
The Handbook of Organizational Economics

Edited by

ROBERT GIBBONS and JOHN ROBERTS

PRINCETON UNIVERSITY PRESS
Princeton and Oxford

Copyright © 2013 by Princeton University Press

Published by Princeton University Press, 41 William Street, Princeton, New Jersey 08540

In the United Kingdom: Princeton University Press, 6 Oxford Street, Woodstock, Oxfordshire OX20 1TW

press.princeton.edu

Jacket art: Stuart Davis (1894–1964), *Yellow Truck*, 1939. Gouache on paper, 15 $\frac{7}{8}$ × 11 $\frac{1}{8}$ in. Gift of Mrs. James Laughlin. 1981.61. Located at the Wadsworth Atheneum Museum of Art, Hartford, CT, USA. Photo credit: Wadsworth Atheneum Museum of Art / Art Resource, NY. Art © Estate of Stuart Davis / Licensed by VAGA, New York, NY.

All Rights Reserved

Library of Congress Cataloging-in-Publication Data

The handbook of organizational economics / edited by Robert Gibbons and John Roberts.

p. cm.

Includes bibliographical references and indexes.

ISBN 978-0-691-13279-2 (hbk. : alk. paper)

1. Industrial organization (Economic theory) 2. Organizational sociology—Economic aspects.

3. Organizational behavior—Economic aspects. I. Gibbons, Robert, 1958— II. Roberts, John, 1945 Feb. 11—

HD2326.H285 2013

338.5—dc23

2012033072

British Library Cataloging-in-Publication Data is available

This book has been composed in Adobe Garamond Premier Pro with Officina Sans display using ZzT_EX
by Princeton Editorial Associates Inc., Scottsdale, Arizona.

Printed on acid-free paper. ⊗

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

CONTENTS

Contributors ix

Introduction 1

Robert Gibbons and John Roberts

PART I FOUNDATIONS

1 Complementarity in Organizations 11

Erik Brynjolfsson and Paul Milgrom

2 Economic Theories of Incentives in Organizations 56

Robert Gibbons and John Roberts

3 Property Rights 100

Ilya Segal and Michael D. Whinston

4 Transaction Cost Economics 159

Steven Tadelis and Oliver E. Williamson

PART II METHODS

5 Clinical Papers in Organizational Economics 193

George P. Baker and Ricard Gil

6 Experimental Organizational Economics 213

Colin F. Camerer and Roberto A. Weber

7 Insider Econometrics 263

Casey Ichniowski and Kathryn Shaw

PART III WITHIN FIRMS: INDIVIDUALS AND GROUPS

8 Employment as an Economic and a Social Relationship 315

James N. Baron and David M. Kreps

- 9 Authority in Organizations 342
Patrick Bolton and Mathias Dewatripont
- 10 Decisions in Organizations 373
Robert Gibbons, Niko Matouschek, and John Roberts
- 11 Leadership and Corporate Culture 432
Benjamin E. Hermalin
- 12 Personnel Economics 479
Edward P. Lazear and Paul Oyer
- 13 Theory and Evidence in Internal Labor Markets 520
Michael Waldman

PART IV WITHIN FIRMS: STRUCTURES AND PROCESSES

- 14 Technological Innovation and Organizations 575
Pierre Azoulay and Josh Lerner
- 15 Hierarchies and the Division of Labor 604
Luis Garicano and Timothy Van Zandt
- 16 Internal Capital Markets 655
Robert Gertner and David Scharfstein
- 17 What Do Managers Do? 680
Robert Gibbons and Rebecca Henderson
- 18 Corporate Governance 732
Benjamin E. Hermalin
- 19 Incentives in Hierarchies 764
Dilip Mookherjee
- 20 Strategy and Organization 799
John Roberts and Garth Saloner

PART V BETWEEN FIRMS

- 21 Vertical Integration and Market Structure 853
Timothy Bresnahan and Jonathan Levin
- 22 Ownership and Organizational Form 891
Henry Hansmann
- 23 Contracts between Legal Persons 918
Lewis A. Kornhauser and W. Bentley MacLeod
- 24 Inter-Firm Contracts 958
Francine Lafontaine and Margaret E. Slade
- 25 Relational Incentive Contracts 1014
James M. Malcomson
- 26 Hybrid Modes of Organization 1066
Claude Ménard

PART VI BEYOND FIRMS

- 27 Corruption 1109
Abhijit Banerjee, Rema Hanna, and Sendhil Mullainathan
- 28 Delegation, Control, and the Study of Public Bureaucracy 1148
Terry M. Moe
- Name Index* 1183
- Subject Index* 1203

- MacLeod, B., and J. Malcolmson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57: 447-480.
- Martimort, D. 1996. "Exclusive Dealing, Common Agency and Multiprincipals Incentive Theory." *RAND Journal of Economics* 27: 1-31.
- Radner, R. 1993. "The Organization of Decentralized Information Processing." *Econometrica* 61: 1109-1146.
- Roberts, J. 2004. *The Modern Firm: Organizational Design for Performance and Growth*. Oxford: Oxford University Press.
- Rotemberg, J., and G. Saloner. 1993. "Leadership Styles and Incentives." *Management Science* 39: 1299-1318.
- Simon, H. 1951. "A Formal Theory of the Employment Relationship." *Econometrica* 19: 293-305.
- Stern, S. 2004. "Do Scientists Pay to Be Scientists?" *Management Science* 50: 835-853.
- Townsend, R. 1979. "Optimal Contracts and Competitive Markets with Costly State Verification." *Journal of Economic Theory* 21: 265-293.
- Van den Steen, E. 2007. "The Cost of Incentives under Disagreement (Can an Employee Be Too Motivated?)." Working paper, Massachusetts Institute of Technology, Cambridge, MA.
- . 2010. "Interpersonal Authority in a Theory of the Firm." *American Economic Review* 100: 466-490.
- Waldman, M. 1990. "Up-or-Out Contracts: A Signaling Perspective." *Journal of Labor Economics* 8: 230-250.
- Wernerfelt, B. 1997. "On the Nature and Scope of the Firm: An Adjustment-Cost Theory." *Journal of Business* 70: 489-514.
- Williamson, O. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.
- . 1979. "Transaction-Cost Economics: The Governance of Contractual Relations." *Journal of Law and Economics* 22: 3-61.
- . 1985. *The Economic Institutions of Capitalism*. New York: Free Press.
- Wilson, J. Q. 1989. *Bureaucracy: What Government Agencies Do And Why They Do It*. New York: Basic Books.

Decisions in Organizations

Robert Gibbons, Niko Matouschek, and John Roberts

1. Introduction

Organizations exist largely to get things done. Determining what should be done, by whom, how, when, and where (and then actually getting it done) requires decisions, and making good decisions depends on the decisionmaker's having the relevant information. However, as emphasized by Hayek (1945), the relevant information is typically dispersed (and may need to be developed). Thus, communication is necessary if decisionmakers are to become informed. If those communicating the information recognize how what they transmit will affect decisions, then they may be led to communicate strategically, withholding information or misrepresenting it. Then the decisionmaker will need to take account of this behavior when making decisions. This chapter investigates these issues and how the negative consequences of strategic behavior can be limited and the quality of decisionmaking improved through design of the decision process and the reward structure of the organization.

Jacob Marschak initiated the formal study of the value of information in decision problems under uncertainty during the 1950s, culminating in Marschak (1971). In parallel, Hurwicz (1960) began exploring the problem of designing communication and decision systems to achieve allocative and informational efficiency. Especially notable in this vein of research is the Marschak and Radner (1972) theory of teams, which analyzed systematically the problems of information acquisition, communication, and decisionmaking in organizations. Notably, this work and the large literature that follows from it assume that all parties share a common objective, such as maximizing the firm's profit. There is thus no shirking, free riding, lying, lobbying, collusion, or politics in these models.¹

We are very grateful for help from Ricardo Alonso, Dan Barron, Nancy Beaulieu, Leshui He, Anton Kolotilin, Hongyi Li, Meg Meyer, Mike Powell, Michael Raich, Heikki Rantakari, Eric Van den Steen, Tommy Wang, Julie Wulf, Luis Zernello, and seminar participants at Oxford. We also gratefully acknowledge research support from the Massachusetts Institute of Technology Sloan School's Program on Innovation in Markets and Organizations.

1. See Garicano and Van Zandt (this volume) for a thorough treatment of common-interests models of organizations.

The formal analysis of such incentive issues in decision and allocation mechanisms began with Hurwicz (1972, 1973), who introduced explicit game-theoretic treatments of strategic communication and first investigated the possibility of designing mechanisms to minimize the effects of strategic behavior induced by the mechanism. Essentially all the literature we discuss in this chapter builds on the foundations that Hurwicz laid.²

The other key element in developing the literature we examine here is contractual incompleteness.³ The parties not only cannot bind themselves to be fully honest in their reports, they also cannot necessarily commit to particular decision rules or future reward schemes. This moves the analysis from the domain of “mechanism design” that has flourished in the wake of Hurwicz’s initial contributions to one where shirking, free riding, lying, lobbying, collusion, or politics are central.

Informal arguments that conflicting interests in organizations are widespread and important for the design and functioning of the organization go back a long way in economics. For example, Knight (1921/1964: 254) observed that the “internal problems of the corporation, the protection of its various types of members and adherents against each other’s predatory propensities, are quite as vital as the external problem of safeguarding the public interests against exploitation by the corporation as a unit.”

Consistent with the importance of these issues, outside economics, the political approach to organizational behavior also has a long pedigree. For example, March (1962/1988: 110–112) described “the business firm as a political coalition” in which “the executive . . . is a political broker” who cannot “solve the problem of conflict by simple payments to the participants and agreement on a superordinate goal.” In a similar spirit, Zald and Berger (1978) described organizational conflict and change in terms of social movements, finding inside corporations analogues of coups d’etat, bureaucratic insurgencies, and mass movements.

As Gibbons (2010: 347) suggests, some of this early work outside economics is close in spirit to current economic modeling. For example, Cyert and March (1963: 30, 32) constructed strategic accounts of organizational behaviors in terms of individuals’ decisionmaking, beginning from the assertions that “people (i.e., individuals) have goals; collectivities of people do not” and “the existence of unresolved conflict is a conspicuous feature of organizations.” More specifically, Cyert and March (1963: 79) anticipated the application of information economics to the study of organizations, arguing first that “where different parts of the organization have responsibility for different pieces of information relevant to a decision, we would expect . . . attempts to manipulate information as a device for manipulating the decision.” They also saw that such strategic behavior would generate attempts by the decisionmaker to undo the attempted manipulation: “We cannot reasonably introduce the concept of communication bias without introducing its obvious corollary—‘interpretive adjustment’” (Cyert and March 1963: 85). In short, compared to the traditional description of an organization in terms of an organization

chart, Cyert and March (1963: 202) suggested that “the kinds of models presented in this book describe the organization as a decision-making process”—and a strategic one, at that.

This chapter surveys the theoretical literature in organizational economics on decisions in organizations, focusing on models where parties’ interests are imperfectly aligned. This literature is mostly recent but is growing quickly. Our overarching perspective is that this literature is very much in the spirit of the approach initiated by March (1962/1988) and Cyert and March (1963).

More specifically, we break our discussion into two parts. In the first half of the chapter (Sections 2 and 3), we study game-theoretic models of communication and decisionmaking under fixed decision processes (e.g., fixed specifications of who controls which decisions). These models categorize various ways in which strategic considerations with respect to information transmission can impose costs on organizations. In the second half (Sections 4 and 5), we then endogenize the decision process (e.g., we ask what allocation of decision rights the parties should choose). More generally, in these sections we explore the tools that organizations have to mitigate the costs identified in Sections 2 and 3.

2. Fixed Decision Processes

This section is the first of two on decisionmaking under fixed decision processes. Here we describe and analyze a basic model that is on one hand quite simple but on the other surprisingly flexible, in that it encompasses a wide range of models from the literature. We use this basic model mostly to introduce ideas. Section 3 then discusses richer analyses.

One of the motivations for our basic model comes from Milgrom and Roberts (1988), who discussed the incentives for people in organizations to distort information collection and transmission to influence decisions to their advantage. They recognized a variety of possible costs of such “influence activity,” including the opportunity costs of time and other resources spent on trying to affect the distribution of costs and benefits without improving efficiency, the possibility of decisions being distorted by the influence, and the costs of altering the organizational design and decision processes to limit influence.

Another motivation comes from Mintzberg (1979), who described a decision process as moving from information to advice to choice to execution, as shown in Figure 1.⁴ We begin by defining our basic model based upon what we see as the influence link in Mintzberg’s schema, from advice to choice. While a surprising array of models can be captured this way, in Section 3 we discuss both a reinterpretation of our basic model as being from choice to execution and an enrichment of the model to being from information to advice to choice.

After describing the basic model in Section 2.1, in Sections 2.2–2.6 we analyze five illustrations of the basic model, motivated by the cheap-talk model in Crawford and Sobel (1982), the signal-jamming models in Holmström (1982/1999) and Gibbons (2005), the signaling models in Austen-Smith and Banks (2000) and Kartik (2007), the verifiable-information models in Grossman (1981) and Milgrom (1981), and the strategic management of public information

4. Fama and Jensen (1983) proposed a related process running from initiation to ratification to implementation to monitoring. This process does not so directly involve communication, so we favor Mintzberg’s for present purposes.

2. See Maskin (1977/1999) and Myerson (1981) for seminal contributions to mechanism design; see Mookherjee (this volume) for applications of these ideas to organizational issues.

3. Williamson (1975) was a leader in insisting on the importance of contractual incompleteness in organizational problems, although, for example, Simon (1951) also depends on incomplete contracting. Grossman and Hart (1986) introduced formal models of incomplete contracts.




Figure 1. The decision process, adapted from Mintzberg (1979).

models in Rayo and Segal (2010) and Kamenica and Gentzkow (2011).⁵ These five illustrations show how even this basic model, although simple, also can generate quite rich examples of strategic political behavior in communication and decisionmaking. We illustrate how different environments for communication and decisionmaking produce different strategic behaviors that impose different costs on the organization.

In Sections 2.2–2.6 we present just the core results that arise from these five illustrations of the basic model. In Section 2.7 we compare, contrast, and interpret these five models and relate their results to prominent claims and observations in the noneconomic literature.

2.1. A Basic Model

Our basic model has only two players, each of whom moves once in the following simple sequence:

1. The state of the world $s \in S$ is drawn from the distribution $f(s)$.
2. Player 1 privately observes the signal θ drawn from the distribution $g(\theta|s)$.
3. Player 1 chooses an *influence action* $a \in A$.
4. Player 2 privately observes the signal σ drawn from the distribution $h(\sigma|s, a)$.
5. Player 2 chooses a *decision* $d \in D$.
6. The players receive payoffs $U_i(s, a, d)$ for $i = 1, 2$.

Note that we assume that payoffs $U_i(s, a, d)$ are independent of the signals θ and σ , capturing the idea that it is really the state $s \in S$ and the players' choices that matter.

For simplicity (and, to some extent, for realism regarding the environments we attempt to capture), we assume that nothing in this model is contractible. For example, player 1's influence action $a \in A$ is not contractible, nor is player 2's decision $d \in D$, nor are the players' payoffs $U_i(s, a, d)$. Thus, there cannot be bargaining over either player's move, nor can there be mechanisms designed to elicit either player's private signal, nor can there be outcome-contingent incentive contracting.

5. Although the terms "signaling," "cheap talk," and perhaps "signal jamming" are well established, the models in Sections 2.5 and 2.6 go by various names, some of which apply to most or all models in this section, such as "persuasion" or "information disclosure." We have therefore chosen names for the models in these two sections that illustrate their core assumption or interpretation: "verifiable information" and "strategic management of public information," respectively.

For definiteness, we take player 1 to be male and player 2 female. As appropriate below, we also refer to the players as "sender" and "receiver," "employee" and "owner," or "leader" and "follower," respectively.⁶

2.2. Cheap Talk

Perhaps the most direct application of our basic model is the seminal theory of "cheap talk" by Crawford and Sobel (1982), in which messages are costless (so talk is not just cheap but actually free). In particular, player 1 learns the state of the world (i.e., $\theta = s$), player 1's action is to send a message to player 2 (i.e., $\sigma = a$), and this message is payoff irrelevant (i.e., payoffs are $U_i(s, d)$). The major result in such models is that the misalignment of preferences typically prevents full revelation of information.

As an informal example of a cheap-talk sender-receiver game, consider a biased local manager advising a CEO about capital allocation to the manager. The manager can say anything he wants to the CEO, at no cost. The CEO wants to gear her decision to the local information that only the manager has, but the manager is known to want the decision to differ systematically from what the CEO would want in any particular state, for instance because of empire building or career concerns. Although the manager is biased, suppose he also has an interest in the CEO's decision being geared to the environment. Then the manager's interest in the CEO's decision being geared to the environment may limit, but not eliminate, his incentive to misrepresent his information. Put differently, in a cheap-talk setting, the CEO cannot accept at face value any claim the manager makes (lest the manager make an inflated claim that would then cause the CEO to take the decision the manager wants), nor can the manager tell the CEO the full truth (lest the CEO then ignore the manager's preferences and impose the decision she prefers). In equilibrium, however, if the manager is not too biased, then partial communication can occur, as follows.

Suppose that the state s is uniformly distributed on $[0, 1]$ and the parties have quadratic payoffs $U_2(s, a, d) = -(d - s)^2$ and $U_1(s, a, d) = -(d - [s + b])^2$, where the imperfect alignment of the parties' interests is captured by the parameter b , which we take to be positive for definiteness. In this setting, the sender's message rule $a(s)$ cannot perfectly reveal the state. Loosely, if the sender's message did reveal the state, then the receiver would choose the decision $d = s$, so the sender would want to send a message that induced the receiver to choose $d = s + b$.

Instead of full revelation of the sender's information, what does happen is that all perfect Bayesian equilibria are economically equivalent to partition equilibria: the state space is partitioned into N intervals $0 = s_0 \leq s_1 \leq \dots \leq s_{N-1} \leq s_N = 1$ and the sender's message reveals which interval $[s_{i-1}, s_i]$ contains the state. For any value of b , there is an uninformative equilibrium in which $N = 1$ (i.e., the sender's message is equivalent to saying that $s \in [0, 1]$, which reveals nothing about the state, so the receiver retains her prior belief that s is uniform on $[0, 1]$).

6. Although all examples that follow exhibit some degree of conflict of interest between the players, this framework can also be used to explore team-theoretic models, which assume common interests. For example, see Section 4.3 of Bolton and Dewatripont (this volume) for a simplified version of Dessein and Santos (2006) that fits our basic framework.

and hence chooses $d = \frac{1}{2}$). For sufficiently small values of b , however, there are also other equilibria with more intervals. For example, if $b < \frac{1}{4}$, then there exists a two-interval equilibrium in which $s_1 = (1 - 4b)/2$. In this equilibrium the sender's message reveals whether $s \in [0, s_1)$ or $s \in [s_1, 1]$, and the receiver's decision is then $\hat{d} = s_1/2$ or $\hat{d} = (1 + s_1)/2$, accordingly. Note that the ideal decision for sender type $s = s_1$ is $\hat{d} = s_1 + b$, so sender type $s = s_1$ is indifferent between the receiver's two decisions in this equilibrium, $\hat{d} = s_1/2$ or $\hat{d} = (1 + s_1)/2$, and this indifference is what makes all sender types $s < s_1$ prefer to say $s \in [0, s_1)$ and all those $s > s_1$ prefer to say $s \in [s_1, 1]$.

For sufficiently small values of b , there are more equilibria—one for each number of intervals from $N = 1$ to $N = N(b)$, where $N(b)$ weakly increases as b falls toward zero. In each such equilibrium, the sender type on the boundary between two intervals, $s = s_i$ between $s \in [s_{i-1}, s_i)$ and $s \in [s_i, s_{i+1})$, is again indifferent between the receiver's equilibrium decisions following these messages, and this indifference is again what makes other sender types strictly prefer to send their equilibrium messages. When b falls, the necessary distance between these two decisions decreases, so it becomes possible to have more interval messages in equilibrium (i.e., $N(b)$ increases).

Interestingly, although Crawford and Sobel's original motivation was from bargaining with asymmetric information, their model has recently become a workhorse in the organizations literature. See Sections 3.1, 3.2, 3.3, and parts of Section 5 for the next steps and Sobel (2012) for a comprehensive survey.

2.3. Signal Jamming

One of the central ideas in the Milgrom and Roberts (1988) work on influence activities is the observation that those in control get lobbied.⁷ This idea also fits naturally in our basic model, such as when an owner controls a decision and would like to have the decision tailored to the state of the world but does not know the state, and an employee has an opportunity to influence the signal about the state that the owner observes (and hence to influence the owner's belief about the state and thus influence the owner's decision). The major result of such models is that resources are spent on lobbying even though, in equilibrium, lobbying does not affect the decision.

More formally, consider the following model from Gibbons (2005), which adapts Holmström's (1982/1999) model of career concerns in labor markets to consider influence activities in organizations. A key point is that these models analyze *signal jamming*. That is, in a departure from both the cheap-talk model in Section 2.2 and the signaling model in Section 2.4, here there is symmetric uncertainty about the state s (i.e., player 1's signal θ in stage 2 of the basic model is uninformative). Nonetheless, the employee (player 1) does have an opportunity to influence the signal that the owner (player 2) observes about the state (σ in stage 4 of the basic model). In particular, $\sigma = s + a + \varepsilon$, where $a \geq 0$ at cost $k(a)$ is the employee's choice of "lobbying" activities in stage 3 of the basic model. We assume that s is normally distributed with mean m and variance $1/h$ (i.e., precision h) and ε is normally distributed with mean zero and precision h_ε .

7. Parts of this section draw on Gibbons (2010), with permission from the *Annual Review of Economics*, Volume 2, © 2010 by Annual Reviews, <http://www.annualreviews.org>.

Because the state s is normal, assume that the owner's decision space is $D = \mathcal{R}$. Finally, let the cost function satisfy $k'(0) = 0$, $k'(\infty) = \infty$, and $k'' > 0$, and suppose that the parties again have quadratic loss functions. The parties' payoffs are then $U_1(s, a, d) = -(d - [s + b])^2 - k(a)$ and $U_2(s, a, d) = -(d - s)^2$, where again $b > 0$ for definiteness.

The signal σ is the crux of the model: even though neither party knows the true state, both parties care about how the eventual decision will relate to the state, so the owner will try to extract from the signal whatever information σ might contain about the state, prompting the employee to try to move the realization of σ upward. In equilibrium, however, the owner correctly anticipates the employee's attempts to influence σ and so correctly accounts for those attempts when interpreting σ as a signal about s . Nevertheless, the employee has an incentive to increase σ . In particular, even though the owner is not fooled, the equilibrium level of lobbying cannot be zero, because if the owner believed the employee to be doing no lobbying, then the employee would have a strong incentive to lobby.

To analyze the model, we work backward. The owner will choose d to solve

$$\max_{d \in D} E_s [U_2(s, d) | \sigma],$$

so the solution is $E_s[s | \sigma]$. If the owner's conjecture about the employee's lobbying is a , then DeGroot (1970: 155–189) shows that

$$E_s[s | \sigma] = \frac{hm + h_\varepsilon(\sigma - \hat{a})}{h + h_\varepsilon}.$$

The employee therefore chooses a to solve

$$\max_{a \geq 0} -E_{s,\varepsilon} \left[\left(\frac{hm + h_\varepsilon(\sigma - \hat{a})}{h + h_\varepsilon} - (s + b) \right)^2 \right].$$

The first-order condition for a is then

$$-2 \left(\frac{h_\varepsilon}{h + h_\varepsilon} \right) \left\{ \frac{h_\varepsilon}{h + h_\varepsilon} (a - \hat{a}) - b \right\} = k'(a),$$

which implicitly defines $a^*(\hat{a})$, the employee's best response to the owner's conjecture.

In equilibrium, the owner's conjecture must be correct, so imposing $a^*(\hat{a}) = \hat{a}$ yields the first-order condition for the equilibrium level of the employee's lobbying activity. Denoting this equilibrium level of lobbying by a^* , we have

$$2 \frac{h_\varepsilon}{h + h_\varepsilon} b = k'(a^*).$$

In words, the equilibrium level of lobbying is positive, even though in equilibrium the owner is not fooled (so the lobbying costs are a deadweight loss). Furthermore, a^* increases with the employee's bias b (because the employee cares more about moving the owner's decision), with

the precision of the public signal h_s (because the owner will put more weight on the signal, so it is more useful to the employee to influence the signal), and with the prior variance of the state $1/h$ (because the owner will again put more weight on the signal, now for lack of useful prior information about the state).

2.4. Signaling

To complement the foregoing cheap-talk and signal-jamming models illustrating the basic model, we turn next to a third category of games, again familiar from information economics: signaling. (See Fudenberg and Tirole (1991: 324–336) for the basic game theory.) As in a cheap-talk game, player 1 has private information (i.e., $\theta = s$ in stage 2 of the basic model), and player 2 observes player 1's action (i.e., $\sigma = a$ in stage 4), so the sender's choice of a may change the receiver's belief about s and hence change the receiver's choice of d in stage 5. As in the signal-jamming game, but unlike the cheap-talk game, player 1's action (a in stage 3) incurs cost. The major result in the models that follow Spence (1973) is that some actions may convey information because only certain types would be willing to take those actions (and they are willing only because these actions will indeed convey information). This may happen either because the costs of the action to the sender or the sender's benefit from the receiver's response are related to the state in such a way that, on one hand, it is worthwhile expending the costs to have the decisionmaker infer the state correctly for some realizations of the state, and on the other, it is not worthwhile for senders in other states to try to fool the decisionmaker.⁸

Hermalin (1998) uses a signaling model to study leadership. He begins from the view that "a defining feature of leadership . . . [is that a] leader is someone with followers" (Hermalin: 1188), from which it follows that "a central question in understanding leadership is how does a leader induce others to follow?" In his model of "leading by example," the leader (here, player 1) observes the state of the world and takes an observable action; the follower (player 2) then draws an inference about the state and subsequently makes a decision. In "good" states, where having the followers actually follow is more important, the leader will take a more costly action to signal that his private information really does indicate that following is particularly valuable. Hermalin (this volume) discusses several models in this vein.

When there are benefits to revealing player 1's information about the state, signaling can occur simply by expending resources in an otherwise useless fashion. Such wasting of resources has come to be called "money burning," but in an organizational context it might be fighting through red tape and bureaucratic sclerosis: being willing to spend one's time on such activities signals one's conviction. For expositional simplicity, we describe a signaling model of money burning adapted from Austen-Smith and Banks (2000) and Kartik (2007); see Kartik et al. (2007) and Kartik (2009) for similar models with lying costs instead of money burning. As in the cheap-talk model, suppose the state is uniformly distributed on $[0, 1]$. As in the signal-jamming model, the parties' payoffs are again $U_1(s, a, d) = -(d - [s + b])^2 - k(a)$ and $U_2(s, a, d) = -(d - s)^2$, with $b > 0$ for definiteness. To capture the idea that player 1's action is burning money, however, assume that $k(a) = a$ for $a \geq 0$. Unlike the cheap-talk model, where

the sender's action could not perfectly reveal the state, here a separating equilibrium can exist, as follows.

If the sender's action perfectly reveals the state, then the receiver's optimal decision is equal to the state that the sender's action revealed, so the sender of type s could choose to induce any decision $d' = s'$ by taking the action intended to be taken by sender type s' . To deter such deviations, the sender's action rule $a^*(s)$ must satisfy

$$-b^2 - a^*(s) \geq -(s' - [s + b])^2 - a^*(s')$$

for all s' . It is straightforward to check that $a^*(s) = 2bs$ satisfies the above incentive-compatibility condition for the sender and thus can be the basis for a separating perfect Bayesian equilibrium in this signaling game. Compared to the signal-jamming game (which has symmetric uncertainty about s and, as a result, the cost of influence activities is $k(a^*)$, regardless of the state), here the cost of money burning increases with the state, because, in a separating equilibrium, each sender of type s needs to distinguish himself from those just below.

2.5. Verifiable Information

Grossman (1981) and Milgrom (1981) introduced the possibility that the sender can withhold information but cannot otherwise misrepresent his information. As opposed to the cheap-talk game (where the sender can say anything) and the signal-jamming and signaling games (where the sender's action is not a direct report about the state of the world), here we call any information the sender presents "hard" or "verifiable." Some such models lead to efficient decisions and outcomes, because the incentives for the sender not to be thought to have very negative information lead to full revelation of the state. Thus, a major result in such models is that strategic considerations may carry no cost in communication and decisionmaking.

Formally, as in the cheap-talk game, player 1 learns the state of the world (i.e., $\theta = s$), player 1's action is to send a message to player 2 (i.e., $\sigma = a$), and this message is payoff irrelevant (i.e., payoffs are $U_i(s, d)$). The question is then how to model the idea that player 1 can withhold but not misrepresent information. The simplest approach is to enrich the basic model slightly, by allowing the feasible set A to depend on the state and to assume that $A = \{s, \phi\}$, meaning that player 1's options are to reveal the state ($a = s$) or say nothing ($a = \phi$).⁹

After receiving the signal σ , player 2 then chooses a decision d , and the parties receive their payoffs. Player 1's payoff is increasing in d and independent of s , whereas player 2's is such that, if player 2 knew s , then her optimal choice of d would be strictly increasing and continuous in s . For now, assume $U_1(s, a, d) = d$ and $U_2(s, a, d) = -(d - s)^2$. As a result, player 1 wants to persuade player 2 that the state is high. For simplicity, suppose that s is distributed on $[0, 1]$.

The first (and striking) result in this model is that any equilibrium has full revelation. The key to this revelation result is an "unraveling" argument. Suppose first that the sender announces $a = \phi$ for all s . Then the receiver will optimize given her prior distribution on the state space, choosing $d = E(s)$. However, when $s = 1$ (and actually for any s close to 1), the sender would

8. This is essentially a matter of the sign of the cross partial derivative of the sender's payoff with respect to the receiver's inference about the state and the sender's action.

9. The same basic results hold if one assumes instead that A is all the subsets of the state space S and player 1 of type s can send any message as long as it contains s .

gain from deviating from announcing the actual state, because this would lead to a higher decision. In fact, if there are any two different states s' and s'' , with $s' > s''$, that both send the same message, then the receiver's decision will be higher if s' reveals itself rather than sending that message, so unraveling will occur. So the equilibrium must have correct revelation (except perhaps at $s = 0$, but this case is immaterial).

Milgrom and Roberts (1986) and Okuno-Fujiwara et al. (1990) extended the basic analysis in a number of directions, such as competition among senders and general Bayesian games (i.e., many players with private types, each taking their own actions). Seidmann and Winter (1997) provide general conditions (on the preferences of players 1 and 2) for existence and uniqueness of a fully revealing equilibrium; in particular, they show that the quadratic-loss, constant-bias payoffs $U_1(s, a, d) = -(d - [s + b])^2$ and $U_2(s, a, d) = -(d - s)^2$ also necessitate full revelation in equilibrium (for overviews, see Milgrom 2008; Sobel 2012).

Finally, although full revelation necessarily occurs in the equilibrium of the basic verifiable-information game above, this result depends on it being common knowledge that the sender is informed (see, for instance, Dye 1985). Suppose instead that in stage 2 there is probability q that $\theta = s$ and probability $1 - q$ that $\theta = \phi$ (i.e., the sender receives no information at all). Straightforward calculations then show that in equilibrium there is a critical value $s^*(q)$ such that

1. if $\theta = s$ and $s > s^*(q)$, then the sender reveals s ;
2. if $\theta = s$ and $s < s^*(q)$, then the sender reports $a = \phi$; and
3. if $\theta = \phi$, then the sender reports $a = \phi$.

To compute $s^*(q)$, suppose that s is uniformly distributed on $[0, 1]$, $U_1(s, a, d) = d$ and $U_2(s, a, d) = -(d - s)^2$. If the sender reveals s , then the receiver chooses $d = s$, whereas if the sender reports $a = \phi$, then the receiver chooses $d = q \times s^*(q)/2 + (1 - q) \times \frac{1}{2}$, so $s^*(q) = (1 - q)/(2 - q)$. Naturally, for $q = 1$, we recover full revelation, whereas for $q = 0$, only the upper half of the type space reveals itself.

2.6. Strategic Management of Public Information

The final illustration of our basic model concerns the strategic management of public information. A major contribution of this emerging literature is to shed light on the conditions under which the ability to manage public information strategically leads to a loss of information.

To introduce this final illustration of our basic model, consider the issue of agenda control. In some settings, for instance, information acquisition and deliberation are followed by voting, but one or more actors can call the question, meaning that the vote must occur immediately, which prevents further information acquisition. For a simple model motivated by Brocas and Carrillo (2007), suppose player 1 has no private information, but he can attempt to influence player 2's belief by affecting what public information she observes. Formally, the state is uniformly distributed on $[\underline{s}, \bar{s}]$, where $\underline{s} < 0 < \bar{s}$, player 1's signal θ is uninformative, player 1's influence action is $a \in A = \{0, 1\}$, and player 2's signal is $\sigma = sa$.

Now suppose that player 2's decision d is binary: either she implements a new project with payoff s for her or she preserves the status quo with payoff 0 for her. Player 1, however, benefits

when the new project is implemented and not otherwise, and player 1's action a is costless. Formally, $D = \{0, 1\}$, $U_1(s, a, d) = d$, and $U_2(s, a, d) = ds$.

In this setting, player 1 has two choices: $a = 0$, in which case player 2 chooses $d = 1$ if and only if $E(s) > 0$, or $a = 1$, in which case player 2 observes $\sigma = s$ and then chooses $d = 1$ if and only if $s > 0$. Clearly, if $E(s) > 0$, then player 1 should choose $a = 0$, and otherwise he should roll the dice by choosing $a = 1$. In a richer model, instead of being uninformative, player 1's signal θ could be publicly observed, in which case analogous reasoning begins by assessing whether $E(s|\theta) > 0$.

Beyond agenda control, now consider a pharmaceutical company designing a clinical trial to influence a regulator. Alternatively, to give an example from inside an organization, imagine a boss, who, as part of rolling out a new program, can commission a study in the hope that its public results will motivate employees to implement the program well. What kind of trial/study should player 1 design?

Formally, the model is as above, except now suppose that player 1's influence action is to choose a critical value $s^* \in [\underline{s}, \bar{s}]$ such that player 2's signal will be $\sigma = (s^*, x)$, where x is a binary random variable: $x = L$ if $s < s^*$, and $x = H$ if $s > s^*$. To keep things interesting, suppose that $E(s) < 0$ (i.e., $|s| > \bar{s}$), in which case player 1's optimal influence action is to choose $s^* = -\bar{s}$, maximizing the probability that player 2 will choose $d = 1$.

Kamenica and Gentzkow (2011) study a general version of this model, which can be interpreted as strategic management of public information by an interested party. This final illustration of our basic model is closer to the signal-jamming than to the cheap-talk, signaling, or verifiable-information illustrations, because we assume that player 1 has no private information and his influence activity is an attempt to change player 2's belief by changing the design of a public signal. Formally, player 1's influence action a is to choose a realization space X and a family of likelihood functions $\{\lambda(x|s)\}_{s \in S}$. Player 2's signal σ is then the family of likelihood functions and a realization $x \in X$. In the clinical trial/corporate study example above, the critical value s^* was all player 2 needed to know to understand the family of likelihood functions $\{\lambda(x|s)\}_{s \in S}$.

Kamenica and Gentzkow (2011) provide necessary and sufficient conditions for player 1's optimal influence action to be uninformative (e.g., X has only one element), and they derive other properties of player 1's optimal influence action, including the case of quadratic loss functions (which we discuss further in Section 2.7). Kolotilin (2012a) analyzes the special case above, where D is binary and player 1's payoff is $U_1(s, a, d) = d$, for which he fully characterizes player 1's optimal influence action. Finally, Rayo and Segal (2010) and Kolotilin (2012b) analyze the case where player 2 has private information, separate from the realization of player 1's likelihood functions. This case is important, both because a single player 2 may indeed have private information and because there may instead be a group of player 2s with known but heterogeneous preferences.

2.7. Summary

We now summarize and assess the models in this section, noting both their similarities and their differences, as well as some of their existing and potential connections to past and future literatures. To emphasize some important similarities, we begin with the special case of our basic

model in which the parties have quadratic loss functions. To encompass the signal-jamming and signaling models, we also include the cost $k(a)$ in player 1's payoff, so the payoff functions are $U_1(s, a, d) = -(d - [s + b])^2 - k(a)$ and $U_2(s, a, d) = -(d - s)^2$.

Under full information about s , the action $a = 0$ and the decision $d = s + (b/2)$ would maximize the parties' total payoff

$$U_1(s, a, d) + U_2(s, a, d) = -(d - [s + b])^2 - k(a) - (d - s)^2,$$

producing a total payoff of $-\frac{1}{2}b^2$. Because d is not contractible, however, in all five models in this section—cheap talk, signal jamming, signaling, verifiable information, and strategic management of public information—player 2's equilibrium decision is $d^*(\sigma) = E_{s|\sigma}(s|\sigma)$. We therefore can write the parties' equilibrium expected total payoff as

$$\begin{aligned} E[U_1(s, a, d) + U_2(s, a, d)] &= E_s[-(d - [s + b])^2 - k(a(s)) - (d - s)^2] \\ &= E_\sigma\{E_{s|\sigma}[-(d - [s + b])^2 - (d - s)^2]|\sigma\} - E_s[k(a(s))] \\ &= -b^2 - 2E_\sigma[Var(s|\sigma)] - E_s[k(a(s))]. \end{aligned}$$

Thus, regardless of the model, the loss the parties suffer between their optimized total payoff under full information, $-\frac{1}{2}b^2$, and their equilibrium expected total payoff in any of our models, $-b^2 - 2E_\sigma[Var(s|\sigma)] - E_s[k(a(s))]$, is

$$\frac{1}{2}b^2 + 2E_\sigma[Var(s|\sigma)] + E_s[k(a(s))].$$

This expression measures the total expected equilibrium loss from strategic communication and decisionmaking in all five models. The first term, $\frac{1}{2}b^2$, captures the loss in total payoff from the noncontractibility of d , so player 2's equilibrium decision is $d^*(\sigma) = E_{s|\sigma}(s|\sigma)$ rather than $E_{s|\sigma}(s|\sigma) + \frac{1}{2}b$. In this sense, the average decision is not optimal, regardless of the state. The parties incur this loss in all five models.

The second term, $2E_\sigma[Var(s|\sigma)]$, measures the loss $E_\sigma[Var(s|\sigma)]$ that each party incurs because the decision is not perfectly adapted to the state, as it would be if the decision were $d = s + (b/2)$ or any other function $d = s + z$ for some constant z . This loss is not about the average decision, but instead about how the decision varies with the state. Because different models produce different equilibrium signals σ , the size of this loss varies across models. For example, the coarse communication in the cheap-talk model does not reduce this term to zero, whereas the full revelation that may occur in the signaling, verifiable-information, and public-information models does, and the way the receiver accounts for the sender's influence activity in the signal-jamming model reduces this term not to zero but at least to as low as it can be, given the noise ε in the signal $\sigma = s + a + \varepsilon$.

Finally, the third term, $E_s[k(a(s))]$, reflects player 1's expected costs of influence activities, so it is zero in the cheap-talk, verifiable-information, and public-information models, but positive in the signal-jamming and signaling models. As we formulated both of the latter models, this cost is a pure waste, undertaken by the sender in equilibrium only to influence the receiver's

beliefs and hence her decision. In alternative formulations of both models, however, the sender's influence activity could also be productive (e.g., in Holmström's (1982/1999) signal-jamming model of wage determination or Hermalin's (1998) signaling model of leadership), but the inefficiency we capture in pure form here would not disappear.

We think the similarities across these models do a nice job formalizing the quotations from Cyert and March (1963) in Section 1 about "attempts to manipulate information as a device for manipulating the decision" and "interpretive adjustment" being an "obvious corollary" of communication bias. Almost two decades later, March (1981: 217) summarized subsequent work (outside economics) by asserting that

information is an instrument of consciously strategic actors. Information may be false; it is always serving a purpose. . . . Thus information itself is a game. Except insofar as the structure of the game dictates honesty as a necessary tactic, all information is self-serving.

Again, this section's models seem consistent with this summary. For example, the full-revelation result in the verifiable-information game could be an example where "the structure of the game dictates honesty as a necessary tactic." Ironically, the models and techniques we have surveyed began to appear in about 1981, whereas this was about the end of this literature outside economics.

Naturally, one of the reasons to build models is to explore the many different environments in which strategic communication and decisionmaking can play out and the many different behaviors and inefficiencies that can then result, so we see the differences among this section's models as useful in that regard. For example, Feldman and March (1981) describe various organizational activities not too distant from money burning in the signaling model: wasteful ways that organization members need to jump through hoops to establish credibility.

More generally, beyond money burning, after presenting several case studies and alluding to many others on information gathering, communication, and decisionmaking in organizations, Feldman and March (1981: 174) summarize these studies with six observations:

1. Much of the information that is gathered and communicated by individuals and organizations has little decision relevance.
2. Much of the information that is used to justify a decision is collected and interpreted after the decision has been made, or substantially made.
3. Much of the information gathered in response to requests for information is not considered in the making of decisions for which it was requested.
4. Regardless of the information available at the time a decision is first considered, more information is requested.
5. Complaints that an organization does not have enough information to make a decision occur while available information is ignored.
6. The relevance of the information provided in the decisionmaking process to the decision being made is less conspicuous than is the insistence on information.

Without taking these small-sample observations as gospel, they fit sufficiently well with our experiences that we suspect there would be steep returns to integrating the kinds of theoretical approaches sketched here with both qualitative and quantitative empirical work on information gathering, communication, and decisionmaking in organizations. To do so, we need to move beyond the relatively simple models described in this section to address more complex and realistic issues in organizational design and performance. In the next section, we begin to undertake this task.

3. Richer Analyses

Having introduced our basic model and its five illustrations, we turn next to several richer analyses. Specifically, we consider

Section 3.1: multidimensional cheap talk;

Section 3.2: more parties (senders, receivers, or both);

Section 3.3: more periods;

Section 3.4: a reinterpretation of the basic model as the choice-to-execution stages of Mintzberg's (1979) timing;

Section 3.5: an enrichment of the basic model to include the information-to-advice-to-choice stages;

Section 3.6: enrichments where some aspect of the model is contractible; and

Section 3.7: models of contests for control, where the right to make decision d cannot be allocated ex ante to player 2 but instead must be secured by one of the parties before d can be chosen in stage 5 of the basic model.

Although one of our main purposes in Section 2 was to specify models clearly enough to analyze and compare them, here we shift gears somewhat, focusing more on sketching applications than on formulating models or deriving results. One exception, however, is Section 3.1, because the recent models of multidimensional cheap talk importantly enrich the basic intuition from Section 2.2.

3.1. Multidimensional Cheap Talk

The cheap-talk model by Crawford and Sobel (1982) captures the fundamental insight that costless communication is less effective when the parties' interests are less aligned. Even our basic model in Section 2.1, however, placed no restrictions on the state space S or the decision space D , and hence allowed the possibility of multidimensional cheap talk. And, interestingly, recent work has shown that multidimensional cheap talk creates new ways for the parties' interests to be aligned and hence costless communication to be effective. Sobel (2012) discusses several papers in this vein; we therefore complement his discussion by giving three additional examples.

Chakraborty and Harbaugh (2007) provide a direct example of how multidimensional cheap talk creates new ways for the parties' interests to be aligned: they show that compar-

ative cheap talk, across dimensions of the state space, can be credible, even if cheap talk about any one dimension would not be. As a simple example, suppose that the state space is $S = S_1 \times S_2$, and the decision space is $D = D_1 \times D_2$, where the state is uniformly distributed on the unit square, each decision is from the real line, and the parties' payoffs are $U_2(s, d) = -(d_1 - s_1)^2 - (d_2 - s_2)^2$ and $U_1(s, d) = -(d_1 - (s_1 + b))^2 - (d_2 - (s_2 + b))^2$. That is, the sender's bias is b on both dimensions, and suppose $b > \frac{1}{4}$, so that cheap talk cannot be effective on one dimension alone.

Chakraborty and Harbaugh (2007) show, however, that comparative cheap talk can be effective. For example, consider the messages " $s_1 > s_2$ " and " $s_2 > s_1$ " (where we ignore equality, because it has zero probability). Because $E[\max\{s_1, s_2\}] = \frac{2}{3}$ and $E[\min\{s_1, s_2\}] = \frac{1}{3}$, in equilibrium, if the sender says " $s_1 > s_2$," then the receiver chooses $(d_1, d_2) = (\frac{2}{3}, \frac{1}{3})$, whereas if he says " $s_2 > s_1$," then she chooses $(d_1, d_2) = (\frac{1}{3}, \frac{2}{3})$. For a given state (s_1, s_2) , the sender's gain from inducing $(d_1, d_2) = (\frac{2}{3}, \frac{1}{3})$ instead of $(d_1, d_2) = (\frac{1}{3}, \frac{2}{3})$ is $\frac{2}{3}(s_1 - s_2)$, so such comparative cheap talk is indeed an equilibrium, and the expression for the sender's gain reveals why: $\frac{2}{3}(s_1 - s_2)$ is independent of the one-dimensional bias b .

Che et al. (2012) also consider comparative cheap talk. For example, suppose the receiver can choose either of two new projects, d_1 or d_2 , or the status quo, d_0 , so $D = \{d_0, d_1, d_2\}$. The parties' payoffs from the new projects depend on the state, but their payoffs from the status quo do not. Unlike in the other models we have discussed, the parties' payoffs are identical for either new project: $U_1(s, d_i) = U_2(s, d_i) > 0$ for $i = 1, 2$, which bodes well for communication. There is a difference in their payoffs from the status quo, however: the sender receives zero from the status quo, but the receiver receives a positive payoff, $0 = U_1(s, d_0) = U_{10} < U_2(s, d_0) = U_{20}$, which complicates communication. In particular, the sender would like to persuade the receiver to choose the superior new project (or, failing that, even the inferior new project), but the receiver might prefer the status quo.

As a numerical example, suppose that the payoff from d_1 is either 1 or 7, the payoff from d_2 is either 4 or 6, and all four payoff combinations are equally likely. In a full-revelation equilibrium, if the sender says " d_1 is better than d_2 ," then the receiver compares $E[U_2(s, d_1) | U_2(s, d_1) > U_2(s, d_2)] = 7$ to U_{20} , whereas if the sender says " d_2 is better than d_1 ," then the receiver compares $E[U_2(s, d_2) | U_2(s, d_2) > U_2(s, d_1)] = 5$ to U_{20} . For full revelation, where the sender's recommendation is always accepted, the receiver's payoff from the status quo must be sufficiently small: $U_{20} \leq 5$. For somewhat larger values of U_{20} , the sender panders to project 1, sometimes recommending it when it is the inferior new project. For example, if $U_{20} = 5.25$, then there is a mixed-strategy equilibrium where the sender recommends the superior new project unless the payoffs are (1, 4), in which case he recommends project 1 with probability $\frac{2}{3}$ (even though it is the inferior project). This strategy by the sender strengthens the receiver's belief about project 2 when it is recommended, so that the expected payoff from project 2 can compete with U_{20} . In response, the receiver always accepts project 1 but rejects project 2 in favor of the status quo with probability $\frac{2}{3}$.

Finally, in an intriguing model of the development and exercise of expertise, Callander (2008b) imagines that one dimension of the state is simple (e.g., what project to adopt), but another is potentially quite complex (e.g., expertise about the probability distribution over outcomes from adopting a project). Callander assumes that an expert can talk about the first dimension but not the second, and he finds that equilibrium communication depends importantly on

the correlation between the two. As we discuss in Section 5.2.1, these ideas have important implications for the feasibility and desirability of delegation. For now, however, we simply provide a third example of multidimensional cheap talk, inspired by Callander's (2008b: 124) critique of existing models in which "a single recommendation by an expert can render an intelligent layperson an expert."

Our setting is more abstract than Callander's. Here, the state space S consists of pairs $s = (s_1, s_2)$. The probability distribution over S puts zero weight on pairs (A, A) , (B, B) , and (C, C) , but puts positive weight on each realization of s_2 : $\Pr(s_2 = X) > 0$, $X = A, B, C$. Player 1 observes $\theta = s_1$, and his cheap-talk message a is an element of $\{A, B, C\}$. Think of it as his claim as to the value of s_1 . (For simplicity, we depart from Callander (2008b) by assuming that player 1 does not observe s_2 .) Player 2 then observes $\sigma = a$ and selects d , an element of $\{A, B, C\}$.

The sender wants the decision d to match the first element of the state, but the receiver wants to match the second. Because the two elements of the state are always different, the parties' interests are always conflicting in this respect. However, neither wants the decision to differ from both s_1 and s_2 , because both would then receive negative payoffs. Formally, if $d = s_1$, then $U_1 = b_1 > 0$ and $U_2 = 0$. Symmetrically, if $d = s_2$, then $U_1 = 0$ and $U_2 = b_2 > 0$. In addition, if d is not equal to either s_1 or s_2 , then both parties receive $K < 0$.

Depending on the probability distribution over states, it may be possible for the receiver to invert the sender's message if he reports honestly, as in Crawford and Sobel (1982). For example, suppose that, conditional on s_1 , the probability distributions over s_2 are degenerate, in each case placing probability zero not just on $s_2 = s_1$ but also on one of the two remaining options. For example, $\Pr(A|A) = \Pr(B|A) = 0$, $\Pr(B|B) = \Pr(C|B) = 0$, and $\Pr(C|C) = \Pr(A|C) = 0$, so that if the first element of the state is A , then the second is C ; if B , then A ; and if C , then B . If the sender announced truthfully, $a = s_1$, then the receiver could immediately determine the value of s_2 and choose accordingly, $d = s_2$. The sender would then have a profitable deviation from the action that induces $d = s_1$, so the equilibrium cannot involve full revelation by the sender.

Alternatively, suppose each of the alternatives to $s_2 = s_1$ is equally likely—for example, $\Pr(B|A) = \Pr(C|A) = \frac{1}{2}$. If the sender is presumed to report honestly, then the receiver faces a choice between accepting the sender's "recommendation" (setting $d = a$, which by hypothesis is equal to s_1) or picking one of the remaining two options, each of which is equally likely to be the receiver's preferred decision. The first alternative yields $U_2 = 0$, whereas the second gives expected value $EU_2 = (b_2 + K)/2$. If K is sufficiently negative, then the first alternative is more attractive, and the receiver simply rubber-stamps the sender's choice, knowing it is surely not the option she would prefer. In this case, a single recommendation by the sender does not render the receiver an expert.

3.2. More Parties

Many organizational settings involve more than the two parties envisioned in the basic model in Section 2.1. Here we briefly discuss three such settings: more senders, more receivers, and more of both (e.g., because each party is both sending and receiving).

3.2.1. More Senders

In a cheap-talk setting, when two senders share the same information but different preferences about how the decision should be tailored to the state, there is an opportunity for their separate reports to discipline each other, together providing more information to the receiver than had either sender reported alone. In applied models, Gibbons (1988) and Gilligan and Krehbiel (1989) study arbitration and legislative committees, respectively. Krishna and Morgan (2001) and Battaglini (2002) analyze more abstract cheap-talk models, and Milgrom and Roberts (1986) analyze competition among senders in a verifiable-information game. See Sobel (2012) for a review of these and subsequent papers.

3.2.2. More Receivers

Complementing the idea that two senders might discipline each other's reports in a cheap-talk game, one can also imagine how two receivers with different preferences about how the decision should be tailored to the state could discipline the report sent by a single sender. In addition, with two receivers, one can analyze whether communication should be in public, in private, or both. Farrell and Gibbons (1989) offer a simple first analysis of these issues; Goltsman and Pavlov (2011) provide a considerable generalization.

Switching from cheap talk to verifiable information, Caillaud and Tirole (2007) develop another model with multiple receivers, asking how a project's sponsor, who is known to benefit if the project is implemented, can recruit necessary support from potentially skeptical group members. Potential applications include flat organizations—such as a partnership, committee, or consortium—as well as such settings as a matrix organization, where an employee has multiple bosses.

In the Caillaud and Tirole model, implementing the project may or may not be in the interest of a given group member. To attempt to gain a member's support, the sponsor can supply partial but verifiable information to the member, who can then further investigate the project at a cost. A member with a sufficiently strong prior that the project is in her interest does not find it worthwhile to investigate (because the resulting information is too unlikely to sway the member's decision), whereas a member with an intermediate prior does find it worthwhile to investigate, and a member with a sufficiently weak prior again does not find it worthwhile to investigate (because the resulting information again is too unlikely to sway the member's decision).

Members potentially differ in the payoffs they would receive if the project were implemented, but they know that their payoffs are correlated, so learning about another member's payoff would be useful. As a result, the sponsor may attempt to induce a "persuasion cascade." For example, the sponsor may supply information to a slightly skeptical first member, in the hope that this member's endorsement may then persuade a more skeptical second member, even if the sponsor herself could not have persuaded this second member directly. Such persuasion by the first member may result in the second member either investigating when she would originally have rejected the proposal outright or supporting the proposal without investigation when she would originally have insisted on investigating. Alternatively, sometimes the optimal cascade for the sponsor is to supply information to the more skeptical member first, because her approval may then convince the less skeptical member to support the proposal without investigation.

3.2.3. More of Both

A third setting may be especially natural and interesting, where some or all parties both have information and control a decision, so they have both sender and receiver roles. Because this literature is fairly new, we start by noting recent contributions from team theory, where the parties share common interests, such as Dessein and Santos (2006) and Calvo-Armengol et al. (2011). This work serves as a benchmark for models that introduce strategic considerations.

Okuno-Fujiwara et al. (1990) analyze communication of verifiable information in a Bayesian game, where different players have private information and control their own actions in the game. Hagenbach and Koessler (2010) and Galeotti et al. (2011) explore related settings with cheap-talk communication. In the latter papers an interesting issue arises because, building on the multisender case, the effectiveness of one party's cheap-talk communication with a second party now depends not only on the divergence in these parties' preferences but also on how communication from a third party may already have changed the second party's beliefs and hence in effect her preferences. This logic has implications for the equilibrium level of communication in a network, for example.

3.3. More Periods

Another stark limitation of our basic model is that it has only one period, thereby eliminating reputation effects and other consequences of learning.

In a two-period model with a unique (or, at least, specified) second-period continuation equilibrium as a function of first-period outcomes, one can substitute the second-period equilibrium payoffs into the first-period analysis, thereby hinting at a dynamic analysis in a one-period model (albeit one with possibly unusual payoff functions). Ottaviani and Sorensen (2006a,b) take this kind of reduced-form approach to analyzing how first-period cheap talk by a professional advisor can affect the advisor's second-period reputation. Prendergast (1993) uses a similar reduced-form modeling strategy to analyze "yes men"—that is, workers whose incentives are to distort their reports in the direction of what they believe their manager wants to hear.

Sobel (1985), Bénabou and Laroque (1992), and Morris (2001) take the next step, building dynamic models that repeat our basic model but also include a permanent type for the sender. Their models analyze cheap talk; more could be done by building on the other approaches discussed in Section 2.

3.4. Reinterpretation: From Choice to Execution

Van den Steen (2010b) makes a fundamental observation about the large literatures on authority and delegation—surveyed, respectively, by Bolton and Dewatripont (this volume) and in Section 5: almost without exception, these models do not involve a principal giving an agent an order. Rather, these literatures focus on what should be called "decision" authority, to distinguish it from what Van den Steen calls "interpersonal" authority, where the former analyzes who should control a decision right and the latter analyzes when orders will be given and whether they will be followed.

As Barnard (1938) emphasizes, orders are not simple—either to enforce or, therefore, to model:

A most significant fact of general observation relative to authority is the extent to which it is ineffective in specific instances. [Barnard 1938: 161]

The decision as to whether an order has authority or not lies with the persons to whom it is addressed, and does not reside in "persons of authority" or those who issue these orders. [Barnard 1938: 163]

I suppose all experienced executives know that when it is necessary to issue orders that will appear to the recipients to be contrary to the main purpose, especially as exemplified in prior habitual practice, it is usually necessary and always advisable, if practicable, to explain or demonstrate why the appearance of conflicts is an illusion. . . . Otherwise the orders are likely not to be executed, or to be executed inadequately. [Barnard 1938: 166]

Motivated by these cautionary tales, we now reinterpret the basic model from Section 2.1 as describing the choice and execution stages of Mintzberg's decision process in Figure 1. Our main purpose in doing so, however, is to surface these issues for further modeling, not to propose this simple model as anything beyond a first attempt.

Although most of our interpretations of the models in Section 2 have so far been that player 1 (who seeks to influence a decision) is the agent and player 2 (who controls the decision) is the principal, we now reverse the roles: player 1 is the principal, whose influence action a is an order (in the suggestive sense of the term emphasized by Barnard), and player 2 is the agent, whose decision d concerns whether or how well to execute the order. In addition to reinterpreting the players, we also impose new structure on the payoffs. For example, it seems important to distinguish between orders and advice (even if the advice comes from the principal), so we interpret the cheap-talk model in Section 2.2, where the principal's action a does not affect the parties' payoffs, as more the latter than the former. Also, the typical order does not seem to be directly costly to the principal, in the sense in which signal jamming and signaling involved a cost $k(a)$ in Sections 2.3 and 2.4, so we do not directly use those models, either.

We therefore imagine a setting where neither the principal's order (a) nor the agent's execution of that order (d) has much impact on the parties' payoffs on its own. Instead, the best payoffs arise when the order and the execution are well matched—to each other and also to the state s . There are of course many ways to formalize ideas like these, including not only Van den Steen (2010b) but also Landier et al. (2009) and Marino et al. (2010). The approach described below is inspired by Blanes i Vidal and Möller (2007). The resulting model does not focus on an order by the principal but instead explores another aspect of the choice and execution stages of Mintzberg's (1979) decision process, where the principal's action is a project choice, and the agent's decision is execution effort.

To keep things simple, suppose that there are only two states, $s = A$ or B . There are also two projects, but the parties can achieve positive payoffs only if the project is appropriate for the state. In addition, the parties are more likely to realize these positive payoffs if the agent's execution effort is high.

The principal's signal θ consists of two parts, public and private, both of which are noisy measures of the state. That is, $\theta = (\theta_{\text{pub}}, \theta_{\text{prt}})$, where $\theta_i \in \{A, B\}$, and θ_{pub} and θ_{prt} are conditionally independent given the state. To keep things interesting, suppose that the principal's private signal is more accurate: $\text{Prob}(\theta_i = s | s) = p_i$ and $p_{\text{prt}} > p_{\text{pub}} > \frac{1}{2}$.

The principal's action is a project choice, $a = A$ or B , and the agent's signal then consists of the public signal and the principal's project choice, $\sigma = (\theta_{\text{pub}}, a)$. After receiving her signal, the agent chooses execution effort $d \in [0, 1]$ at convex cost $c(d)$. Finally, the project's output is $y = L$ or H , where $y = L$ if $a \neq s$ but $y = H$ with probability d if $a = s$. For simplicity, the parties split the project's output, so the payoffs are $U_1 = \frac{1}{2}y$ and $U_2 = \frac{1}{2}y - c(d)$.

Because the principal's private signal is more accurate, the first-best project choice is for the principal to follow his private signal, $a = \theta_{\text{priv}}$. Because the public and private signals are conditionally independent, the first-best execution effort by the agent then depends on both signals. In particular, first-best effort is higher when the signals agree with each other and with the principal's project choice.

In equilibrium, however, the principal may not follow his private signal in choosing the project, exactly because the agent's execution effort is higher when she thinks that the project and both signals all agree. Instead, if the private signal is not too much more accurate than the public one, in equilibrium the principal may choose the project indicated by the public signal when the public and private signals disagree. Of course, the agent understands this and chooses her effort based on a correct belief about whether the project matches the private signal. Such a model thus takes a first step toward Barnard's (1938: 166) observation that contrary "orders are likely not to be executed, or to be executed inadequately."

3.5. Enrichment: From Information to Advice to Choice

Casual observations suggest that shareholders often rubber-stamp the board's decisions, the board rubber-stamps the CEO's decisions, the CEO rubber-stamps the division managers' decisions, and so on. Aghion and Tirole (1997) model such outcomes by distinguishing between formal authority held by an owner versus real authority held by an employee. Their model expands the basic model from Section 2.1 to include information acquisition as well as advice and choice.¹⁰

Suppose there are three possible projects, indexed by $k = 1, 2, 3$. Project k delivers benefits B_k to the owner and b_k to the employee. One project is terrible for both parties: $B_k = b_k = -\infty$. The other two projects deliver benefits of 0 and $B > 0$ to the owner and 0 and $b > 0$ to the employee. With probability α , the payoffs from the latter two projects are (B, b) and $(0, 0)$; with probability $1 - \alpha$, the payoffs from these projects are $(B, 0)$ and $(0, b)$. Thus, a higher value of α means that the owner's and the employee's interests are more likely to be aligned.

The problem is that, initially, neither the owner nor the employee knows which project is which. That is, each party can see the three possible projects, but neither party knows which project is the terrible one or which is the good one for him or her. Because of this uncertainty (and the severity of the terrible project), if no information is collected about which project is which, neither party will want any project to be chosen.

Both the owner and the employee can try to collect information about which project is which, but at a cost. If the owner incurs the cost $c_B(E)$, then the owner learns her own payoff on each project with probability E but learns nothing with probability $1 - E$. Similarly, if the employee

incurs the cost $c_s(e)$, then the employee learns his own payoff on each project with probability e but learns nothing with probability $1 - e$ (where these events for the employee are independent of those just described for the owner).

To analyze the parties' incentives to collect information, consider what happens in the following three situations. First, if the owner becomes informed about her payoffs on the three projects, then she will choose the project that pays her B . Second, if the owner remains uninformed but the employee becomes informed, then the employee will recommend the project that pays him b ; even though the owner cannot observe the projects' payoffs (and, by assumption, the employee cannot communicate them), the owner will accept the employee's proposed project, because its expected payoff to the owner is $\alpha B + (1 - \alpha) \times 0 = \alpha B > 0$. Finally, if neither the owner nor the employee becomes informed, then neither party will want any project to be chosen, as described above.

From these three situations, we can write the expected payoffs to the owner and the employee if the owner incurs the cost $c_B(E)$ and the employee incurs the cost $c_s(e)$:

$$U_B = EB + (1 - E)e\alpha B - c_B(E)$$

and

$$U_s = E\alpha b + (1 - E)eb - c_s(e),$$

respectively. From these expected payoffs, we can compute the Nash equilibrium choices of E and e . Assuming that the cost functions have the usual properties (namely, convexity, zero slope at zero, and infinite slope at one), the Nash equilibrium (E^*, e^*) satisfies:

$$B(1 - e\alpha) = c'_B(E)$$

and

$$b(1 - E) = c'_s(e).$$

In our view, the Aghion and Tirole (1997) model is an elegant and insightful explanation of equilibrium rubber-stamping: with probability E^* , the owner becomes informed and chooses the project; but with probability $(1 - E^*)e^*$, only the employee becomes informed, in which case he proposes a project, and the owner rubber-stamps the proposal. This rubber-stamping can be seen as the difference between formal and real authority: the owner does not cede formal authority, but rather cedes real authority when the employee has superior information and sufficiently similar preferences. (To see the role of similar preferences, imagine revising the model to allow projects with payoffs $B_k = -\infty$ and $b_k = b$: now the owner would not dare to cede real authority.) For further approaches to this section's information-acquisition issue, see the discussions of Che and Kartik (2009) and Rantakari (2012) in Section 5.1.4.1.

3.6. Enrichments: Contractibility

We emphasized in Section 2.1 that nothing in our basic model was contractible: not player 1's action, player 2's decision, or the parties' payoffs, so there cannot be bargaining over either player's move, mechanisms to elicit their signals, or outcome-contingent incentive contracts. As we explained, we made this assumption not only for simplicity but also to some extent for

10. Parts of this section draw on Gibbons (2010), with permission from the *Annual Review of Economics*, Volume 2, © 2010 by Annual Reviews, <http://www.annualreviews.org>.

realism, given that our focus is on organizations. That is, if everything were contractible, then any governance structure would suffice, as argued by Coase (1960). And, loosely considering settings somewhat shy of that extreme, if contracts were only slightly imperfect, then markets might do fine at inspiring and implementing the acquisition and dissemination of information, as argued by Hayek (1945). As Arrow (1974: 33) puts it, however, "organizations are a means of achieving the benefits of collective action in situations where the price system fails." And recall from Section 1 March's (1962/1988: 110–112) view of "the business firm as a political coalition" in which "the executive . . . is a political broker" who cannot "solve the problem of conflict by simple payments to the participants and agreement on a superordinate goal."

In this section, therefore, we begin to relax our stark assumption in Section 2.1 that nothing is contractible, but we hasten to add that this must be done with care. That is, if too much is contractible, then the transaction should probably be conducted in a market rather than in an organization. Analyzing this macro-level choice of governance structure, between organization and market, is of course both important and well beyond the scope of this chapter; see Bresnahan and Levin (this volume) and Tadelis and Williamson (this volume).

As a first exploration of contracting issues, we return to the information-acquisition problem of Section 3.5. Lambert (1986) and Demski and Sappington (1987) were the first to analyze models in which player 1 does not have a direct interest in player 2's decision (i.e., player 1's payoff is independent of d) but can exert effort to improve the quality of his information. To induce such effort, the parties can write contracts that specify payment schemes for the expert that are contingent on the outcome of the decision. As in the standard agency framework, the only source of friction in these models is that the expert must be induced to exert effort. In contrast to the standard framework, however, here effort generates information rather than output. As Dewatripont and Tirole (1999) show, when effort is about acquiring information, the Monotone Likelihood Ratio Property is lost. A major insight of this literature is that, to provide better incentives for information acquisition, a contract may have to create an endogenous conflict of interest regarding the decision to be made; see, for example, Inderst and Klein (2007), Malcomson (2009, 2011), and Zermeno (2012).

Another organizational literature that assumes some contractibility studies committees. Indeed, such models are close to those described in Section 3.2, where more parties are added to our basic model. For example, each committee member might receive private information, then engage in cheap talk, and finally vote; see Feddersen and Pesendorfer (1998), Persico (2004), Levy (2007), and Visser and Swank (2007) for related models. A key assumption in such models, however, is that the eventual decision is a predictable consequence of the votes, such as by majority rule, thus raising the question of whether the parties might instead use direct-revelation mechanisms, mapping type reports into decisions, rather than use voting schemes. See Mookherjee (this volume) for more on this approach.

Another approach is to assume that, although the state s and the decision d are non-contractible, messages are. For example, in a cheap talk equilibrium with a finite number of messages, one could imagine payments to the sender to be contingent on the message he sends. Krishna and Morgan (2008) examine such a model, assuming that the sender has limited liability (i.e., cannot be forced to pay money to the receiver). They then show that full revelation is always feasible but never optimal. Instead, for the standard example with uniform state s , quadratic loss functions, and sender bias $b > 0$, they show that, because of the standard argu-

ment that integrates the local incentive-compatibility constraints, inducing revelation by high types requires large payments to lower types if they are to reveal themselves. As a result, the optimal payment and communication scheme has positive payments and full revelation for low states but no payment and no communication for high states. We see this kind of analysis, where limited contractibility relaxes the stark assumption in Section 2.1 that nothing is contractible, as a productive vein for further work.

3.7. Contested Control

As a transition to our discussion in Sections 4 and 5 of decision architectures that are designed ex ante, here we consider situations where decision rights are not allocated ex ante but are contested ex post. Tullock (1980) provided an elemental model of rent seeking in such an environment. Two contestants (A and B) purchase lottery tickets in amounts x_A and x_B , respectively, each hoping to win a prize of \$100. Contestant i 's expected payoff is then $x_i / (x_i + x_j) \times 100 - x_i$, so the symmetric Nash equilibrium has positive expenditures on lottery tickets, but in this equilibrium the contestants have the same probabilities of winning as if neither had spent anything—namely, one-half. Skaperdas (1992) and Rajan and Zingales (2000) used similar models to analyze battles for control (in organizations and other settings).

As a simple version of some of Skaperdas's ideas, suppose that each of two parties ($i = 1, 2$) is endowed with resources r_i that can be spent either on producing output (y_i) or on developing "arms" (a_i) that can be used in a later conflict over dividing the aggregate output. Specifically, let $y_i = y(r_{yi})$ and $a_i = a(r_{ai})$, where $r_{yi} + r_{ai} = r_i$, and suppose that the payoff to party i is $p(a_i, a_j)[y_i + y_j]$, where $p(a_i, a_j)$ can be interpreted as the probability that party i wins the conflict and so receives the aggregate output. Skaperdas derives conditions under which the equilibrium allocations of the parties' resources into outputs and arms are efficient (i.e., both parties put all their resources into output) or partially so (one party puts all its resources into output). More often, however, as in Tullock's model, the equilibrium is inefficient and sometimes dramatically so.

Rajan and Zingales analyze a two-period model that can be seen as having fully contractible trade in the first period and Skaperdas's model in the second. Trade in the first period endogenizes heterogeneous resource endowments (r_i) and heterogeneous production functions (now $y_i(r_{yi})$) for the second period. The second-period consequences of this heterogeneity can be so severe that the parties forgo a contractible and seemingly Pareto-improving trade in the first period. For example, consider two divisions in a company (or two adjacent countries). Suppose that party i has a machine that it could sell to party j , and suppose that party j is already more productive than party i and so could make better use of the machine ($y'_j > y'_i$). Selling the machine increases i 's endowment for the second period (r_i increases) but reduces its productivity (y'_i falls), and the reverse holds for j . What then happens to the equilibrium allocations of the parties' resources into outputs and arms? Loosely, reducing i 's productivity increases i 's incentive to invest in arms, and increasing j 's productivity (more precisely, j 's output) also increases i 's incentive to invest in arms, but an increase in a_i then increases j 's incentive to invest in arms, and so on. In equilibrium, these increases in arms (and reductions in output) can be severe enough to overwhelm the supposed efficiency advantages from allocating the machine to the more productive party j , so no trade occurs in the first period.

In all these rent seeking models, control is valuable, because it allows one party to consume the prize in question—\$100 in Tullock's (1980) model and aggregate output in those of Skaperdas (1992) and Rajan and Zingales (2000). As a slight variant on this approach, suppose instead that control allows one party to take a decision (e.g., about the organization's strategy). The argument is then similar, with the new feature that different parties may place different values on being able to take the decision in question, thus creating different incentives for gaining control.

As a concrete example, consider the following game (which is intended to illustrate some of the central governance challenges in decision architectures involving shared control, e.g., joint ventures, consortia, and matrices): (1) the parties publicly observe the state of the world, $s \in S$; (2) the parties engage in a battle for control (specifically, to keep things simple, we follow Tullock—each party expends x_i for probability $x_i/(x_i + x_j)$ of winning control); (3) the party who wins control then chooses a decision $d \in D$; and (4) the parties receive payoffs $U_i(s, d)$, gross of their expenditures in stage 2. Now the prize to party i for winning control is not \$100 but rather $U_i(s, d_i^*(s)) - U_i(s, d_j^*(s))$, where $d_i^*(s)$ maximizes $U_i(s, d)$ over $d \in D$. Because the prizes are asymmetric, so are the equilibrium expenditures in the battle for control: the party with the larger prize fights harder. This example is of course extremely stark, but it nonetheless captures some of the tensions that commonly arise under shared control.

4. Endogenous Architectures: Fixed Decision Allocations

In the previous sections we explored the inefficiencies that can arise in organizations from strategic communication under fixed decision processes and (except for Section 3.7) a fixed allocation of decision rights among the interested parties. One reason to have done so was to focus on situations where the decision rights are inalienable. For instance, the right to appoint a CEO in a corporation is vested in the board of directors. The board can seek advice on its choice, and it may even sometimes rubber-stamp the advisors' recommendation, but it cannot formally delegate the actual decision. One can also view our discussion in the previous sections as having been about the second stage of a two-stage game in which the owner of the organization allocates decision rights in the first stage and the different members of the organization take their actions in the second.

In this section we turn to actions that the owner can take in the first stage to mitigate the inefficiencies of organizational communication and decisionmaking in the second stage. In particular, in Section 4.1 we explore three suggestions put forward by Milgrom and Roberts (1988) aimed at controlling influence activities:

1. the owner could close communication channels, at least for some decisions;
2. she could commit not to respond to information provided by others; or
3. she could try to align incentives through compensation, promotion, and similar policies, including altering the decision processes and even the boundaries of the firm itself.

In Section 4.2, we move away from the rich organizational descriptions that characterize the discussion in Section 4.1 to a more abstract formulation, based on our model from Section 2, that allows us to explore just what responses may be feasible and effective.

Alternatively, of course, the owner could simply delegate the decision directly to the employee. Although delegation might be considered an instance of the Milgrom-Roberts's third option (change the organizational design), it has been the focus of most of the recent formal literature on responses to inefficiency in decisionmaking. Thus we make it the focus of Section 5.

4.1. Organizational Responses in the Case of Fixed Decision Rights

The first published work on influence costs is Milgrom (1988). A simplified version of the model he considered is the following. There are a manager and an employee, both risk neutral, and the former contracts with the latter to deliver unobservable effort. The employee can allocate a fixed amount T of effort between productive activities t and influence x . (Effort is thus not directly costly to him.) The manager may be able to make a change decision that gives benefits k to the employee that are, in effect, a transfer from another (unmodeled) employee. Thus, k enters the employee's utility function but not the aggregate welfare. Implementing the change brings benefits $r(x)I$ to the manager.

If the manager is allowed discretion over adopting the change, then the probability that she does adopt it is a function $q(x)$ that is assumed to be increasing. The employee's productive effort, $t = T - x$, produces a benefit π with a probability $p(t)$ that is increasing in t , with $p(0) = 0$. The problem is to determine a wage for the agent as a function of whether the benefit π was realized and whether the change was implemented, with the objective to maximize total surplus, namely, $p(t)\pi + q(x)r(x)I$, subject to t and x maximizing $p(t)q(x)w_{11} + p(t)(1 - q(x))w_{10} + (1 - p(t))q(x)w_{01} + (1 - p(t))(1 - q(x))w_{00} + q(x)k$ given $t + x = T$. Meanwhile, if the manager does not have discretion, then the payoff is simply $p(T)\pi$, because the agent has no incentive to exert influence and so can be assumed to devote all his time to production.

The key result is that it may be optimal to deny the decisionmaker discretion over implementing the change: there are values of I and k such that eliminating discretion (i.e., fixing $q(x)$ at zero rather than using pay to deter influence) is optimal, and this is true for all higher values of k (when redistribution effects are greater) and lower values of I (when the change is less valuable). This is an example of Milgrom and Roberts's (1988) second suggestion: limiting discretion can be valuable.

The logic behind the q and r functions is not explored in Milgrom (1988). Note, however, that delegation might be an attractive alternative here, especially if $r(x)$ is not increasing. In this case, if the implementation decision could be and were delegated, then the first best would be achieved.

Milgrom and Roberts (1988, 1999) considered a problem with two employees, each of whom has to allocate a fixed amount of time between influence activity and productive effort, with the principal unable to observe their choices. The principal values the output that results (probabilistically) from their efforts, but she also wants to promote one of the two workers to another job, where his underlying, unknown ability will be valuable. (Ability does not influence productivity in the current job.)

The new job is desirable to the workers in that it must, for competitive reasons, carry a wage that exceeds the outside option of the person appointed to it. There is symmetric information about the workers' abilities, which are identically distributed. The impact of influence activity

and determination of its equilibrium level are not modeled in detail. Instead, the influence activity is assumed to "build credentials," making the worker's ability appear greater but also providing information about ability. One possibility that Milgrom and Roberts (1999) suggest is a test on which promotion might be based, with the employees' influence activities consisting of diverting time to prepare for the test. Simply taking the test generates information, but once the test is scheduled and its results will be used in the promotion decision, there will be incentives to try to increase one's performance on the test.

In this context, the principal has to decide on wage payments for the current period, wages to pay to each employee once the promotion decision is made, and the criteria to use in promotion.

Milgrom and Roberts (1988) consider only a limited class of possible decision rules, to any of which the employer can commit. Even so, a full solution to this design problem has not been obtained. However, several points can be established. First, if the importance of assigning the better person to the new job is sufficiently low and the productivity of effort is high, then the employer will choose not to use apparent qualifications at all in the promotion decision. In this case, both employees will devote full time to work. Further, no incentive pay will be needed (because the work is not itself costly, and there is no reward to diverting time to influence). Promotion will be random. This can be interpreted as an instance of option 1, cutting communication.

If the marginal impact of employee quality on performance in the new job is sufficiently high, however, then the employer will want to use apparent qualifications in the promotion choice. The danger is now that, unless there is a reward to productive effort, the employees will concentrate solely on building credentials. This in turn results in several possible adjustments. One is to pay a bonus for individual output, even though there is no inherent cost of effort. A second is to raise the wage of the employee who is not promoted closer to the level paid the promoted employee. This adjustment reduces the return to being promoted and thus the incentives to exert influence to get a raise (this works only once there is pay for performance). Finally, it may be worthwhile to promote partially on current performance, even though it is not informative about the candidate's actual ability, because this lowers the amount of performance pay necessary to induce productive effort. Thus, all aspects of the organizational design can be altered to limit the cost of influence activities, in line with the third suggestion of Milgrom and Roberts (1988).¹¹

An important organizational decision that has been much studied, both theoretically and empirically, is the allocation of capital inside firms. In particular, Meyer et al. (1992) modeled a situation in which failing businesses in a multidivisional firm exert influence to try to obtain more corporate resources to keep themselves alive. In equilibrium, the top decisionmaker is not fooled by these attempts, so the allocation of capital is not distorted. As in the signal-jamming model in Section 2.3, however, resources are wasted on attempts at influence. The solution Meyer et al. suggested is to spin off units that are prone to attempt influence.

More recently, several papers, beginning with Rajan et al. (2000), Scharfstein and Stein (2000), and Wulf (2002), have examined in more detail situations where lower level managers seek to use strategic communication so that more capital is allocated to them than is optimal for the firm. In these models, the allocation rules are distorted to reduce the costs of such strategic behavior. The resulting misallocation is costly. In addition, Wulf (2002, 2007) investigated the

interplay between the design of capital budgeting processes and the basis for and extent of performance pay in light of influence activities, and in Wulf (2009) she documents patterns in these variables consistent with limiting the costs of influence. For a discussion of this literature, including newer contributions, see Gertner and Scharfstein (this volume).

Finally, Powell (2012) analyzed a rich formal model of influence, building on the signal-jamming models in Holmström (1982/1999) and Gibbons (2005) discussed in Section 2.3. In Powell's framework, there are two parties and two decisions. Before these decisions can be made, the parties decide who will be empowered to make each decision, but once this allocation of decision authority is set, contracting difficulties prevent renegotiation. If the other party is to make a decision, then the disempowered party may have incentives to try to affect the decision by undertaking (unobserved) influence activities.

Given an arbitrary allocation of decision rights, Powell solves for the equilibrium level of influence. In equilibrium the party making a decision will correctly anticipate the amount of influence the other party is undertaking and, using that correct estimate, will back out the effect of the influence and make the decision that is best for him. So, again, the influence does not distort decisions, but it does occur and it imposes costs.

The optimal allocation of decision rights maximizes the sum of utilities, given that the parties will choose their levels of influence in the equilibrium fashion. This allocation may consist of either divided control, with each controlling one decision (which is advantageous, because the costs of influence are convex, and so it is better to have the total amount of influence split between the two parties), or unified control under the party who cares more about the outcome.

Powell then examined some suggestions made by Milgrom and Roberts (1988) for limiting influence activities and mitigating influence costs. A "closed-door" policy of limiting communication is modeled as no public signal being observed (because the parties agreed not to hire the needed staff, loaded their calendars so they were too busy, and limited their meetings). This policy eliminates the disempowered party's incentive to attempt influence, but it also deprives the party in control of useful information. However, if the initial uncertainty was not great, then not much is lost, and the closed-door policy may be preferred to the alternative of playing the game as originally specified.

Also, if there is a lot of disagreement about how to tailor decisions to the state of the world, then again it may become worthwhile to adopt the closed-door policy. Then a selection story applies: as disagreement increases, both the open- and closed-door policies do worse, but the closed door does relatively better at high levels of disagreement. Thus, the rigid, apparently bad management practice of limiting communication obtains when circumstances are bad, and observed firms with these practices are worse performers than those that are observed adhering to open-door policies. However, the closed-door firms perform better in the environments they inhabit than would firms with apparently "better" management practices. See the right-hand side of Figure 1 in Gibbons and Roberts ("Economic Theories of Incentives in Organizations," this volume) for a graphical representation of this selection argument.

4.2. An Abstract Approach

Having analyzed potential inefficiencies in communication and decisionmaking under fixed decision processes in Sections 2 and 3, our focus here in Section 4 has been on ways the parties might design their environment *ex ante* so as to ameliorate some of these inefficiencies. Similar to

11. Note that delegation is not a real option in this model.

the way that the Milgrom and Roberts (1988) focus on the costs of influence activities animated many models we discussed in Section 2, their suggestions of strategies for reducing influence inefficiencies motivated the organizational responses described in Section 4.1.

Consistent with Milgrom and Roberts (1988), the models in Section 4.1 are detailed enough to suggest and analyze concrete organizational responses to influence costs. But each considers a separate organizational environment, making it hard to assess whether the full range of possible interventions has been considered. In this section, therefore, we pursue an abstract approach that we see as complementary to the more grounded analyses in Section 4.1. This abstract approach allows us, for a simple environment, to characterize a set of responses to influence inefficiencies that are collectively exhaustive and, in a sense, mutually exclusive.

The environment we consider in this section is the basic model from Section 2.1. For this model we follow Kolotilin et al. (2012) by describing five ways that the parties might respond to influence activities by reshaping their environment *ex ante*. (Kolotilin et al. do this for a cheap-talk model, but we apply their ideas to our basic model, thus covering not only cheap talk but also the other four illustrations discussed in Section 2.) As we discuss below, these five potential responses to influence inefficiencies overlap partially with both the three Milgrom-Roberts approaches described in Section 4.1 and the delegation models described in Section 5, but they also include other options, so we discuss all five here, as well as note papers that investigate these five options.

Imagine enriching our basic model to add a nonstrategic player who accepts an informational input and returns a (possibly stochastic) output according to some prespecified mapping. This nonstrategic player can be seen as a machine, which could be designed by the parties *ex ante* and then inserted into various moments of our basic model. Proceeding sequentially through the timing of our basic model, there could then be the following five interventions (after each of which the timing of the basic model resumes):

- a. "information control," where the new player, instead of the sender, observes the sender's signal θ in stage 2 of the model in Section 2.1 and delivers a signal θ' to the sender in stage 2.5, according to the distribution $\gamma(\theta'|\theta)$, after which the sender chooses the influence action a in stage 3;
- b. "influence control," where the new player replaces the sender, receiving the signal θ in stage 2 and choosing the action $a \in A$ in stage 3 according to the distribution $q(a|\theta)$, after which the receiver receives the signal σ in stage 4;
- c. "communication control," where the new player, instead of the receiver, observes the receiver's signal σ in stage 4 and then delivers a signal σ' to the receiver in stage 4.5 according to the distribution $\eta(\sigma'|\sigma)$, after which the receiver chooses a decision $d \in D$ in stage 5;
- d. "decision control," where the new player replaces the receiver, receiving the signal σ in stage 4 and choosing a decision $d \in D$ in stage 5 according to the distribution $r(d|\sigma)$, after which the parties receive their payoffs; and
- e. "execution control," where the receiver observes the signal σ in stage 4 and then in stage 5 recommends a decision $d \in D$ to the new player, who then chooses a decision $d' \in D$ in stage 5.5 according to the distribution $\rho(d'|d)$, after which the parties receive their payoffs.

To interpret just one of these five interventions, "execution control" includes the case where the parties agree *ex ante* that the receiver can choose from only a specified set of decisions *ex post*. In this abstract approach, there are no natural limits on the size or composition of such a decision set; deriving such limits is one of the complementary benefits one can get from the more grounded models in Section 4.1.

We note that although these five interventions are collectively exhaustive (in that there are no other moments in the timing of our basic model at which one could intervene), they are in some sense not mutually exclusive, because influence control and decision control, where the new player replaces the sender or receiver, weakly dominate the other options, where the new player is injected into the information flow of the model. In particular, influence control weakly dominates information control, and decision control weakly dominates both communication control and execution control.¹² However, it may be impossible to replace a player in intervention b or d, but it may still be possible to intervene as envisioned in a, c, or e, so we present all five interventions here.

Compared with the approaches that Milgrom and Roberts (1988) consider, communication control encompasses their first approach (reducing communication channels), and execution control encompasses their second (limiting the receiver's discretion to respond), but their third approach (changing the alignment of interests by changing organizational policies) does not fit easily in this framework.

One way to interpret their third approach is as changing the parties' payoff functions, such as modifying the bias b in the quadratic-loss functions illustrated throughout Section 2. Although literally changing utility functions is perhaps outside the standard toolkit of economics, their third approach can instead be interpreted in more standard terms by imagining that there are multiple periods, as described in Section 3.3, and the parties' payoff functions at the end of this basic model reflect the continuation equilibrium for the remainder of the game. Under this interpretation, the Milgrom-Roberts ideas of changing organizational policies for the remainder of the game could well be represented in our basic model as changing the bias b . Absent a grounded model like those in Section 4.1, however, it is unclear what changes in b can be accomplished and at what costs, so we sidestep this question for the remainder of this section.

In addition to the three Milgrom-Roberts approaches discussed in Section 4.1, we also discuss delegation as a response to influence inefficiencies in Section 5, which relates to decision control, and influence control is akin to the strategic management of public information model described in Section 2.6. Finally, there is information control, which is related to our discussion in Section 3.5 of the information-to-advice-to-choice parts of Mintzberg's (1979) decision process. In particular, whereas Aghion and Tirole (1997) envision keeping the receiver busy, so as to increase her marginal cost of becoming informed and thus empower the sender, here one can imagine keeping the sender busy, so as to increase his marginal cost of becoming informed and then being motivated to engage in influence activities.

There are several papers that analyze some of interventions a–c for cheap-talk games, including Green and Stokey (2007) and Ivanov (2010) on information control; Krishna and Morgan (2004), Blume et al. (2007), and Goltsman et al. (2009) on communication control;

12. For example, to see that delegation weakly dominates limited authority, note that if the receiver's strategy under limited authority is $d(\sigma)$, then delegation can mimic limited authority by choosing the rule $r(d'|\sigma) = \rho(d'|\sigma)$.

and Kolotilin et al. (2012) on execution control, as well as many of the papers on delegation discussed in Section 5. It would be interesting to explore some of these interventions in the other settings illustrated in Section 2, beyond cheap talk. Regardless of the setting, we see this framework as complementing Section 4.1 in two ways: it exhaustively categorizes what interventions are possible, and it allows analysis of the maximum reduction in influence inefficiencies that any one of these interventions can achieve.

5. Endogenous Architectures: Delegation

In Sections 2 and 3 we explored the inefficiencies that can arise in organizations from strategic communication under fixed decision processes. In Section 4 we then discussed possible organizational responses that attempt to reduce these inefficiencies. One such response—which was ignored in Section 4 but is prominent in the literature—is delegation of the decision right to the employee. In this section we survey this literature on delegation.

We start by looking at an organization that consists of an owner and an employee. As in the previous section, the organization has to make a decision that is both *ex ante* and *ex post* noncontractible. The parties, however, can now contract over the right to make the decision. At the beginning of the game, the owner offers the employee a contract that specifies a wage and an allocation of the decision right. If the employee rejects the offer, both parties realize a zero outside option. For the time being, we assume that neither the wage nor the decision right can be made contingent on any information that is revealed during the game. Moreover, we assume that if the employee accepts the contract, it cannot be renegotiated at any point.

In the next section we discuss some of the main reasons for delegation that have been examined in the literature. The models that we discuss in that section focus on the choice between centralization and delegation, and they abstract from some factors that complicate delegation in real organizations. The models, for instance, assume that the owner is able to delegate decision rights formally and irrevocably to her employee. In practice, however, legal barriers make full delegation in organizations impossible. The models also abstract from actions the owner can take to control the employee, such as constraining the decisions he can make or aligning his incentives through incentive contracts. We discuss these additional issues in Sections 5.2 and 5.3.

Throughout, we focus on the economics literature on delegation in firms. We therefore do not discuss the related political science literature on the delegation of public policymaking from elected politicians to bureaucrats (see, e.g., Epstein and O'Halloran 1999; Huber and Shipan 2006) and from legislatures to standing committees (e.g., Gilligan and Krehbiel 1987, 1989; Krishna and Morgan 2001).

5.1. Why Delegate?

In this section we discuss some of the main reasons for delegation that have been examined in the literature.

5.1.1. To Economize on Limited Resources

Probably the most common reasons for assigning a task to another person that one could do oneself are to add capacity to carry out the task and to tap into comparative advantage. In the

context of delegated decisionmaking, both factors can be at work. Gathering information and analyzing it to support good decisions is a time-consuming activity. Sharing the load allows decisions to be made more quickly. This in turn may be desirable because “time is money” or because, in a dynamic context, decisions are made on the basis of more up-to-date information. At the same time, there may be efficiency gains from using someone who is especially efficient in developing information, analyzing it, and making the needed choices. This efficiency could come from experience or ability. In addition, as Jensen and Meckling (1992) accentuated, delegation of decisions might be driven by the difficulties of communication that limit the possibility of informing the original decisionmaker, a phenomenon they referred to as “specific information.” We discuss this motivation for delegation in Section 5.1.3.

The most prominent model of delegation of decision rights to increase capacity is Geanakoplos and Milgrom (1991). They consider the problem of allocating a target output among n production facilities whose costs are random. The amount of information a manager can obtain about these costs is an increasing function of the time he spends and his ability. Geanakoplos and Milgrom show that a hierarchical structure, in which targets tumble down to the next level, where new targets for the succeeding level are calculated, is an efficient approach to the problem. However, the model is formulated as a team problem, so there are none of the incentive issues that are the focus of this chapter. Moreover, this model is analyzed in detail in Garicano and Van Zandt (this volume). Thus we will not develop it further here.

The models of organization based on distributed computing of an associative operation that begin with Mount and Reiter (1990) can also be thought of as delegation to increase the speed of arriving at a solution. And although the interpretation of the knowledge-based hierarchy in Garicano (2000) as delegation seems a stretch, his model does feature the use of multiple agents to add processing power. Again, incentive issues are ignored in these literatures, so we simply refer the reader to the extensive discussions in Garicano and Van Zandt (this volume).

One literature in which incentive issues are central that can be interpreted as involving delegation to increase capacity (although it is not usually formulated this way) concerns using an agent to select a risky investment. This literature begins with Lambert (1986), who posed the problem of motivating a risk-averse agent to gather information about the returns to a risky investment and then choose between this investment and a safe alternative. To think of this as a delegation problem, we could imagine that the principal could make the choice herself, and perhaps even develop the information, but assume that the agent is more efficient at doing so. Lambert (1986) noted an interesting nonmonotonicity: increasing the sensitivity of the agent's pay to the performance of the investment does not always increase his effort, because at some point he bears too much risk and switches to choosing the safe investment. Even under the optimal contract, the agent does not always choose the investment the principal would prefer. For details of this model, see Gibbons and Roberts (“Economic Theories of Incentives in Organizations,” this volume).

A second literature that can be interpreted in the same way as involving delegation to a more efficient agent concerns delegated portfolio management. The models here, beginning with Bhattacharya and Pfleiderer (1985) and Admati and Pfleiderer (1997), have assumed that the manager to whom selection of a portfolio of safe and risky assets is delegated knows *ex ante* if he is good or bad at his task. In addition, there is a moral hazard problem in inducing him to invest in information about the risky assets. Clean, positive results in these models are difficult to obtain (see Stracca 2006 for a survey).

A third context involves "referrals" (Garicano and Santos 2004), where a party must decide whether to delegate a task to another party who may be better suited to doing it. Referring the task means, however, that the referring party will not get the rents that attach to carrying it out. Garicano and Santos (2004) analyze the complex incentives that arise when the original party is privately informed about the returns to the task and effort exerted in carrying it out is unobservable. They suggest that ex ante sharing arrangements (partnerships) may arise.

Finally, Athey et al. (1994) investigate a model of management by exception, where managers delegate decisions to lower level employees to preserve managerial time, even if the managers might make better decisions than the employees. Which level is better equipped to make a decision depends on the state of the world. As the probability distribution over states shifts to increase the likelihood of decisions where the manager is the better decisionmaker, the range of decisions delegated to the employees increases to free up managerial bandwidth.

5.1.2. To Pay the Employee Less

Many employees value the right to make decisions. Employees may, for instance, believe that being in charge allows them to make decisions that are better for them personally. Or they may believe that being in charge is good for their careers, for instance, because it allows them to signal their abilities or develop their human capital. And, of course, they may simply value the right to make decisions per se, even if they end up making the same decisions that their superiors would have made. Whatever their reasons, if employees value the right to make decisions, firms may be able to charge them for this right by cutting wages. Delegation is then optimal if the reduction in wages is larger than any adverse effect that delegated decisionmaking may have on gross profits. This basic motivation for delegation has been explored in several papers, including (the first part of) Aghion and Tirole (1997) and Bester (2009). Our discussion below is based on the first part of Bester (2009).

To explore this motivation for delegation, recall the cheap-talk model discussed in Section 2.2, but suppose now that both parties can observe the state. Also, payoffs are now given by

$$U_1(d, w) = K - k_1(d - s - b)^2 + w \quad (1)$$

and

$$U_2(d, w) = K - k_2(d - s)^2 - w, \quad (2)$$

where w is the wage, K is a constant, and k_1 and k_2 are parameters that measure the importance of the decisions to the employee and the owner, respectively. We assume throughout that K is sufficiently large for joint surplus to be positive. Moreover, we now assume that the state is drawn from a uniform distribution with support $[-L, L]$.

If the owner has the right to make the decision, she will always make her preferred decision $d^C = s$, where C stands for centralization. To ensure participation, she then has to offer the employee a wage $w^C = -(K - k_1 b^2)$. If, instead, the owner delegates the decision right to the employee, the employee will always make his preferred decision $d^D = s + b$, where D stands for delegation. The owner then has to offer the employee only $w^D = -K$ to ensure participation. On one hand, delegation therefore reduces gross profits by $k_2 b^2$. On the other hand, it also allows the owner to cut wages by $k_1 b^2$.

In the case of quadratic preferences, delegation is therefore optimal if the employee's payoff is more sensitive to the decision than is the owner's, that is, if $k_1 \geq k_2$. More generally, delegation is optimal if the decision is more important to the employee than to the owner, in the sense that his willingness to pay for the right to make the decision $E[U_1(d^D, w) - U_1(d^C, w)]$ is larger than the owner's willingness to pay $E[U_2(d^C, w) - U_2(d^D, w)]$. Notice that this is equivalent to saying that delegation is optimal if the expected surplus $S(d) \equiv E[U_1(d, w) + U_2(d, w)]$ is larger when the employee makes the decision than when the owner makes it herself. In the next few sections we explore different reasons this may be the case.

Before we do so, it is worth making two observations. First, our focus on the efficient allocation of decision rights is different from some papers in the literature that instead focus on the allocation that maximizes expected profits gross of wages. This difference does not affect the main insights from these papers, but it does affect some of the comparative statics.

Second, the prediction that organizations allocate decision rights to whomever values them the most is not uncontroversial. Rotemberg (1993), in particular, argues that firms instead allocate decision rights to those employees they most want to retain. To explain why this may be so, he points to a difference between paying with cash and paying with control that we abstracted from above. In particular, he observes that some benefits of being allowed to make a decision today accrue to an employee only if he is still with the firm tomorrow. Rotemberg therefore views control as a form of deferred compensation. He then shows that if the manager cannot commit to future wages, and the determination of future wages is hindered by asymmetric information, the owner may find it optimal to give decision rights to the employees she most wants to retain rather than to the ones who value decision rights the most.

5.1.3. To Make Better Use of the Employee's Information

A common argument for delegation is that it allows the owner to make better use of her employee's information. Two early papers that explore this argument are Holmström (1977, 1984) and Jensen and Meckling (1992). Both papers focus on the trade-off between the benefit of making better use of the employee's information and the costs that arise because the employee's decisionmaking is biased toward decisions that are good for him but not necessarily good for the organization as a whole. The papers differ, however, in their explanation for why the owner cannot retain the decision right and simply have the employee communicate his information to her. In particular, whereas Jensen and Meckling (1992) focus on communication costs that arise because of bounded rationality, in Holmström (1977, 1984) communication is costly because of the employee's incentives to use his information strategically.

Below we start by discussing Dessein (2002), who provides a more recent exploration of these issues. Similar to Holmström (1977, 1984), Dessein focuses on communication costs that arise because of the employee's incentives to use his information strategically. In contrast to both Holmström (1977, 1984) and Jensen and Meckling (1992), however, Dessein does not allow the owner to constrain the employee's decision or align his preferences through incentive contracts. Dessein's focus on the choice between centralization and unconstrained delegation is in line with the other models that we discuss in this section. As mentioned above, we discuss models that allow the owner to control the employee by constraining his decisions or by using incentive contracts in Section 5.3.

A feature of Dessein (2002) is that there is a single decision, and all information that is relevant for this decision is concentrated in a single employee. We first discuss this case. We then turn to an extension of Dessein (2002) in which information is vertically dispersed, in the sense that the employee knows more about some factors that are relevant for the decision and the owner knows more about others. Finally, we discuss recent work by Alonso et al. (2008) and Rantakari (2008), in which there are multiple decisions that have to be coordinated and relevant information is dispersed horizontally among multiple employees.

5.1.3.1. Concentrated Information. For an illustration of the main arguments in Dessein (2002), recall the model we described in the previous section, but suppose now that the owner cannot observe the state of the world s . Suppose also that the decision is more important to the owner than to the employee, that is, $k_2 > k_1$. It follows from our discussion in the previous section that if this were not the case, delegation would always be optimal.

If the owner delegates the decision right to the employee, the employee will make his preferred decision, just as in the previous section. The difference between first-best expected surplus and expected surplus under delegation is then given by

$$S(d^{FB}) - S(d^D) = \frac{k_2^2}{k_1 + k_2} b^2, \tag{3}$$

where the right-hand side is the inefficiency that arises because the employee's preferences are different from first best.

If the owner does not delegate the decision right to the employee, the employee sends the owner a cheap-talk message, after which the owner makes her preferred decision. Under centralization, the model is therefore the same as the cheap-talk model we examined in Section 2.2, except that the owner now has to pay the employee a wage to ensure participation. In Section 2.2 we saw that if $b > 0$, communication is imperfect. Moreover, the quality of communication is measured by the residual variance

$$V(b) = \frac{\text{Var}(s)}{N(b)^2} + \frac{N(b)^2 - 1}{3} b^2, \tag{4}$$

where $\text{Var}(s)$ is the variance of s , and $N(b)$ is the maximum number of partitions given by the smallest integer larger than or equal to

$$-\frac{1}{2} + \frac{1}{2} \sqrt{1 + \frac{4L}{b}}, \tag{5}$$

where $2L$ is the length of the support. The difference between first-best expected surplus and expected surplus under centralization is then given by

$$S(d^{FB}) - S(d^C) = \frac{k_1^2}{k_1 + k_2} b^2 + (k_1 + k_2) V(b), \tag{6}$$

where the first term on the right-hand side is the distortion that arises because the owner's preferences are different from first best, and the second term is the distortion that arises because the owner is imperfectly informed about the state.

It follows from (3) and (6) that the owner will delegate to the employee if and only if

$$(k_2 - k_1) b^2 \leq (k_1 + k_2) V(b). \tag{7}$$

One implication of this expression is that delegation can be optimal even though the decision is more important to the owner than to the employee. The reason, of course, is that delegation now has the additional benefit of allowing the owner to make more efficient use of the employee's information. Another implication of the above expression is that the optimal allocation of the decision right depends on the quality of communication between the employee and the owner. To understand the effect of changes in the environment on the efficient allocation of decision rights, one therefore has to understand how such changes affect the flow of information in organizations.

In particular, consider a reduction in the divergence parameter b . On one hand, such a reduction leads to more efficient decisionmaking under delegation. On the other hand, it also improves communication under centralization. A reduction in b therefore favors delegation only if it has a more beneficial effect on decisionmaking under delegation than on communication under centralization. To see that this is indeed the case, consider (4), (5), and (7). It follows from these expressions that delegation is optimal if and only if

$$(k_2 - k_1) b^2 \leq (k_1 + k_2) \text{Var}(s), \tag{8}$$

where we are using the fact that delegation is always optimal when informative communication is feasible. Delegation is therefore optimal if the divergence in preferences is small relative to both the employee's informational advantage and the importance of the decision to the owner. Moreover, more congruent preferences can lead to a switch from centralization to delegation but not the reverse.

5.1.3.2. Vertically Dispersed Information. A key implication of Dessein (2002) is that the efficient allocation of decision rights depends on the endogenous flow of information. Extending the model by allowing for dispersed information further highlights the importance of taking into account the endogenous flow of information. Our discussion of this case is based on Harris and Raviv (2005).

For an illustration of their main arguments, suppose that there is another state \hat{s} , which is independently drawn from a uniform distribution with support $[-\hat{L}, \hat{L}]$ and variance $\text{Var}(\hat{s}) = \hat{L}^2/3$. The state \hat{s} is observed by the owner but not by the employee. Under either organizational structure, the party without the decision right sends a single cheap-talk message to the decisionmaker, after which the decisionmaker makes his or her preferred decision. Finally, the owner's and the employee's payoffs are now given by

$$U_1(d, w) = K - k_1(d - s - \hat{s} - b)^2 + w \tag{9}$$

and

$$U_2(d, w) = K - k_2(d - s - \hat{s})^2 - w. \tag{10}$$

Proceeding as we did above, one can show that delegation is now optimal if and only if

$$(k_2 - k_1) b^2 \leq (k_1 + k_2)(V(b) - \hat{V}(b)), \quad (11)$$

where the residual variance $\hat{V}(b)$ characterizes the quality of communication under delegation and can be obtained by replacing $\text{Var}(s)$ with $\text{Var}(\hat{s})$ in (4) and L with \hat{L} in (5). As in the model with concentrated information, a reduction in the divergence parameter improves both decisionmaking under delegation and communication under centralization. In addition, however, a reduction in b now improves communication under delegation. As a result, such a reduction can actually lead to a switch from delegation to centralization. Although this result contradicts the prediction of the model with concentrated information, it underlines that model's central insight that the efficient allocation of decision rights depends on the endogenous flow of information.

5.1.3.3. Horizontally Dispersed Information. So far we have focused on a setting with a single decision. Additional issues arise when there are multiple decisions that have to be coordinated and the information that is relevant for the decisions is dispersed among multiple employees. Coordinated decisionmaking then requires aggregation of the dispersed information. Some recent papers explore the allocation of decision rights in such a setting (Alonso et al. 2008; Rantakari 2008; Dessein et al. 2010; Friebl and Raith 2010). In line with the models we just discussed, these papers focus on communication costs that arise because of the employees' incentives to use their information strategically, which is why we concentrate on this branch of the literature below. For related models in which communication is not strategic, see, in particular, Aoki (1986), Hart and Moore (2005), Dessein and Santos (2006), and Crémer et al. (2007).

To illustrate the main insights in Alonso et al. (2008), suppose now that there are two employees. We think of the owner as the CEO of a multidivisional firm and the employees as the division managers.

Profits depend on two decisions, one of which is associated with Division 1 and the other with Division 2. For each decision, the firm faces a trade-off between coordination and adaptation. On one hand, each decision should be adapted to the business environment faced by the relevant division; on the other hand, decisions should also be coordinated with each other. Suppose, for instance, that the firm is a car manufacturer and that each division serves a particular country. Moreover, the firm has to decide what car it should produce for, and sell in, each country. The firm then faces a trade-off between tailoring the car types to the local tastes in each market—and thus increasing revenue—and coordinating the car types across markets—and thus reducing costs. This trade-off is captured in the profit functions

$$\pi_1 = K - (d_1 - s_1)^2 - \delta (d_1 - d_2)^2 \quad (12)$$

and

$$\pi_2 = K - (d_2 - s_2)^2 - \delta (d_1 - d_2)^2, \quad (13)$$

where K is a constant and $\delta \geq 0$ is the parameter that measures the relative importance of adapting each decision to the state and setting them equal to each other. Consistent with the

notation in the previous sections, d_1 and d_2 are the two decisions, and s_1 and s_2 are the states that summarize the business environment faced by each division.

All managers are risk neutral and thus care about expected profits. They differ, however, in how much weight they put on the expected profits of the different divisions. In particular, although the CEO cares about the firm's overall profits, each division manager puts weight λ on the profits of his own division and only $1 - \lambda$ on that of the other division. The parameter $\lambda \in (1/2, 1]$ therefore measures each division manager's own-division bias.

In line with the models in Section 5.1, we are interested in comparing the efficiency of two organizational structures, centralization (in which case both decisions are made by the CEO) and decentralization (in which case each division manager makes his decision without knowing what decision the other division manager has made). Under centralization the division managers communicate vertically with the CEO, and under decentralization they communicate horizontally with each other. In either case, communication takes the form of a single cheap-talk message that each division manager sends to either the CEO (in the case of centralization) or to his counterpart (in the case of decentralization). These messages are always sent simultaneously.

The endogenous quality of communication is crucial for the main result of the model. To see this, suppose first that the quality of communication is exogenously given. This would be the case, for instance, if communication were imperfect because of physical constraints rather than because of the division managers' strategic incentives to manipulate their information. And to keep things simple, suppose also that the exogenously given quality of communication is the same across organizational structures.

In such a setting, it is not immediately obvious whether the firm should centralize or decentralize. On one hand, the division managers know more about the states of the world than the CEO; on the other, their incentives to make the right decisions are distorted. If the quality of communication were perfect, centralization would therefore be optimal. If communication were entirely uninformative, however, decentralization would always be optimal. And in between these two extremes, the choice between the two organizational structures is determined by the need for coordination δ : whatever value λ takes, centralization is always optimal provided that the need for coordination is sufficiently strong. This conclusion is also illustrated in Figure 2. In the limit, however, in which δ goes to ∞ , expected profits under the two organizational structures converge to the same level. The reason is that as coordination becomes all important, the division managers understand that setting decisions equal to each other is all important, and they act accordingly. In spite of this convergence, however, it is still the case that for any large but finite δ , centralization is optimal. And this result, of course, is in line with the received wisdom that if coordination is important, centralization outperforms decentralization.

The main insight of the model is that this intuitive result depends crucially on the assumption that communication is exogenous. To see this, we need to understand the division managers' incentives to misrepresent their information under the two organizational structures. Consider, then, the incentives of one of the division managers, say, Manager 1, the division manager in charge of Division 1. From his perspective, when the CEO makes decision d_1 , she puts too little weight on adapting it to s_1 and too much on coordinating it with d_2 . Thus if Manager 1 would truthfully communicate his state, he would expect the CEO's choice of d_1 to be too close to zero. To induce the CEO to make a more extreme decision, Manager 1 therefore exaggerates his state and reports that its absolute value is larger than it actually is. Crucially, the more important is coordination, the more weight the CEO will put on setting d_1 equal to d_2 and

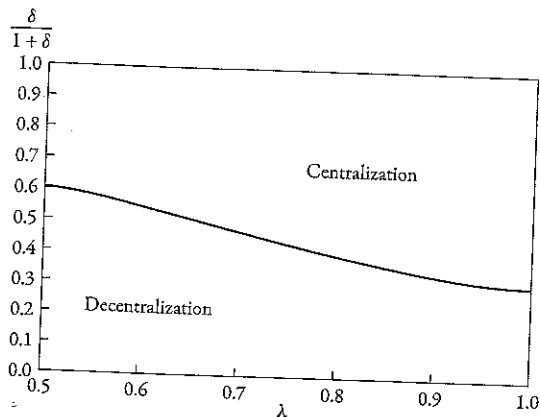


Figure 2. Exogenous communication.

the less weight she will put on setting d_1 equal to s_1 . Under centralization, therefore, the more important is coordination, the more the division managers will exaggerate their states, and the less information will be communicated in equilibrium.

This comparative static is the exact opposite under decentralization. To see this, consider Manager 1's incentives to misrepresent his information under decentralization. From his perspective, Manager 2 puts too little weight on coordinating d_2 with d_1 and too much weight on adapting d_2 to s_2 . Thus if Manager 1 truthfully communicated his state, he would expect Manager 2's choice of d_2 to be too close to zero. To induce Manager 2 to make a more extreme decision, Manager 1 will therefore once again exaggerate his state. In contrast to the centralized structure, however, Manager 1's incentive to exaggerate his state is weaker, the more important coordination is. Essentially, the more important is coordination, the more weight Manager 2 will put on coordinating d_2 with d_1 and the less weight he will put on adapting d_2 to s_2 . Under decentralization, therefore, the more important coordination, the less the division managers will exaggerate their states, and thus the more information will be communicated in equilibrium.

This difference between horizontal and vertical communication has important implications for the choice between the two organizational structures. In particular, with endogenous communication it is no longer the case that centralization is optimal when coordination is sufficiently important. This is illustrated in Figure 3. As with exogenous communication, decentralization is optimal if coordination is sufficiently unimportant. If λ is small, however, decentralization now remains optimal even if coordination becomes arbitrarily important. This model therefore once again highlights the importance of taking into account the endogenous nature of communication when predicting how changes in the environment affect the optimal allocation of decision rights.

5.1.4. To Motivate the Employee

Another common argument for delegation is that it motivates employees to provide more effort (Aghion and Tirole 1997). In line with our discussion in Section 3.5, it is useful to

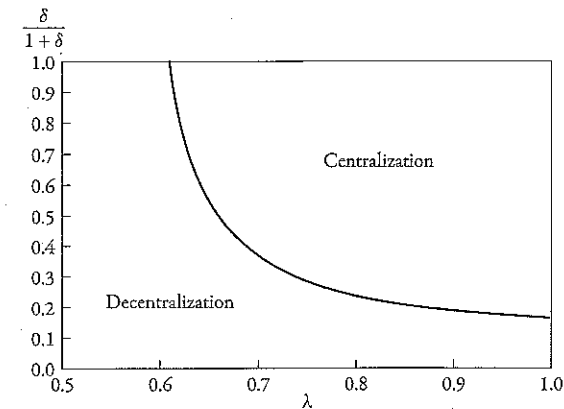


Figure 3. Endogenous communication.

distinguish between motivating the employee to put more effort into information acquisition and motivating him to put more effort into execution.

5.1.4.1. Motivating Information Acquisition. Delegation not only allows the owner to make more efficient use of the employee's information, as we saw in the previous section. It may also allow the employee to use his information more effectively in pursuing his own goals. One would therefore expect delegation to increase the value of information to the employee and thus increase his incentives to become informed. In this section we first explore this motivational effect of delegation in the context of Aghion and Tirole (1997). We then return to Che and Kartik (2009), who show that in some cases delegation may actually reduce the value of information to the employee and thus weaken his incentives to become informed.

To explore the motivational effect of delegation, we make three changes to the version of Aghion and Tirole (1997) described in Section 3.5. First, at the beginning of the game, the owner now offers the employee a contract that specifies a wage and an allocation of the decision right. Second, we assume that the decision is more important to the owner than to the employee, that is, $B > b$. If this were not the case, delegation would be optimal even if both parties were perfectly informed. Third, to focus on the employee's incentives, we assume for now that the probability E with which the owner becomes informed is fixed. This would be the case, for instance, if her expertise were based on past experience rather than on current efforts to acquire information.

In Section 3.5 we saw that if the owner retains the right to make the decision, the employee benefits from becoming informed only if the owner is uninformed. Specifically, the employee's marginal benefit of getting informed is given by $(1 - E)b$.

If, instead, the owner delegates the decision right to the employee, the employee benefits from becoming informed even if the owner is informed too, because he can then insist on his preferred decision being made rather than having to accept the owner's. As a result, under delegation,

the employee's marginal benefit of becoming informed is given by $(1 - E)b_E + E(1 - \alpha)b = b(1 - \alpha E)$.

Delegation therefore increases the employee's marginal benefit of becoming informed, which in turn induces him to provide more effort. Delegation is then optimal if this motivational effect is sufficiently large relative to the cost of having decisions made by someone whose preferences are further from first best. In particular, delegation is optimal if

$$S^D - S^C = \left[(1 - E)(e^D - e^C)(\alpha B + b) - (c(d^D) - c(d^C)) \right] - Ee^D(1 - \alpha)(B - b) > 0, \quad (14)$$

where the first term on the right-hand side captures the benefit of a more motivated employee, and the second term captures the loss of control.

So far we have taken the owner's effort as exogenous. If effort is endogenous, delegation discourages information acquisition by the owner, for the same reason that it encourages information acquisition by the employee. As in Grossmann and Hart (1986) and Hart and Moore (1990), a change in residual control rights therefore leads to changes in the equilibrium levels of effort or investment.

The observation that delegation may involve a trade-off between losing control and motivating the employee is a central insight of Aghion and Tirole (1997). Although this trade-off is intuitive, some papers have explored conditions under which delegation may actually discourage the employee. The model in Rantakari (2012), for instance, can be viewed as a multitasking version of Aghion and Tirole (1997) in which the employee can choose both how much effort he puts into information acquisition and the type of decisions that he investigates. In such a setting, delegation motivates the employee to put effort into finding decisions that are good for him personally, but it discourages him from finding decisions that are good for the owner.

Another model in which delegation can discourage the employee is Che and Kartik (2009). They consider a model in which the owner and the employee have the same preferences over decisions but different priors about the state of the world and thus about what decision to make. One motive for the employee to acquire information is then to persuade the owner of his own opinion (see also Van den Steen 2002). Under delegation, the employee no longer needs to persuade the owner and may thus be less motivated to acquire information. This complacency effect of delegation can arise even if the employee and the owner have the same priors. To see this, notice that in Aghion and Tirole (1997) the allocation of the decision right does not matter if both parties are uninformed, because they then agree to make the default decision. Suppose in contrast that there are two default decisions, one that gives the owner zero but costs the employee $c > 0$ and another that gives the employee zero and costs the owner c . Under centralization, the employee's marginal benefit of becoming informed is then given by $(1 - E)(b + c)$, whereas under delegation it is still given by $b(1 - \alpha E)$. Delegation then discourages the employee, provided that c is sufficiently large. In the terminology of Che and Kartik (2009), delegation makes the employee "complacent," because it makes it less costly for him to be uninformed.

5.1.4.2. Motivating Execution. Delegation might motivate the employee to put more effort not only into information acquisition but also into implementing and executing decisions. The employee may, for instance, believe that his preferred decision has a higher chance of

succeeding than that of the owner (Zábojník 2002; Van den Steen 2006). Or he may simply receive a larger private benefit from successfully implementing his preferred decision (Bester and Krämer 2008).

Delegation can motivate execution; however, the need to motivate execution may actually make delegation less efficient relative to centralization (Bester and Krämer 2008). The reason is that if the owner has to motivate execution, she needs to take the employee's preferences into account when she makes a decision. The need to motivate execution therefore acts as a disciplinary device on the owner that makes centralized decisionmaking more efficient relative to delegated decisionmaking.

To see this more clearly, consider a simplified version of the model in Bester and Krämer (2008). In particular, recall the model in Section 5.1.2, but suppose now that after a decision has been made, the employee has to choose his implementation effort $e \in \{0, 1\}$. The decision succeeds if and only if the employee provides effort and success is not contractible. The parties' payoffs are then given by

$$U_1(d, e) = e \left[K - k_1(d - \theta - b)^2 \right] - ce + w \quad (15)$$

and

$$U_2(d, e) = e \left[K - k_2(d - \theta)^2 \right] - w, \quad (16)$$

where c is the employee's cost of effort. Suppose also that effort costs satisfy

$$K - k_1b^2 < c < K - k_1E \left[(d^{FB} - \theta - b)^2 \right], \quad (17)$$

where d^{FB} is the first-best decision if the employee does provide effort. Finally, suppose that $k_1 > k_2$, so that the decision is more important to the employee than to the owner.

If effort is contractible, motivating execution is not a concern. As in the model in Section 5.1.2, it is then efficient to give the decision right to the party that values this right more highly, which in this case is the employee. If effort is not contractible, however, motivating execution does become a concern. Under centralization, the owner then cannot simply make her preferred decisions, because if she did, the employee would not put any effort into execution. To motivate execution, the owner therefore has to bias her decision toward the employee's preferred decision, which makes her decision more efficient. In contrast, if the owner delegates, the employee will still make his preferred decision. The need to motivate execution therefore makes centralized decisionmaking more efficient relative to delegation. As a result, the owner may find it optimal to centralize when she has to motivate execution, even though she would delegate if motivation were not a concern.

5.1.5. To Learn More about the Employee

Delegation can also be a tool to test the employee and to learn more about his ability or preferences. For instance, a CEO may delegate control over a division to a senior executive to test the executive's ability to run his own business. Or she may delegate to the executive to find out more about his preferences, such as whether he is an empire builder or is focused on generating

profits. Aghion et al. (2002, 2004) explore this second possibility. In doing so, they also challenge the assumption in most of the literature that the allocation of decision rights is deterministic. In this section we first discuss the use of delegation as a means to learn more about the employee's preferences and then discuss contingent allocations of decision rights.

5.1.5.1. Learning about the Employee's Preferences. To see some of the insights in Aghion et al. (2002, 2004), consider the following two-period model. In the first period an organization has to decide on the type of project it wants to design, and in the second it has to decide on whether to implement the project. In the first period Aghion et al. make three changes to the model in Aghion and Tirole (1997), which we described in the previous section. First, wages are fixed at zero. Second, the owner and the employee both know the payoffs of all the projects, that is, $E = e = 1$. Third, the employee is privately informed about his type, either bad (which happens with probability $\mu < \frac{1}{2}$) or good (which happens with probability $1 - \mu$). If the employee is a good type, the congruence parameter α is equal to one, and if he is a bad type, it is equal to zero. For reasons that will become apparent, we assume that $b > 1$.

In the second period, the owner decides whether to implement the project that was designed in the first. If the owner does not implement the project, both parties receive a zero payoff, regardless of whether the employee is a good or a bad type. But if the owner does implement the project, she gets +1 if the employee is a good type and -1 if he is a bad type, while the employee obtains +1 in either case. Implementing the project is therefore in the owner's interest only if the employee is a good type.

In this setting, if the owner does not delegate the design decision, she will choose her preferred design in the first period and then implement it in the second period. In this case, the owner's expected payoff is therefore given by $E[U_2^C] = B + 1 - 2\mu$. If, instead, the owner does delegate the design decision, the employee will always choose his preferred design and thus reveal his type. In the second period, the owner then implements the design if and only if the employee is a good type. In this case, the owner's expected payoff is therefore given by $E[U_2^D] = (1 - \mu)(B + 1)$.

The gain from delegation is then given by $E[U_2^D - U_2^C] = \mu - \mu B$, which is positive if $B < 1$. In this setting, delegation is therefore a tool that allows the owner to learn more about the employee's preferences. For the owner to make use of this tool, however, her future benefit of knowing the employee's preferences must outweigh current costs of learning them, that is, $B < 1$. At the same time, the employee's current benefit of using authority in his favor must outweigh the future costs of a ruined reputation, that is, $b > 1$.

5.1.5.2. Contingent Allocations of Decision Rights. The models we have examined so far assume that, although decision rights are contractible, their allocation cannot be made contingent on any information that is revealed after the initial contract is signed. This assumption is typically justified informally by arguing that the reallocation of decision rights in organizations takes too long relative to the time frame in which most decisions have to be made. In practice, reallocations certainly appear to be rare relative to the frequency with which decisions are made.

In the confines of the models that we have examined so far, however, there is no reason the owner should not be able to commit to contingent allocations of the decision rights. And as

shown in Aghion et al. (2002, 2004) and Krämer (2006), doing so can indeed make the owner better off.

To see why contingent allocations can make the owner better off, consider the setting we sketched above, and suppose that the allocation of the first-period decision right is made contingent on a message sent by the employee. In particular, if the employee admits to being a bad type, he is rewarded by receiving the decision right with probability $1/b$; if, however, he claims to be a good type, the owner retains the decision right for sure. The probability $1/b$ with which the decision right is delegated to the employee ensures that a bad employee is just indifferent between revealing his type (in which case he expects to make $(1/b)b$ today and zero tomorrow) and claiming to be a good type (in which case he makes zero today and +1 tomorrow). In this setting, contingent delegation therefore allows the owner to learn the employee's type at a lower cost and thus makes her better off.

If contingent delegation can make the owner better off, why would she ever delegate to the employee unconditionally? One reason is that the contingent allocation requires the owner to delegate to the employee when she has learned that the employee is a bad type. One can imagine situations in which the owner is unable to commit herself to such an ex post inefficient action, even if it is ex ante efficient to do so.

Notice also that in a slightly more general setting than the one we examined above, the optimal revelation mechanism may call for the unconditional transfer of the decision right to the employee. In particular, suppose that the good employee does not have the same preferences as the owner. Aghion et al. (2002, 2004) show that if the owner's preferences are sufficiently different from those of both types of employees, it is no longer possible to induce the employee to reveal his type. The owner can then no longer be made better off by making delegation contingent on the employee's messages.

Finally, as we discuss in Section 5.2, legal barriers prevent organizations from committing to contracts that transfer decision rights from the top to others in the organization. Below we argue that organizations may still be able to implement deterministic allocations through other means that do not require formal contracts. It is unclear, however, whether these other mechanisms would also allow for the implementation of contingent allocations of decision rights. For these reasons, in the remainder of this chapter we return to the assumption that the allocation of decision rights is deterministic.

5.1.6. Anything Else?

Alfred Sloan, the longtime President and CEO of General Motors, is widely regarded as one of the inventors of the multidivisional organization that is now common in large corporations. A major theme in his autobiography is the importance for corporate performance of properly allocating decision rights between headquarters and the divisions. As he puts it in the last chapter of his autobiography:

It has been a thesis of this book that good management rests on the reconciliation of centralization and decentralization, or "decentralization with coordinated control." . . . From decentralization we get initiative, responsibility, development of personnel, decisions close to the facts, flexibility—in short, all the qualities necessary for an organization to adapt to new conditions. From coordination we get efficiencies and economies. It must be apparent

that co-ordinated decentralization is not an easy concept to apply. There is no hard and fast rule for sorting out the various responsibilities and the best way to assign them. The balance which is struck between corporate and divisional responsibility varies according to what is being decided, the circumstances of the time, past experience, and the temperaments and skills of the executives involved. [Sloan 1964: 429]

Much of what Sloan says in this quotation is reflected in the literature on delegation that we have discussed so far. Just like Sloan, for instance, the literature has focused on "initiative," "flexibility," and "decisions close to the facts" as primary benefits of a decentralized structure. And the trade-off between adaptation and coordination that Sloan describes has been at the center of much recent work. The quotation, however, also points to some issues that we have not yet addressed.

For instance, Sloan is not the only practitioner who believes in the importance of delegation for the "development of personnel." Nevertheless, little, if any, work has been done on the role of delegation in the development of managerial human capital. More generally, the literature has so far focused largely on static models of delegation. The dynamics of delegation is an important and potentially fruitful area of future research.

The models that we have discussed so far are also quite stark in that they force the owner to choose between centralization and unconstrained delegation and do not allow her to control the employee through different means. As such, these models do not capture Sloan's policy of "decentralization with coordinated control."

In the rest of this chapter we explore some of these additional issues. In particular, in the next section we discuss the issues that arise if the owner cannot commit to contracts that transfer decision rights to the employee. This is one area of the literature in which dynamic models of delegation do play a prominent role. In Section 5.3 we then discuss models that allow the owner to better control the employee by either constraining the decisions he can make or by using incentive contracts to align his preferences.

5.2. How to Delegate

The models we have discussed so far ask *why* the owner may find it optimal to delegate decision rights, but they do not examine *how* the owner may do so. This is potentially problematic, because, by law, formal authority resides at the top of organizations and the "business judgment rule" makes it impossible to write legal contracts that transfer formal authority to other parties in the organization (see Kornhauser and MacLeod, this volume).

One way for the owner to transfer formal authority to an employee is to sell him the assets that are necessary to make the decision and thus turn him into an independent subcontractor (Grossman and Hart 1986; Hart and Moore 1990). Given our focus on internal organization, however, we are interested in how the owner can delegate formal authority to an employee without changing the boundaries of the organization.

In this section we review some papers that show that even if the owner cannot write legal contracts delegating formal authority to another party in her organization, she may be able to commit to behaving as if formal authority had been delegated. We first discuss how the owner may be able to do so by committing to an information structure in which it is optimal for her

to rubber-stamp any recommendation that the employee might make. We then discuss models in which reputational concerns induce the owner to rubber-stamp the employee's recommendations, even if it is not in her short-term interest to do so.

5.2.1. Delegation through Ignorance

The old saying that knowledge is power suggests that information can lead to the separation of real and formal authority, which is a central observation in Aghion and Tirole (1997). We saw above that in their model an uninformed owner finds it optimal to rubber-stamp her employee's recommendation, even if she knows that the employee is selfishly recommending his own preferred decision. This suggests that an owner who cannot commit to delegation contractually may still be able to do so by committing to an appropriate information structure.

For an owner to delegate by committing to an appropriate information structure, two conditions have to hold. First, it must be the case that even if the owner knows that the employee is simply recommending his preferred decision, she cannot invert the recommendation and thus infer his information. Second, there must be an underlying congruence in the two parties' preferences: the owner must prefer the employee's preferred decision to her either doing nothing or making a decision at random.

Both conditions are satisfied in Aghion and Tirole (1997). First, if the employee recommends a decision, the owner cannot infer the payoffs associated with the other decisions. Second, because the owner expects a positive payoff from the employee's preferred decision, she prefers that decision to her either doing nothing, in which case she gets zero, or making a decision at random, in which case she expects to be wiped out. In this setting, the owner can therefore transfer real authority to the employee by committing to not becoming informed. Aghion and Tirole (1997) suggest that the owner may be able to do so, for instance, by becoming overloaded with other decisions.

In Aghion and Tirole (1997) the employee's recommendation is entirely noninvertible. Callander (2008b) instead develops a model in which an employee's recommendations are partially invertible. To see a simple version of his model, suppose that an owner and an employee have quadratic preferences over an outcome $x \in \mathcal{X}$, with the employee preferring a larger outcome than does the owner. An outcome function maps the decision d into the outcome. The crucial feature of the model is that the outcome function is generated by Brownian motion. Both parties know the drift and variance of the Brownian motion, and they know the outcome associated with a default decision. However, only the employee knows the full realization of the Brownian motion.

In this setting, if the employee recommends the decision that leads to his preferred outcome, and the owner is aware of this, the owner can infer some information about the outcomes associated with other decisions. Callander (2008b) shows that the owner finds it optimal to rubber-stamp the employee's selfish recommendations if the variance of the Brownian motion is above a threshold, in which case the owner is very uncertain about the outcomes associated with the other decision. If the variance is below the threshold, however, the owner may find it optimal to use the information that she can infer from the employee's recommendation to overrule the employee and make another decision. This model suggests that the owner's ability to delegate through ignorance depends on the complexity of the decision, where the complexity is captured by the variance of the Brownian motion.

5.2.2. Informal Delegation

We now turn to papers that explore the feasibility and optimality of informal delegation. The first section examines the conditions under which an owner is able to maintain a reputation for rubber-stamping her employee's recommendations, even if it is not in her short-term interest to do so. In the second section we then ask whether an owner who is able to maintain such a reputation finds it optimal to do so. The reason she may not is that if she can implement delegation, she may be able to implement other decision processes that do even better than either centralization or delegation.

5.2.2.1. Feasibility of Informal Delegation. We first examine the conditions under which an owner is able to resist the short-term temptation of overruling her employee and thus is able to engage in informal delegation. Baker et al. (1999) show that these conditions can be different from, and sometimes the opposite of, the conditions under which the owner would want to delegate to the employee in the first place.

To capture some of their insights, we make four changes to the version of Aghion and Tirole (1997) that we discussed in Section 5.1.4. First, we now assume that the owner is never informed about the decisions, that is, $E = 0$. Second, the employee's cost of effort is now given by $c(e) = ce^2$. Third, if the employee recommends a decision, the owner learns the payoffs from that decision. Information is therefore now assumed to be hard. Finally, we continue to assume that with probability α one of the decisions pays the owner B and the employee b , whereas the other pays each party zero. With probability $1 - \alpha$, however, one of the decisions now pays the owner B and the employee zero, whereas the other pays the owner and the employee $-B$ and b , respectively. In this case, the owner therefore prefers the default decision to the employee's preferred decision.

Before we turn to the infinitely repeated version of this game, consider a once-off interaction in which the owner is able to commit to delegation contractually. If the owner does delegate to the employee, the employee will always make his preferred decision. Effort and expected surplus will then be given by

$$e^D = \arg \max_e w + eb - ce^2 \quad (18)$$

and

$$S^D = e^D[\alpha(B + b) + (1 - \alpha)(-B + b)] - c(e^D)^2. \quad (19)$$

Suppose now that the owner does not delegate to the employee. If the employee becomes informed, he will then recommend his preferred decision if it coincides with the owner's preferred decision, but he will recommend the default decision if their preferred decisions do not coincide. Effort and expected surplus will then be given by

$$e^C = \arg \max_e w + eab - ce^2 \quad (20)$$

and

$$S^C = e^C\alpha(B + b) - c(e^C)^2. \quad (21)$$

We are interested in the case in which delegation is more efficient than centralization, but the owner is not able to commit to delegation contractually. For the rest of this section, we therefore

suppose that $S^D > S^C$, which will be the case if the congruence parameter α is sufficiently large.

Consider then an infinitely repeated version of the game in which a long-lived owner faces a series of short-lived employees, each of whom she interacts with only once.¹³ We focus on trigger strategies in which the owner promises to rubber-stamp any recommendation an employee might make and in which the parties revert to the best static equilibrium if the owner breaks her promise.

Standard arguments then show that the owner can implement delegation informally if and only if

$$S^D - S^C \geq rB, \quad (22)$$

where r is the discount rate. As one would expect, the owner is therefore able to replicate delegation informally, provided that she is sufficiently patient. More interestingly, the model predicts that the feasibility of informal delegation depends on the owner's renegeing temptation, which, in turn, depends on both the average and the extreme realizations of payoffs. Suppose, for instance, that we increase B but also adjust b and α in such a way that $S^D - S^C$ stays constant. Even though such a change does not affect the efficiency of delegation, it makes it more difficult to implement it, because the change increases the owner's renegeing temptation. Finally, notice that when preferences are highly congruent, delegation is efficient but cannot be implemented informally, because the owner's fallback option is then too attractive. This result is in stark contrast to models with contractible control, many of which predict that the owner will engage in delegation when preferences are highly congruent (e.g., see Dessein 2002; Alonso et al. 2008).

5.2.2.2. Optimality of Informal Delegation. We now turn to whether a long-lived owner who is able to replicate delegation informally finds it optimal to do so. Baker et al. (1999) and Alonso and Matouschek (2007) show that she may not, even if delegation is more efficient than centralization. As mentioned above, the reason is that if the owner is able to engage in informal delegation, she may also be able to commit to other, more efficient decision processes.

To understand the argument in Alonso and Matouschek (2007), consider an infinite repetition of the version of Dessein (2002) model that we described in Section 5.1.3. In line with the model in the previous section, suppose that the owner is long-lived and that she faces a series of short-lived employees. Delegation can then be implemented informally if and only if

$$S^D - S^C \geq rb^2, \quad (23)$$

where S^D and S^C are given by (18) and (20), and we are again assuming that if the owner reneges on her promise, the parties revert to centralization. Notice that the maximum renegeing temptation is given by b^2 , because when an employee truthfully reveals the state, the owner can always implement her preferred decision rather than that of the employee.

13. The main insights of the model would be unchanged if the long-lived owner interacted with a single, long-lived employee. The above formulation highlights that we are focusing on the owner's reputational concerns rather than those of the employees.

The key observation now is that if the owner is patient enough to engage in delegation, she is also patient enough to engage in threshold delegation. Under this scheme, the owner makes the employees' preferred decision if they report a state below the threshold $s^* = L - 2bk_2/(k_1 + k_2)$, and she implements $d = s^* + b$ otherwise. Threshold delegation involves the same maximum renegeing temptation as delegation, but it is more efficient.

If an owner is sufficiently patient to replicate delegation informally, she may therefore also be able to commit to decision processes that perform better than either centralization or delegation. On one hand, this provides a challenge to our analysis, because it requires us to consider many more decision processes than the two that we have focused on so far. On the other hand, it also brings the analysis closer to the real world, because firms clearly do use many different types of decision processes. In the next section we go beyond the simple example in this section and examine optimal decision processes in more general environments.

5.3. How to Control the Employee

The models we have discussed so far focus on the choice between centralization and delegation and mostly abstract from any actions that the owner might be able to take to control the employee. The ability of the owner to take actions to control the employee in many settings has been noted by Holmström (1977, 1984) and Jensen and Meckling (1992), among others. For instance, the owner may be able to constrain the decisions the employee can make. Or she may be able to use incentive contracts to align the employee's incentives. In this section we discuss these two options in turn.

5.3.1. Optimal Decision Processes

In the previous section we saw that even if decisions are *ex ante* and *ex post* noncontractible, a sufficiently patient owner can implement decision processes other than just centralization and delegation. This raises the question of what is the optimal process among all deterministic decision processes. A related question is what the owner should do if she is able to constrain the decisions the employee can make. It turns out that these two questions are actually the same, at least if one restricts attention to deterministic decision rules. In particular, Holmström (1977, 1984) observes that searching for the optimal deterministic decision rule is equivalent to figuring out how to optimally constrain the set of decisions from which the employee is allowed to choose. In this section we discuss Holmström (1977, 1984) and the literature built on his contributions.

Holmström (1977, 1984) starts by posing a general class of problems in which an owner needs to make a decision that affects her payoff and that of her employee. The employee privately observes the state of the world that determines each party's preferred decision. In contrast to the models we have examined so far, Holmström allows the owner to commit to any deterministic outcome function that maps the employee's recommendations to decisions. He does, however, continue to rule out contingent payments between the two parties.

In this setting, the Revelation Principle implies that one can restrict attention to direct mechanisms in which the employee truthfully reveals the state. Any deterministic direct mechanism in turn can be implemented through a constrained delegation scheme in which the owner of-

fers the employee a set of decisions from which the employee can then choose. The contracting problem therefore reduces to finding the optimal delegation set. However, the optimal delegation set can take essentially any form, the only requirement being that the employee's choice is well defined.

In practice, constrained delegation of this type is very common. Alfred Sloan, for instance, delegated the pricing decisions of the different types of cars to his division managers but allowed them to choose only prices between certain thresholds (Sloan 1964).

However, notice that although constrained delegation is weakly optimal in this setting, there may be other decision processes that might perform equally well. Mylovanov (2008), for instance, shows that when the optimal delegation set takes the form of a single interval, and the owner's and employee's preferred decision functions intersect at most once, the optimal delegation set can be implemented through veto delegation.

After posing the general problem and showing that a solution exists, Holmström (1977, 1984) shows that in the standard uniform-quadratic-constant-bias case the optimal delegation set takes the form of a single interval. As mentioned above, some papers have built on Holmström's work. For instance, Melumad and Shibano (1991), Martimort and Semenov (2006), Alonso and Matouschek (2008), and Amador and Bagwell (2011) characterize the optimal delegation set for more general settings. Koessler and Martimort (2012) and Frankel (2012) extend the standard setting by allowing for multidimensional decisions. Szalay (2005) endogenizes information acquisition by the employee. And Ambrus and Egorov (2009) allow for the owner to impose decision-contingent costs on the employee. All these papers focus on deterministic mechanisms. In contrast, Kovac and Mylovanov (2009) allow for stochastic mechanisms and provide sufficient conditions under which the optimal mechanism is in fact deterministic.

One insight from this literature is that even if the owner is able to commit to essentially any decision process, delegation can arise as the optimal one. And if delegation is not optimal, other commonly observed decision processes often do arise. For instance, Alonso and Matouschek (2008) and Amador and Bagwell (2011) provide necessary and sufficient conditions for the optimal delegation set to consist of a single interval. Even though commonly observed decision processes can arise as optimal, however, some of the widely held beliefs about how changes in the environment affect decisionmaking may no longer hold. For instance, Alonso and Matouschek (2008) show that if one allows for any decision process, the owner may give less discretion to an employee whose preferences are more aligned with the owner's or to one with a greater informational advantage.

To illustrate an application of this literature, albeit in a different context, consider the classic problem of regulating a monopolist who is privately informed about his costs (Baron and Myerson 1982). But in contrast to Baron and Myerson, suppose that transfers between the regulator and the monopolist are ruled out by law, as is often the case in practice. In particular, suppose that the monopolist can produce $q \geq 0$ units of a good at costs θq and faces a linear demand curve $d = A - Bq$, where d is the price, and A and B are both strictly positive parameters. Marginal costs s are drawn from a unimodal distribution with support $[0, 1]$. The monopolist's preferred price is then given by $d_s(s) = (A + s)/2$, whereas the regulator would like to set the price equal to marginal costs. Finally, suppose that the maximum willingness to pay A is larger than the largest marginal cost, that is, $A > 1$. Our discussion of this example follows Alonso and Matouschek (2008).

A first question is whether the regulator should give the firm any discretion over pricing. The answer is that the regulator should do so if and only if $A/2 < E(s)$. This condition is more likely to be satisfied the smaller is A , which is intuitive, because a reduction in A reduces the firm's preferred price. However, notice that the condition depends only on the mean of the cost distribution and not on any measure of its dispersion. The magnitude of the firm's informational advantage therefore has no effect on whether the regulator should give it any discretion over pricing.

The second question is then what prices the firm should be allowed to choose from, if the above condition is indeed satisfied. The answer turns out to be very simple: for any unimodal cost distribution, the regulator cannot do better than to let the firm choose any price below a threshold. This result is consistent with the widespread use of price-cap regulation in the United Kingdom and the United States (Armstrong and Sappington 2007). Finally, it can be shown that the price cap is decreasing in A and that the effect of a change in the firm's informational advantage is in general ambiguous. Even though this example deals with the organization of a regulatory system, it does suggest that the results from this literature can be fruitfully applied to similar issues that arise in firms, such as the design of capital budgeting rules and the design of transfer pricing schemes.

The general problem that this literature examines has been further explored in the context of various economic problems. For instance, Athey et al. (2005) study the optimal degree of discretion for a monetary authority. Amador et al. (2006) explore how an agent may want to restrict his future decisions in a setting where the agent expects to receive relevant information in the future but also knows that he suffers from self-control problems. Martimort and Semenov (2006) examine a delegation problem with multiple employees to analyze the organization of lobbying by interest groups. And in a setting with a different information structure than the one shared by the models that we have described so far, Armstrong and Vickers (2010) develop a model of delegated project choice and apply it to the optimal design of competition policy.

5.3.2. Endogenous Incentives

In addition to constraining the decisions the employee can make, the owner might be able to control the employee through the use of incentive contracts. Most of the literature on delegation assumes that the owner cannot rely on incentive contracts to align the employee's incentives, just as most of the literature on incentive contracts assumes that the owner cannot contract on the allocation of decision rights. One reason for ruling out incentive contracts is that they can be difficult to reconcile with the assumption that contracts are highly incomplete (see, for instance, Mookherjee, this volume), which is central to all the models of delegation discussed in this chapter. As observed by Jensen and Meckling (1992) and others, however, in practice the design of incentive contracts and the allocation of decision rights are interdependent. Indeed, there is a strong sense that the use of incentive contracts and delegation are not only interdependent but also complementary. For instance, Milgrom and Roberts (1992: 17) observe that "incentives and delegated authority are complements: each makes the other more valuable." In this section we discuss papers that explore this interdependence. In line with our discussion above, we distinguish between models in which there is only one employee and one decision and others in which there multiple employees and decisions.

5.3.2.1. Incentive Contracts for a Single Employee. The trade-off between risk and incentives is a central insight of the literature on incentive schemes. But as Prendergast (2002) observes, there is only weak empirical support for this trade-off. Indeed, if anything, firms appear to be using steeper incentives in more, rather than less, uncertain environments. To provide an explanation for this empirical pattern, Prendergast develops a model in which incentives and the allocation of decision rights are jointly determined. In the model, the more uncertain is the environment, the more discretion an owner gives her employee over what he can do on the job. And the more discretion the employee has over what he can do on the job, the more important it is that his pay be closely tied to his performance. In line with the observation above, incentives and discretion are therefore complements.¹⁴ And because incentives and discretion are complements, an increase in uncertainty that leads to more discretion also leads to steeper incentives.¹⁵

Krishna and Morgan (2008) also allow for endogenous incentives in a setting with a single employee. In their model, however, incentives and delegation are substitutes. Specifically, they consider a version of the model in Dessein (2002) in which the owner can make wages contingent on messages sent by the employee. They then show that even if wages cannot be negative, contingent wages improve communication under centralization. Because message-contingent wages do not improve decisionmaking under delegation, the owner's ability to commit to incentive contracts then makes centralization more likely. In this setting, incentive contracts and delegation are therefore substitute tools that the owner can use to make decisions more sensitive to the state.

A key difference between Prendergast (2002) and Krishna and Morgan (2008) is that in Prendergast's model incentive contracts improve decisionmaking under delegation, whereas in Krishna and Morgan's model they improve communication under centralization. In practice, incentive contracts may improve both decisionmaking and communication. For incentive contracts and delegation to be complements, as is often argued, incentive contracts must more effectively improve decisionmaking under delegation than they do communication under centralization. This observation once again highlights the importance of taking into account the endogenous nature of communication for the analysis of efficient organization.

5.3.2.2. Incentive Contracts for Multiple Employees. Some papers also explore the joint determination of incentives and the allocation of decision rights in a setting with multiple employees, who are typically thought of as division managers.

In Athey and Roberts (2001) there are two division managers who have two tasks: to put effort into increasing their divisions' profits and to select a project. A central insight of the paper is that there is a trade-off between motivating effort provision and project selection. In particular, effort provision is best motivated through relative pay, which filters out shocks that are common to both divisions; however, relative pay distorts project selection, because it gives

14. The literature has identified additional reasons for why incentives and delegation may be complements. For instance, Van den Steen (2010a) explores a setting in which the owner and the employee have different priors about what decision is most likely to lead to the successful conclusion of a project. In his setting, income and decision rights should be concentrated in the same party, because the party with the income right cares the most about the decision being correct, and the party with the decision right has the highest valuation of the project.

15. For a further discussion on why more uncertainty may lead to steeper incentives, see Raith (2008).

the division managers poor incentives to internalize externalities. Instead of letting the division managers decide on projects, it may then be better to move this decision to a third employee who gets paid a small fraction of overall profits. The model therefore provides a rationale for the incentives that are assumed in Alonso et al. (2008), Rantakari (2008), and other papers. As discussed in Section 5.1.3, these models assume that the CEO cares about overall profits, whereas division managers are biased toward their divisions.

In a related paper, Dessein et al. (2010) develop a model in which it is optimal to bias functional managers toward standardization and division managers toward local adaptation. They then explore the interdependence between the endogenously created incentive conflict and the allocation of decision rights. Friebel and Raith (2010) also explore the joint determination of incentives and the allocation of decision rights in a multidivisional firm. However, they focus on the boundary of the firm, which they allow to be endogenously determined. Rantakari (2011) takes the boundary of the firm as given but endogenizes the degree of operational integration, in addition to incentives and the allocation of decision rights. He argues that stable environments give rise to tightly integrated and centralized organizations with firm-wide incentives for division managers, whereas volatile environments induce firms to adopt more loosely integrated and decentralized organizational structures with division-level incentives for division managers.

6. Conclusion

A number of the first scholars to examine organizations (both within and outside the field of economics) identified both the crucial role of decisionmaking and the inherently strategic nature of communication. In the intervening years, a large volume of economic research has explored these ideas (although often without being aware of their origins). In this chapter we have surveyed this work.

It should be clear that there remain many open theoretical questions. For example, each individual line of inquiry could be deepened, as we described the cheap-talk literature moving from its origins in Section 2.2 to recent work in Section 3.1. Furthermore, there are likely to be advances from combining some of the approaches described here, such as by adding limited opportunities for contracting to several of the existing research streams, as we began to describe in Section 3.6. Finally, the question of what an organization might do to ameliorate some of the inefficiencies we describe has only just begun to be analyzed, so there is much work to do in continuing that described in Sections 4 and 5.

Probably more important than all these opportunities for more theory, however, we believe that—as is true throughout most of organizational economics—the greatest need is for empirical work in this area. We hope this need will begin to be met soon.

REFERENCES

- Admati, Anat, and Paul Pfleiderer. 1997. "Does It All Add Up? Benchmarks and the Compensation of Active Portfolio Managers." *Journal of Business* 70: 323–350.
- Aghion, Philippe, and Jean Tirole. 1997. "Formal and Real Authority in Organizations." *Journal of Political Economy* 105:1–29.
- Aghion, Philippe, Mathias Dewatripont, and Patrick Rey. 2002. "On Partial Contracting." *European Economic Review* 46(papers and proceedings): 745–753.
- . 2004. "Transferable Control." *Journal of the European Economic Association* 2: 115–138.
- Alonso, Ricardo, and Niko Matouschek. 2007. "Relational Delegation." *RAND Journal of Economics* 38: 1070–1089.
- . 2008. "Optimal Delegation." *Review of Economic Studies* 75: 259–293.
- Alonso, Ricardo, Wouter Dessein, and Niko Matouschek. 2008. "When Does Coordination Require Centralization?" *American Economic Review* 98: 145–179.
- Amador, Manuel, and Kyle Bagwell. 2011. "The Theory of Optimal Delegation with an Application to Tariff Caps." Working paper, Stanford University, Stanford, CA.
- Amador, Manuel, Ivan Werning, and George-Marios Angeletos. 2006. "Commitment vs. Flexibility." *Econometrica* 74: 365–396.
- Ambrus, Attila, and Georgy Egorov. 2009. "Delegation and Nonmonetary Incentives." Working paper, Harvard University, Cambridge, MA.
- Aoki, Masahiko. 1986. "Horizontal vs. Vertical Information Structure of the Firm." *American Economic Review* 76: 971–983.
- Armstrong, Mark, and David Sappington. 2007. "Recent Developments in the Theory of Regulation," in Mark Armstrong and Robert Porter (eds.), *Handbook of Industrial Organization*, volume 3. Amsterdam: North-Holland, pp. 1557–1700.
- Armstrong, Mark, and John Vickers. 2010. "A Model of Delegated Project Choice." *Econometrica* 78: 213–244.
- Arrow, Kenneth. 1974. *The Limits of Organization*. New York: W. W. Norton.
- Athey, Susan, and John Roberts. 2001. "Organizational Design: Decision Rights and Incentive Contracts." *American Economic Review* 91: 200–205.
- Athey, Susan, Joshua Gans, Scott Schaefer, and Scott Stern. 1994. "The Allocation of Decisions in Organizations." Stanford University Graduate School of Business Research Paper 1322, Stanford, CA.
- Athey, Susan, Andrew Atkeson, and Patrick Kehoe. 2005. "The Optimal Degree of Discretion in Monetary Policy." *Econometrica* 73: 1431–1475.
- Austen-Smith, David, and Jeffrey Banks. 2000. "Cheap Talk and Burned Money." *Journal of Economic Theory* 91: 1–16.
- Baker, George, Robert Gibbons, and Kevin J. Murphy. 1999. "Informal Authority in Organizations." *Journal of Law, Economics, and Organization* 15: 56–73.
- Barnard, Chester. 1938. *The Functions of the Executive*. Cambridge, MA: Harvard University Press.
- Baron, David, and Roger Myerson. 1982. "Regulating a Monopolist with Unknown Costs." *Econometrica* 50: 911–930.
- Battaglini, Marco. 2002. "Multiple Referrals and Multidimensional Cheap Talk." *Econometrica* 70: 1379–1401.
- Bénabou, Roland, and Guy Laroque. 1992. "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility." *Quarterly Journal of Economics* 107: 921–958.
- Bester, Helmut. 2009. "Externalities, Communication and the Allocation of Decision Rights." *Economic Theory* 41: 269–296.
- Bester, Helmut, and Daniel Krämer. 2008. "Delegation and Incentives." *RAND Journal of Economics* 39: 664–682.
- Bhattacharya, Sudipto, and Paul Pfleiderer. 1985. "Delegated Portfolio Management." *Journal of Economic Theory* 36: 1–25.
- Blanes i Vidal, Jordi, and Marc Möller. 2007. "When Should Leaders Share Information with Their Subordinates?" *Journal of Economics and Management Strategy* 16: 251–283.

- Blume, Andreas, Oliver J. Board, and Kohci Kawamura. 2007. "Noisy Talk." *Theoretical Economics* 2: 395–440.
- Brocas, Isabelle, and Juan Carrillo. 2007. "Influence through Ignorance." *RAND Journal of Economics* 38: 931–947.
- Caillaud, Bernard, and Jean Tirole. 2007. "Consensus Building: How to Persuade a Group." *American Economic Review* 97: 1877–1900.
- Callander, Steven. 2008a. "Political Motivations." *Review of Economic Studies* 75: 671–697.
- . 2008b. "A Theory of Policy Expertise." *Quarterly Journal of Political Science* 3(2): 123–140.
- Calvó-Armengol, Antoni, Joan de Martí, and Andrea Prat. 2011. "Communication and Influence." Working Paper 1-60, London School of Economics, London.
- Chakraborty, Archishman, and Rick Harbaugh. 2007. "Comparative Cheap Talk." *Journal of Economic Theory* 132: 70–94.
- Che, Yeon-Koo, and Navin Kartik. 2009. "Opinions as Incentives." *Journal of Political Economy* 117: 815–860.
- Che, Yeon-Koo, Wouter Dessein, and Navin Kartik. 2012. "Pandering to Persuade." *American Economic Review*, forthcoming.
- Coase, Ronald. 1960. "The Problem of Social Cost." *Journal of Law and Economics* 3: 1–44.
- Crawford, Vincent, and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50: 1431–1451.
- Crémer, Jacques, Luis Garicano, and Andrea Prat. 2007. "Language and the Theory of the Firm." *Quarterly Journal of Economics* 122: 373–407.
- Cyert, Richard, and James March. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice-Hall.
- DeGroot, Morris H. 1970. *Optimal Statistical Decisions*. New York: McGraw-Hill.
- Demski, Joel, and David Sappington. 1987. "Delegated Expertise." *Journal of Accounting Research* 25: 68–89.
- Dessein, Wouter. 2002. "Authority and Communication in Organizations." *Review of Economic Studies* 69: 811–838.
- Dessein, Wouter, and Tano Santos. 2006. "Adaptive Organizations." *Journal of Political Economy* 114: 956–995.
- Dessein, Wouter, Luis Garicano, and Robert Gertner. 2010. "Organizing for Synergies." *American Economic Journal: Microeconomics* 2(4): 77–114.
- Dewatripont, Mathias, and Jean Tirole. 1999. "Advocates." *Journal of Political Economy* 107: 1–39.
- Dye, Ronald. 1985. "Disclosure of Nonproprietary Information." *Journal of Accounting Research* 23: 123–145.
- Epstein, David, and Sharyn O'Halloran. 1999. *Delegating Powers: A Transaction Cost Politics Approach to Policy Making under Separate Powers*. Cambridge: Cambridge University Press.
- Fama, Eugene, and Michael Jensen. 1983. "Separation of Ownership and Control." *Journal of Law and Economics* 26: 301–325.
- Farrell, Joseph, and Robert Gibbons. 1989. "Cheap Talk with Two Audiences." *American Economic Review* 79: 1214–1223.
- Feddersen, Timothy, and Wolfgang Pesendorfer. 1998. "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting." *American Political Science Review* 92: 23–35.
- Feldman, Martha, and James G. March. 1981. "Information in Organizations as Signal and Symbol." *Administrative Science Quarterly* 26: 171–186.
- Frankel, Alex. 2012. "Delegating Multiple Decisions." Working paper, University of Chicago Booth School of Business.
- Friebel, Guido, and Michael Raith. 2010. "Resource Allocation and Organizational Form." *American Economic Journal: Microeconomics* 2(2): 1–33.
- Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. Cambridge, MA: MIT Press.
- Galeotti, Andrea, Christian Ghiglino, and Francesco Squintani. 2011. "Strategic Information Transmission in Networks." Economics Discussion Paper 668, University of Essex, UK.
- Garicano, Luis. 2000. "Hierarchies and the Organization of Knowledge in Production." *Journal of Political Economy* 108: 874–904.
- Garicano, Luis, and Tano Santos. 2004. "Referrals." *American Economic Review* 94: 499–525.
- Geanakoplos, John, and Paul Milgrom. 1991. "A Theory of Hierarchy Based on Limited Managerial Attention." *Journal of the Japanese and International Economies* 5: 205–225.
- Gibbons, Robert. 1988. "Learning in Equilibrium Models of Arbitration." *American Economic Review* 79: 896–912.
- . 2005. "Four Formal(izable) Theories of the Firm?" *Journal of Economic Behavior and Organization* 58: 200–245.
- . 2010. "Inside Organizations: Pricing, Politics, and Path Dependence." *Annual Review of Economics* 2: 337–365.
- Gilligan, Thomas, and Keith Krehbiel. 1987. "Collective Decisionmaking and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures." *Journal of Law, Economics, and Organization* 3: 287–335.
- . 1989. "Asymmetric Information and Legislative Rules with Heterogeneous Committees." *American Journal of Political Science* 33: 459–490.
- Goltsman, Maria, and Gregory Pavlov. 2011. "How to Talk to Multiple Audiences." *Games and Economic Behavior* 72: 100–122.
- Goltsman, Maria, Johannes Hörner, Gregory Pavlov, and Francesco Squintani. 2009. "Mediation, Arbitration and Negotiation." *Journal of Economic Theory* 144: 1397–1420.
- Green, Jerry, and Nancy Stokey. 2007. "A Two-Person Game of Information Transmission." *Journal of Economic Theory* 135: 90–104.
- Grossman, Sanford. 1981. "The Role of Warranties and Private Disclosure about Product Quality." *Journal of Law and Economics* 24: 461–483.
- Grossman, Sanford, and Oliver Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy* 94: 691–719.
- Hagenbach, Jeanne, and Frédéric Koessler. 2010. "Strategic Communication Networks." *Review of Economic Studies* 77: 1072–1099.
- Harris, Milton, and Artur Raviv. 2005. "Allocation of Decision-Making Authority." *Review of Finance* 9: 353–383.
- Hart, Oliver, and John Moore. 1990. "Property Rights and the Nature of the Firm." *Journal of Political Economy* 98: 1119–1158.
- . 2005. "On the Design of Hierarchies: Coordination versus Specialization." *Journal of Political Economy* 113: 675–702.
- Hayek, Friedrich. 1945. "The Use of Knowledge in Society." *American Economic Review* 35: 519–530.
- Hermalin, Benjamin. 1998. "Toward an Economic Theory of Leadership: Leading by Example." *American Economic Review* 88: 1188–1206.
- Holmström, Bengt. 1977. *On Incentives and Control in Organizations*. PhD thesis, Stanford University, Stanford, CA.
- . 1982/1999. "Managerial Incentive Problems: A Dynamic Perspective," in Björn Walroos, Bo-Göran Ekholm, Alf-Erik Lerviks, and Hendrik Meyer (eds.), *Essays in Economics and Management in Honor of Lars Wahlbeck*. Helsinki: Swedish School of Economic and Business Administration, pp. 209–230. Republished in *Review of Economic Studies* 66: 169–182.

- Holmström, Bengt. 1984. "On the Theory of Delegation," in M. Boyer and R. Kihlstrom (eds.), *Bayesian Models in Economic Theory*. New York: North-Holland, pp. 115–141.
- Huber, John, and Charles Shipan. 2006. "Politics, Delegation, and Bureaucracy," in Barry Weingast and Donald Wittman (eds.), *The Oxford Handbook of Political Economy*. Oxford: Oxford University Press.
- Hurwicz, Leonid. 1960. "Optimality and Informational Efficiency in the Design of Resource Allocation Processes," in Kenneth Arrow, Samuel Karlin, and Patrick Suppes (eds.), *Mathematical Methods in the Social Sciences 1959*. Palo Alto: Stanford University Press, pp. 27–46.
- . 1972. "On Informationally Decentralized Systems," in Charles McGuire and Roy Radner (eds.), *Decision and Organization: A Volume in Honor of Jacob Marschak*. Amsterdam: North-Holland, pp. 279–336.
- . 1973. "The Design of Mechanisms for Resource Allocation." *American Economic Review* 63: 1–30.
- Inderst, Roman, and Manuel Klein. 2007. "Innovation, Endogenous Overinvestment, and Incentive Pay." *RAND Journal of Economics* 38: 881–904.
- Ivanov, Maxim. 2010. "Communication via a Strategic Mediator." *Journal of Economic Theory* 145: 869–884.
- Jensen, Michael, and William Meckling. 1992. "Specific and General Knowledge and Organizational Structure," in Lars Werin and Hans Wijkander (eds.), *Contract Economics*. Oxford: Blackwell, pp. 251–274.
- Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101: 2590–2615.
- Kartik, Navin. 2007. "A Note on Cheap Talk and Burned Money." *Journal of Economic Theory* 136: 749–758.
- . 2009. "Strategic Communication with Lying Costs." *Review of Economic Studies* 76: 1359–1395.
- Kartik, Navin, Marco Ottaviani, and Francesco Squintani. 2007. "Credulity, Lies, and Costly Talk." *Journal of Economic Theory* 134: 93–116.
- Knight, Frank. 1921/1964. *Risk, Uncertainty, and Profit*. New York: Hart, Schaffner, and Marx. Reprinted by New York: Harper and Row.
- Koessler, Frederic, and David Martimort. 2012. "Optimal Delegation with Multi-Dimensional Decisions." *Journal of Economic Theory*, forthcoming.
- Kolotilin, Anton. 2012a. "Experimental Design to Persuade." Working paper, Massachusetts Institute of Technology, Cambridge, MA.
- . 2012b. "Optimal Information Disclosure to an Informed Decision-Maker." Working paper, Massachusetts Institute of Technology, Cambridge, MA.
- Kolotilin, Anton, Hao Li, and Wei Li. 2012. "Optimal Limited Authority for Principal." Working paper, Massachusetts Institute of Technology, Cambridge, MA.
- Kovac, Eugen, and Tymofiy Mylovanov. 2009. "Stochastic Mechanisms in Settings without Monetary Transfers: The Regular Case." *Journal of Economic Theory* 144: 1373–1395.
- Krähmer, Daniel. 2006. "Message-Contingent Delegation." *Journal of Economic Behavior and Organization* 60: 490–506.
- Krishna, Vijay, and John Morgan. 2001. "Asymmetric Information and Legislative Rules: Some Amendments." *American Political Science Review* 95: 435–452.
- . 2004. "The Art of Conversation: Eliciting Information from Experts through Multi-Stage Communication." *Journal of Economic Theory* 117: 147–179.
- . 2008. "Contracting for Information under Imperfect Commitment." *RAND Journal of Economics* 39: 905–925.
- Lambert, Richard. 1986. "Executive Effort and Selection of Risky Projects." *RAND Journal of Economics* 17: 77–88.
- Landier, Augustin, David Sraer, and David Thesmar. 2009. "Optimal Dissent in Organizations." *Review of Economic Studies* 76: 761–794.
- Levy, Gilat. 2007. "Decision Making in Committees: Transparency, Reputation, and Voting Rules." *American Economic Review* 97: 150–168.
- Malcomson, James. 2009. "Principal and Expert Agent." *B.E. Journal of Theoretical Economics* 9 (contributions): Article 17.
- . 2011. "Do Managers with Limited Liability Take More Risky Decisions? An Information Acquisition Model." *Journal of Economics and Management Strategy* 20: 83–120.
- March, James. 1962/1988. "The Business Firm as a Political Coalition." *Journal of Politics* 24: 662–678. Reprinted in James March (ed.), *Decisions and Organizations*. Oxford: Basil Blackwell, pp. 101–115.
- . 1981. "Decisions in Organizations and Theories of Choice," in Andrew Van de Ven and William Joyce (eds.), *Perspectives on Organization Design and Behavior*. New York: Wiley Interscience, pp. 205–244.
- Marino, Anthony, John Matsusaka, and Jan Zábajník. 2010. "Disobedience and Authority." *Journal of Law, Economics, and Organization* 26: 427–459.
- Marschak, Jacob. 1971. "The Economics of Information Systems," in Michael Intriligator, (ed.), *Frontiers of Quantitative Economics*. Amsterdam: North-Holland, pp. 32–107.
- Marschak, Jacob, and Roy Radner. 1972. *Economic Theory of Teams*. New Haven, CT: Yale University Press.
- Martimort, David, and Aggey Semenov. 2006. "Continuity in Mechanism Design without Transfers." *Economics Letters* 93: 182–189.
- Maskin, Eric. 1977/1999. "Nash Equilibrium and Welfare Optimality." Mimeo, Massachusetts Institute of Technology, Cambridge, MA. Reprinted in *Review of Economic Studies* 66: 23–38.
- Melumad, Nahum, and Toshi Shibano. 1991. "Communication in Settings with No Transfers." *RAND Journal of Economics* 22: 173–190.
- Meyer, Margaret, Paul Milgrom, and John Roberts. 1992. "Organizational Prospects, Influence Costs, and Ownership Changes." *Journal of Economics and Management Strategy* 1: 9–35.
- Milgrom, Paul. 1981. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics* 12: 380–391.
- . 1988. "Employment Contracts, Influence Activities, and Efficient Organization Design." *Journal of Political Economy* 96: 42–60.
- . 2008. "What the Seller Won't Tell You: Persuasion and Disclosure in Markets." *Journal of Economic Perspectives* 22(2): 115–132.
- Milgrom, Paul, and John Roberts. 1986. "Relying on the Information of Interested Parties." *RAND Journal of Economics* 17: 18–32.
- . 1988. "An Economic Approach to Influence Activities in Organizations." *American Journal of Sociology* 94: S154–S179.
- . 1992. *Economics, Organization and Management*. Englewood Cliffs, NJ: Prentice-Hall.
- . 1999. "The Internal Politics of the Firm," in Samuel Bowles, Maurizio Franzini, and Ugo Pagano (eds.), *The Politics and Economics of Power*. London: Routledge, pp. 46–62.
- Mintzberg, Henry. 1979. *The Structuring of Organizations: A Synthesis of the Research*. Englewood Cliffs, NJ: Prentice-Hall.
- Morris, Stephen. 2001. "Political Correctness." *Journal of Political Economy* 109: 231–265.
- Mount, Kenneth, and Stanley Reiter. 1990. "A Model of Computing with Human Agents." Discussion Paper 890, Northwestern University Center for Mathematical Studies in Economics and Management Science, Evanston, IL.

- Myerson, Roger. 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6: 58–73.
- Mylovannov, Tymofiy. 2008. "Veto-Based Delegation." *Journal of Economic Theory* 138: 297–307.
- Okuno-Fujiwara, Masahiro, Andrew Postlewaite, and Kotaro Suzumura. 1990. "Strategic Information Revelation." *Review of Economic Studies* 57: 25–47.
- Ottaviani, Marco, and Peter Sørensens. 2006a. "Reputational Cheap Talk." *RAND Journal of Economics* 37: 155–175.
- . 2006b. "The Strategy of Professional Forecasting." *Journal of Financial Economics* 81: 441–466.
- Persico, Nicola. 2004. "Committee Design with Endogenous Information." *Review of Economic Studies* 71: 165–191.
- Powell, Michael. 2012. "Influence Cost Models of Firm's Boundaries and Control Structures." Working paper, Northwestern University Kellogg School of Management, Evanston, IL.
- Prendergast, Canice. 1993. "A Theory of 'Yes Men.'" *American Economic Review* 83: 757–770.
- . 2002. "The Tenuous Trade-off between Risk and Incentives." *Journal of Political Economy* 110: 1071–1102.
- Raith, Michael. 2008. "Specific Knowledge and Performance Measurement." *RAND Journal of Economics* 39: 1059–1079.
- Rajan, Raghuram, and Luigi Zingales. 2000. "The Tyranny of Inequality." *Journal of Public Economics* 76: 521–558.
- Rajan, Raghuram, Henri Servaes, and Luigi Zingales. 2000. "The Cost of Diversity: The Diversification Discount and Inefficient Investment." *Journal of Finance* 55: 35–80.
- Rantakari, Heikki. 2008. "Governing Adaptation." *Review of Economic Studies* 75: 1257–1285.
- . 2011. "Organizational Design and Environmental Volatility." *Journal of Law, Economics, and Organization*, forthcoming.
- . 2012. "Employee Initiative and Managerial Control." *American Economic Journal: Microeconomics* 4(3): 171–211.
- Rayo, Luis, and Ilya Segal. 2010. "Optimal Information Disclosure." *Journal of Political Economy* 118: 949–987.
- Rotemberg, Julio. 1993. "Power in Profit-Maximizing Organizations." *Journal of Economics and Management Strategy* 2: 165–198.
- Scharfstein, David, and Jeremy Stein. 2000. "The Dark Side of Internal Capital Markets: Divisional Rent-Seeking and Inefficient Investment." *Journal of Finance* 55: 2537–2564.
- Seidmann, Daniel, and Eyal Winter. 1997. "Strategic Information Transmission with Verifiable Messages." *Econometrica* 65: 163–169.
- Simon, Herbert. 1951. "A Formal Theory of the Employment Relationship." *Econometrica* 19: 293–305.
- Skaperdas, Stergios. 1992. "Cooperation, Conflict, and Power in the Absence of Property Rights." *American Economic Review* 82: 720–739.
- Sloan, Alfred P. 1964. *My Years with General Motors*, John McDonald (ed.). Garden City, NJ: Doubleday.
- Sobel, Joel. 1985. "A Theory of Credibility." *Review of Economic Studies* 52: 556–573.
- . 2012. "Giving and Receiving Advice," in Daron Acemoglu, Manuel Arellano, and Eddie Dekel (eds.), *Advances in Economics and Econometrics*. Cambridge: Cambridge University Press, forthcoming.
- Spence, Michael. 1973. "Job Market Signaling." *Quarterly Journal of Economics* 87: 355–374.
- Stracca, Livio. 2006. "Delegated Portfolio Management: A Survey of the Theoretical Literature." *Journal of Economic Surveys* 20: 823–848.
- Szalay, Dezso. 2005. "The Economics of Extreme Options and Clear Advice." *Review of Economic Studies* 72: 1173–1198.
- Tullock, Gordon. 1980. "Efficient Rent Seeking," in James Buchanan, Robert Tollison, and Gordon Tullock (eds.), *Toward a Theory of the Rent-Seeking Society*. College Station: Texas A&M University Press, pp. 97–112.
- Van den Steen, Eric. 2002. "Disagreement and Information Collection." Working paper, Massachusetts Institute of Technology, Cambridge, MA.
- . 2006. "Limits of Authority: Motivation versus Coordination." Working paper, Massachusetts Institute of Technology, Sloan School of Management, Cambridge, MA.
- . 2010a. "Disagreement and the Allocation of Control." *Journal of Law, Economics, and Organization* 26: 385–426.
- . 2010b. "Interpersonal Authority in a Theory of the Firm." *American Economic Review* 100: 466–490.
- Visser, Bauke, and Otto Swank. 2007. "On Committees of Experts." *Quarterly Journal of Economics* 122: 337–372.
- Williamson, Oliver. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.
- Wulf, Julie. 2002. "Internal Capital Markets and Firm-Level Compensation Incentives for Division Managers." *Journal of Labor Economics* 20: S219–S262.
- . 2007. "Authority, Risk and Performance Incentives: Evidence from Division Manager Positions inside Firms." *Journal of Industrial Economics* 55: 169–196.
- . 2009. "Influence and Inefficiency in the Internal Capital Market." *Journal of Economic Behavior and Organization* 72: 305–321.
- Zábojník, Jan. 2002. "Centralized and Decentralized Decisions in Organizations." *Journal of Labor Economics* 20: 1–22.
- Zald, Mayer, and Michael Berger. 1978. "Social Movements in Organizations: Coup d'Etat, Insurgency, and Mass Movements." *American Journal of Sociology* 83: 823–861.
- Zermeño, Luis. 2012. "A Principal-Expert Model and the Value of Menus." Working paper, Instituto Tecnológico Autónomo de México, Mexico City.