have a greater incentive to invest substantial effort in research than it would if it could be confident that the overseer would provide a meaningful second look. The intuition here is that the agency knows it is operating without a safety net and so is likely to devote more effort to information acquisition than it would if the overseer's research effort were exogenous.[142] Thus, a sequential decision setting may induce an equilibrium in which there is more intense investigation at the first stage, and laxer investigation at the

---

[140] *Cf.* Ethan Bueno de Mesquita & Matthew Stephenson, *Informative Precedent and Intrajudicial Communication*, 96 AM. POL. SCI. REV. 755, 764 (2002) (developing a model in which a judge can change precedent incrementally in order to preserve the informational value of the prior line of precedent while also shifting the legal rule in the direction of the judge's ideal).

[141] *Cf.* VERMEULE, *supra* note 4, at 75–77 (arguing that judges have an incentive to conserve decisionmaking costs by following precedent, which in turn can lead to cascades and informational free riding).

[142] The Supreme Court has adverted to (though not fully embraced) similar logic to explain the value of judicial deference to a magistrate judge's decision that probable cause existed to support a warrant application, when the court is called upon after the fact to assess that determination. *See* Franks v. Delaware, 438 U.S. 154, 167 (1978) ("The less final, and less deference paid to, the magistrate's determination of [the] veracity [of the statements contained in the warrant application], the less initiative will he use in that task. Denigration of the magistrate's function would be imprudent insofar as his scrutiny is the last bulwark preventing any particular invasion of privacy before it happens."). This example is especially interesting in that an important action — the search — may occur in between the decisions of the first and second agents. This timing may suggest an even greater cost to diminishing the first-moving agent's incentive to do costly investigation.

second stage, than one would observe in a world where information is exogenous.

There is a second possibility, however. Suppose the overseer can observe not only the agency's decision, but also the agency's research effort. If the agency invests very little in research, the overseer's optimal response will be to discount the agency's decision and to do its own investigation of the problem. Because the agency can anticipate this response, it may be rational for the agency to do little or no research, because the agency knows the overseer will pick up the slack. Thus, instead of an equilibrium in which the agency invests heavily in research and the overseer acts as a rubber stamp, there may be an equilibrium in which the agency shirks its responsibilities and the overseer ends up doing all the work.[143] Both of these equilibria deviate considerably from the "second look" or "accumulated wisdom" idea that often undergirds arguments for sequential decisionmaking (whether in the form of oversight or in the form of common law decisionmaking).[144]

The above logic further suggests a nonmonotonic relationship between the degree of preference divergence among agents and the amount that each agent will invest in research. When the preferences of the agency and the overseer are closely aligned, the agency knows that if it shirks on research, the overseer will pick up the slack and select a policy that is very close to the policy that the agency would prefer if the agency were well informed. When the preferences of the agency and the overseer are far apart, the agency knows that even if it invests extensively in research, the overseer will not trust the agency's policy judgment but will instead conduct independent research into

---

[143] One illustration of this sort of phenomenon is Professor James Bradley Thayer's argument that judicial review causes legislators to pay less attention to investigating and deliberating about potential constitutional problems with proposed legislation. *See* James B. Thayer, *The Origin and Scope of the American Doctrine of Constitutional Law*, 7 HARV. L. REV. 129, 155–56 (1893); *see also* MARK TUSHNET, TAKING THE CONSTITUTION AWAY FROM THE COURTS 57–58 (1999); VERMEULE, *supra* note 2, at 261–62; Vermeule, *Second Opinions*, *supra* note 4, at 28. This problem might be attenuated, however, if the legislature suffers some reputational cost when a court reverses the legislature's decision. *See* Fox & Stephenson, *supra* note 94 (manuscript at 17–20).

[144] Modern expositions of the idea that the common law reflects an accumulated stock of wisdom typically presume that each successive judge's decision is informed by new information that is integrated into the line of precedent. *See, e.g.*, SUNSTEIN, MANY MINDS, *supra* note 4, at 36–41; Lawrence B. Solum, *The Supreme Court in Bondage: Constitutional Stare Decisis, Legal Formalism, and the Future of Unenumerated Rights*, 9 U. PA. J. CONST. L. 155, 191–92 (2006); Strauss, *supra* note 136, at 891–92; Todd J. Zywicki & Anthony B. Sanders, *Posner, Hayek, and the Economic Analysis of Law*, 93 IOWA L. REV. 559, 581–83 (2008). Likewise, most analyses of the "error-correction" benefits associated with hierarchical appellate systems assume that the accuracy of the information of both the initial evaluator (for example, the trial court) and the overseer (for example, the appellate court) is exogenous; the focus of these analyses is typically on the selection of disputes for appeal. *See, e.g.*, Cameron & Kornhauser, *supra* note 77, at 178; Daughety & Reinganum, *supra* note 77, at 512–13; Spitzer & Talley, *supra* note 77, at 658, 660–62.

the issue.[145] The agency will therefore again shirk on research, leaving both the research work and the ultimate decision to the overseer. When the preferences of the agency and the overseer diverge somewhat, but not too much, the outcome may be quite different: In this case, although the overseer knows that its preferences differ from those of the agency, the overseer will be willing to defer to the agency's decision so long as the agency invests extensively in researching the issue. Although the overseer can improve the expected quality of the final decision by doing its own research, the research costs may not justify the expected gains. The agency, knowing this, will recognize that it can get its ideal policy outcome if it invests sufficient research effort. Under some circumstances, an agency with policy preferences that diverge moderately from those of the overseer will be willing to bear the research costs in order to secure a more favorable policy outcome, while the overseer will be willing to accept a somewhat less desirable policy outcome in order to avoid the costs of research. If the overseer is the principal (or a perfect agent of the principal), the above analysis implies another deviation from the ally principle: the principal may prefer an agent with moderately divergent preferences over a perfect ally, so that the agent, rather than the principal, will bear the research costs associated with the policy decision.

C. *Competitive Information Provision*

To this point, the analysis of multi-agent systems has emphasized the collective action problem that may arise when several agents are responsible for producing decision-relevant information. If it were possible to compensate agents directly for their research efforts or to base their compensation on the observed quality of the final decision, then mitigating this problem would be relatively straightforward. But as discussed earlier, in many public policy contexts this is not feasible. However, it may be feasible to base each agent's reward on the *content* of the final decision, and to do so in a way that induces desirable competition among agents to produce useful information. There are two main types of competitive systems that an institutional designer interested in improving research incentives might employ: advocacy systems and tournament systems.[146]

---

[145] *See* Gailmard, *supra* note 85, at 550–53.

[146] This Article focuses on institutional arrangements that seek to induce competition among government agents, bracketing a related set of questions about how governments might harness information dispersed among private citizens through competitive market mechanisms (such as privatization of certain government functions or the expanded use of prediction markets). *See, e.g.*, ABRAMOWICZ, *supra* note 30, at 283–84 (discussing how prediction markets can improve government decisionmaking). It is again worth noting that the government's decision of whether or how to exploit private competition to improve public policy is itself a public decision that may require research into the efficacy of market mechanisms.

In an advocacy system, the principal (or an adjudicator acting on the principal's behalf) specifies an array of possible policy choices and assigns a different agent to advocate for each one. An advocate receives a reward if the adjudicator selects the alternative that the advocate represents.[147] As long as each advocate can improve her odds of winning (that is, of having her policy alternative adopted) by providing useful information that supports her side, then an advocacy system can induce greater research than would occur if a single "neutral" agent were charged with investigating the issue. This advantage may sometimes outweigh the disadvantages associated with turning key aspects of the policy process over to "biased" advocates.[148] Although the neutral agent will weigh the evidence appropriately ex post, she will also recognize that there is some chance the evidence she acquires will point in different directions. Such conflicting evidence would lead her to make approximately the same decision she would have made had she not invested so much in research, thus dampening her incentive to conduct extensive research in the first place.[149] The benefit associated with stimulating research through competitive advocacy comes at a price, of course. Biased advocates will sometimes suppress useful information, and in some cases, the advocacy system will stimulate too much (possibly redundant) research. It is also possible that the incentives an advocacy system creates will stimulate undesirable behaviors — including behaviors that worsen information gathering, such as fabrication or distortion of evidence. But in many cases, these costs may be outweighed by the greater research incentives that the competitive advocacy scheme induces.[150]

The above analysis assumes that the pieces of information gathered by competing advocates are (partial) substitutes, or at least that the complementarities among pieces of evidence are not very strong. When there are strong complementarities among different pieces of information, however, the case for assigning the task of researching these different types of information to competing advocates, rather than to a

---

[147] Mathias Dewatripont & Jean Tirole, *Advocates*, 107 J. POL. ECON. 1, 13–15 (1999).

[148] *See id.* at 16–17, 25, 34; *see also* David Austen-Smith & John R. Wright, *Competitive Lobbying for a Legislator's Vote*, 9 SOC. CHOICE & WELFARE 229, 245 (1992) (arguing that legislators are more likely to vote correctly if they are lobbied by advocates on only one side of an issue than if they are not lobbied at all); Hongbin Cai, *Costly Participation and Heterogeneous Preferences in Informational Committees*, 40 RAND J. ECON. 173, 174 (2009) (arguing that, "under certain conditions, heterogeneous preferences [among biased agents] provide [agents] incentives to gather information"); Krishna & Morgan, *supra* note 86, at 767 (2001) (explaining that decisionmakers can derive full information even from experts with "opposing biases" under circumstances that allow for full debate and rebuttal). *But cf.* Morten Bennedsen & Sven E. Feldmann, *Informational Lobbying and Political Contributions*, 90 J. PUB. ECON. 631, 632 (2006) (arguing that certain "information externalit[ies]" reduce the incentive for biased agents to engage in robust research).

[149] *See* Dewatripont & Tirole, *supra* note 147, at 11–12.

[150] *Id.* at 33–34.

single neutral party, may weaken or collapse. To see this possibility, suppose that a legislature is considering two alternatives to the status quo, A and B, and that there are two pieces of information, X and Y, that might bear on this decision. The legislature would prefer to select one of the alternatives to the status quo only if both pieces of evidence support that decision; one piece of evidence alone, or contradictory pieces of evidence, would be insufficient to convince the legislature to adopt a new policy. If a single neutral agent is assigned the task of researching both X and Y, she may have an incentive to do so, as such research might lead to a desirable change from the status quo (to either A or B), even though the agent's incentives are dampened by the fact that if the pieces of evidence conflict (say, if X supports A but Y supports B), she will have wasted her research effort. But if the legislature assigns the tasks of researching X and Y to competing advocates, one of whom supports A and one of whom supports B, then neither advocate will invest in research. Each advocate knows that she alone cannot persuade the principal to implement any policy other than the status quo; any information she provides is useless unless the other advocate provides complementary information. But each advocate also knows that her adversary has no incentive to do so. Thus, dividing the task of producing complementary forms of information between competing advocates may have perverse effects.

A tournament system, like an advocacy system, seeks to exploit competitive pressures among agents to improve research incentives, but in a somewhat different way. While in an advocacy system each agent is matched with a preexisting policy option, in a tournament system each agent may design her own policy response to some problem and present her solution to the principal (or an assessor who acts on the principal's behalf). The agent whose proposal is selected then receives some reward.[151] If the likelihood that a proposal is selected is positively correlated with the quality of the evidence developed in support of that proposal, then this reward serves as an incentive for each agent to do extensive research. The advantage of a tournament system over an advocacy system is that even modest improvements in each agent's proposal may lead to a very large reward, which, at least under the right conditions, tends to spur greater effort.

The similarity of the tournament system to the advocacy system, at least in terms of the effect on agents' research incentives, is immediately apparent. Indeed, many of the costs and benefits are similar and need not be restated. Tournament systems, like advocacy systems, are

---

[151] *See* Edward P. Lazear & Sherwin Rosen, *Rank-Order Tournaments as Optimum Labor Contracts*, 89 J. POL. ECON. 841, 844–47 (1981); Barry J. Nalebuff & Joseph E. Stiglitz, *Prizes and Incentives: Towards a General Theory of Compensation and Competition*, 14 BELL J. ECON. 21, 26–30 (1983); Sherwin Rosen, *The Economics of Superstars*, 71 AM. ECON. REV. 845 (1981).

desirable insofar as they induce competition on socially beneficial dimensions — of greatest interest here, the production of useful information about policy alternatives. But tournaments may also entail serious costs, especially if they induce competition along undesirable dimensions (for example, by inducing greater distortion of information or attempts to sabotage competitors' efforts) or if they divide up responsibility and expertise for investigating complementary aspects of a problem.[152] These systems are also costly insofar as they involve socially wasteful duplication of effort.

CONCLUSION

One of the great and lasting accomplishments of the Legal Process school was bringing issues of institutional design to the forefront of the study of public law. Professors William Eskridge and Philip Frickey succinctly summarize this key tenet of the Legal Process school:

> In a government seeking to advance the public interest, each organ has a special competence or expertise, and the key to good government is not just figuring out what is the best policy, but figuring out which institutions should be making which decisions and how all the institutions should interrelate.[153]

More recent public law scholarship — a "New Legal Process" school[154] — has built on this insight, combining traditional legal analysis with theories and methods drawn from other disciplines to assess how best to allocate decisionmaking authority in light of the relative competence of different public institutions.[155]

However, the Legal Process insight that institutional design choices must take into account the relative competence of different government agents is incomplete, and potentially misleading, because it neglects the extent to which institutional choices may *change* the relative competence of different government agents. Institutional choices that appear prudent when government agents' expertise is treated as exogenous may turn out to be counterproductive when such agents' expertise is endogenous. Likewise, seemingly inefficient institutional choices may turn out to be desirable when the impact on information-gathering incentives is considered. Attention to the endogeneity of in-

---

[152] One possible difference between tournament systems and advocacy systems — and an advantage of the former over the latter — is that outcomes need not be specified in advance. This feature may also facilitate the participation of a larger number of potential competitors, though it is possible to have a "tournament" of two competitors, or an adversarial hearing with three or more competing advocates.

[153] Eskridge & Frickey, *supra* note 2, at lx.

[154] *See generally* Rubin, *supra* note 2; Daniel B. Rodriguez, *The Substance of the New Legal Process*, 77 CALIF. L. REV. 919 (1989) (reviewing WILLIAM N. ESKRIDGE, JR. & PHILIP P. FRICKEY, CASES AND MATERIALS ON LEGISLATION (1988)).

[155] *See* sources cited *supra* notes 2–4.

formation highlights the fact that in many contexts, feasibility constraints force an institutional designer to try to achieve two separate goals (such as efficient ex post use of information and efficient ex ante acquisition of information) with only one tool, in a way that creates an unavoidable trade-off.

Assessing the impact of any particular institutional design choice on the government's acquisition of decision-relevant information would require an in-depth contextual analysis well beyond the scope of this Article. Yet this Article has tried to show that there are a few useful basic principles, as well as some midlevel generalizations about the effects of particular classes of legal and institutional rules, that might prove helpful in guiding such a context-specific inquiry. At the most general level, an institutional change is likely to increase a government agent's research incentives if it either lowers the agent's marginal research cost or increases her marginal research benefit. The latter consideration — the focus of the analysis in this Article — can be further decomposed into two factors: the agent's default payoff (her expected utility from declining to invest additional resources in information gathering) and her research payoff (her expected utility if she does make such an investment). Lowering the former, or raising the latter, strengthens an agent's research incentives.

This seemingly simple observation turns out to have a great deal of utility: As this Article has shown, a great variety of legal and institutional choices can be assessed with respect to their likely impact on an agent's default payoff and research payoff. One can lower an agent's default payoff by, for example, prohibiting the agent from taking her ex ante preferred action or imposing more stringent procedural or evidentiary requirements on that action. One can also worsen an agent's default payoff by removing safety nets — such as the input of other agents or a second look by an oversight body — that increase the odds that the agent will realize an acceptable outcome even if she does very little work. One can raise an agent's research payoff by expanding her discretion and autonomy or by raising the relative attractiveness of decisions that seem undesirable ex ante. Under some circumstances, one can raise the research payoff more directly, by offering a reward — perhaps in the form of more discretion — if the agent demonstrates substantial research investment. One might further raise the research payoff by dividing research tasks among separate agents or by stimulating competition among agents with respect to overlapping research tasks. These and other observations suggest the usefulness of more fully integrating considerations of endogenous government expertise into New Legal Process research on optimal institutional design in public law.