

A Theory of Rule Development

Glenn Ellison and Richard Holden*

October 12, 2007

Abstract

This paper develops a model with endogenously coarse rules. A principal hires an agent to take an action. The principal knows the optimal state-contingent action, but cannot communicate it perfectly due to communication constraints. The principal can use previously realized states as examples to define rules of varying breadth. We analyze how rules are chosen under several assumptions about how rules can be amended. We explore the inefficiencies that arise and how they depend on the ability to refine rules, the principal's time horizon and patience, and other factors. Our model exhibits path dependence in that the efficacy of rule development depends on the sequence of realizations of the state. We interpret this as providing a foundation for persistent performance differences between similar organizations and explore the role of different delegation structures in ameliorating the effects of bounded communication.

*Ellison: Massachusetts Institute of Technology, Department of Economics, E52-380a, 50 Memorial Drive, Cambridge MA 02142. E-mail: gellison@mit.edu. Holden: Massachusetts Institute of Technology, Sloan School of Management, E52-410, 50 Memorial Drive, Cambridge MA 02142. E-mail: rholden@mit.edu. We are deeply indebted to Bob Gibbons for numerous discussions and suggestions. We also thank Oliver Hart and Bengt Holmstrom, and seminar participants at Harvard and MIT for helpful comments.

1 Introduction

Firms and bureaucracies cannot function without rules. These rules lead them to take inefficient decisions in some circumstances. One long-recognized reason why organizations will be inefficient is that the world is highly complex and it is practically impossible to provide agents with the information-contingent plans they would need to take the act optimally.¹ In this paper, we develop a model of endogenously coarse rules due to imperfect communication. The agency problem in our model is not one of incentives and private benefits: we assume the agent follows any rules that are set out for him. The difficulty is instead that it will typically be impossible to describe the first-best decision rule to the agent. We discuss the inefficiencies that arise and how they are related to the assumptions about the environment.

We model rule development as a constrained-optimal process in a world with two types of communication constraints. The first of these is that the organization's agents must react to events they observe before communicating with the principal. The second is that it is impossible for the principal to communicate a complete contingent plan to the agent. Our motivation is that we think of states of nature as being very complex objects and it would be difficult to describe any state completely or provide a list of all possible states in advance of the principal and agent having shared experiences. A novel element of our model is how we approach partial describability: we assume that the only classes of states that can be described are sets of states that are similar to a previously realized states, but that there is a commonly understood notion of distance that allows the breadth of a rule to be a choice variable.

Section 2 describes our model. There are two players: a principal and an agent. There is a discrete set of time periods. In each period a new state of the world arises. The agent observes the state and must take an action. Inefficient decisions are sometimes taken not because of any incentive issues, but because of one communication constraint: the agent must act before he can communicate with the principal (and has no prior knowledge of the mapping from states to optimal actions). We assume that the agent will always follow

¹Some examples of the early literature on the difficulties of reacting to changing circumstances include Knight 1921; Barnard 1938; Hayek 1945.

any rule that has been communicated to him, but these will typically incomplete and/or suboptimal because of the other communication constraint: the principal is restricted to the use of analogy-based rules specifying that an action be taken whenever the new state is within some distance of a previously realized state. The principal does observe the state and have an opportunity to promulgate new rules at the end of each period, so rulebooks will become more complex and efficient over time. We often compare outcomes under three different assumptions about the extent to which rules can be changed over time. We call these the no overwriting, incremental overwriting, and vast overwriting cases.

Section 3 discusses the simplest case in which rules exist: a two period model in which a rule for second-period behavior can make reference to the first-period state. A primary result is that rules have excess breadth and sometimes produce incorrect decisions. Indeed, at the margin the first-period rule is completely worthless and must produce as many incorrect decisions as correct decisions. Section 4 examines a three period model. A primary assumption here is that option-value considerations lead the principal to narrow rules. In the no overwriting case, this may even include refusing to expand the breadth of a rule even though it would be also be correct in all marginal cases. Section 5 examines the infinite-horizon case. Here, nearly optimal rules must develop eventually if the principal is very patient. Whether exactly optimal rules develop depend on the assumptions about how rules can be revised.

Section 6 builds on the model to discuss governance structures in firms and how they may evolve. There are now two tasks, the principal has access to two agents, and can also perform a task himself. Time constraints force the principal to choose between delegating both tasks with communication, or performing one task personally and delegating the other without communication. Section 7 develops a connection to the literature on endogenous categorization. Section 8 concludes.

Our paper relates to a number of literatures. We noted above that there is a long history of studying adaptation. Knight (1921), in further developing what has come to be known as Knightian uncertainty, emphasized the role that entrepreneurs have in overcoming “the mania of change and advance which characterizes modern life”. In developing the first real theory of incentives in management, Barnard (1938) emphasized that coordinated

adaptation must be facilitated by management in a “conscious, deliberate and purposeful” manner. Hayek (1945) claimed that the fundamental economic problem is “adaptation to changes in the particular circumstances of time and place”. In contrast to Barnard, for Hayek the genius of the invisible hand of the price mechanism is that it adapts quickly to change, and that it economizes on information in doing so. As he puts it, the “marvel of the market...[is in] how little the individual participants need to know to be able to take the right action.” For Williamson, “adaptation to disturbances [is] the central problem of economic organization” and “markets and hierarchies...each possess distinctive strengths and weaknesses, where hierarchy enjoys the advantage for managing cooperative adaptations, and the market for autonomous adaptations.” (Williamson (2005))

We are far from being the first to consider the issue of communication formally. There is a vast literature, pioneered by Marschak and Radner (1972), on optimal communication structures within organizations. Arrow (1974, 1984) introduced the notion of a *code* in an organization as “all the known ways, whether or not inscribed in formal rules, for conveying information.” He focuses on the role of codes on returns to scale, and on the irreversible nature of the investment in developing a code. More recently, Cremer et al. (2007) formalize certain aspects of Arrow (1974). They analyze a model with boundedly rational agents where a common language reduces the cost of irrationality, but comes at the expense of being narrowly tailored to the state of the world. They explore the implications of this tradeoff for optimal scope of organizations.

Our model is also related to the literature on thinking via categories (Mullainathan (2002), Fryer and Jackson (2007), Mullainathan et al. (2008)). This literature—for the most part—documents the implications of a decision maker having a smaller number of mental categories than is required to perfectly partition events. One way to think of our paper is providing—in a specific context—a foundation for where these categories come from. In our model, the categories are the rules which the principal gives the agent, and they arise as an endogenous response to the bounded communication.

Our discussion of governance structures relates to the much larger theory on why some transactions take place in the market, while others take place inside firms.² A complete

²Coase (1937), Klein et al. (1978), Williamson (1971), Williamson (1979), Grossman and Hart (1986),

formal treatment of the tradeoff between the market and the firm in adapting to changed circumstances would involve the following. Fix a particular *transaction* and then compare allocative efficiency in the market and the firm. To do this, however, one needs a model of price formation. Without such a model it seems impossible to talk about degrees of efficiency of the price mechanism. We return to this issue in the final section, but we do not intend to provide such a model in this paper. Instead our section 6 just tackles a more modest task: highlighting a particular friction that occurs within the authority relationship and analyzing how different delegation structures may be used in response to the evolution of this friction. In this regard, our paper is related to Simon (1951), which analyzes a model in which a principal and an agent have preferences over a state contingent action. The state is common knowledge, but contracts are incomplete so that the Pareto efficient state contingent contract is not enforceable. Simon provides conditions under which it is optimal to specify a non-contingent action ex ante rather than allow the principal to mandate an action after observing the state. Baker et al. (2006) introduce cash flow rights which can be allocated ex ante and multiple decisions rights into Simon's framework in order to analyze how allocation of control can facilitate relational contracting.

2 The Model

2.1 Statement of the Problem

There are two players: the P(incipal) and the A(gent). P hires A to react to idiosyncratic situations and take actions on her behalf in each period $t = 1, \dots, T$. At the start of each period, the principal issues a set of rules R_t . The Agent then observes the state of nature $\omega \in \Omega$ and chooses an action $a \in \{-1, 1\}$. The principal then receives a payoff $\pi(a, \omega)$ that depends on A's action and the state of nature. The principal's payoff in the full game is the discounted sum of his per period payoffs: $V = \delta^t \sum_{t=1}^T \pi(a_t, \omega_t)$.

There are no hidden action/private benefit issues in our model. The first best would obtain if the principal at each time t issued rules instructing the agent to choose $a_t \in \text{Argmax}_a \pi(a, \omega_t)$. Constraints we impose on the communication, however, will usually

Hart and Moore (1990).




Figure 1: Model timing

make this impossible.

Our model of limited communication is built around a distance function $d : \Omega \times \Omega \rightarrow \mathbb{R}$. A rule r is a quadruple $r = (\omega, d, a, p)$. This is interpreted as prescribing that action a should be chosen in period t if $|\omega_t - \omega| < d$. The extra parameter p will be used to give some rules higher priority than others. We will sometimes write $\omega(r)$, $d(r)$, $a(r)$, and $p(r)$ for the function giving the relevant component of the rule r .

A rule book R is a finite set of rules with a well defined precedence order: $r, r' \in R \implies p(r) \neq p(r')$. We assume that the agent follows any rule book given to him. If two or more rules conflict, the one with the highest priority is followed. If no rules apply the agent randomizes 50:50. Formally, write $D(R)$ for the set of states covered by at least one rule, $D(R) = \{\omega' \in \Omega \text{ s.t. } |\omega' - \omega| < d \text{ for some } (\omega, d, a, p) \in R\}$. We assume that the agent's action choice is

$$a_t = \begin{cases} a & \text{if } (\omega, d, a, p) \in R, |\omega_t - \omega| < d, \text{ and } p = \max_{r: \omega_t \in D(\{r\})} p(r) \\ \frac{1}{2} \cdot -1 + \frac{1}{2} \cdot 1 & \text{if } \omega_t \notin D(R_t). \end{cases}$$

Note that we have made assumptions directly on the agent's behavior rather than giving him a utility function. There are two possible interpretations of this. The first is that the agent as a robot who mechanically follows rules. For example, the agent could be a low skilled employee who simply does what he is told, but lacks the sophistication to. The second interpretation is that agents are short-run players with unknown private benefits that will result in their apparently choosing actions randomly if they have no financial incentives, and in which the feasible incentive contracts are exactly those that punish agents if they fail to follow a pre-specified rule. That is, other contracts are not enforceable.

The restriction to distance-based rules is by itself not much of a limitation on the principal: the first best could always be approximated (under appropriate regularity conditions)

by a long finite list of rules. We will, however, also impose two other major restrictions on the set \mathfrak{R}_t of feasible rule books. First, we assume throughout the paper that only rules based on past observations are feasible and that at most one rule may be based on any past observation:

$$\begin{aligned} r \in \mathfrak{R}_t &\implies \omega(r) = \omega_{t'} \text{ for some } t' < t \\ r, r' \in \mathfrak{R}_t &\implies \omega(r) \neq \omega(r'). \end{aligned}$$

Second, we will sometimes impose restrictions on how rules are changed over time. We consider three variants of our model

1. *No Overwriting.* The only feasible change to the rule book is to add a single rule that references the most recent state and applies to a previously uncovered domain, i.e. $R_t \in \mathfrak{R}_t \implies R_t = R_{t-1} \cup \{r\}$ for some rule r with $\omega(r) = \omega_{t-1}$ and $D(\{r\}) \cap D(R_{t-1}) = \emptyset$.
2. *Incremental Overwriting.* Once again the feasible changes are to add a single rule that references the most recent state. In this specification, however, the new rule is allowed to overlap with one or more previous rules and takes precedence if the new rules conflict, i.e. $R_t \in \mathfrak{R}_t \implies R_t = R_{t-1} \cup \{r\}$ for some rule r with $\omega(r) = \omega_{t-1}$ and $p(r) > p(r')$ for all $r' \in R_{t-1}$.
3. *Vast Overwriting.* In this specification there is no additional restriction on R_t . The principal can design an entirely new rule book in each period.

We have stated the model for an arbitrary state space Ω and have in mind applications where Ω is a complicated set whose elements would be difficult to describe without reference to a previous example. Our analyses, however, will all take place in a simpler setting: we assume henceforth that $\Omega = [0, 1]$ and define distances as if the state space was a circle, $d(x, y) = \min\{|x - y|, 1 - |x - y|\}$. Write $f(\omega)$ for the payoff gain/loss from implementing action 1 at ω rather than letting the Agent choose randomly,

$$f(\omega) = \pi(1, \omega) - \left(\frac{1}{2}\pi(1, \omega) + \frac{1}{2}\pi(-1, \omega) \right).$$

We assume that $f(\omega)$ is continuous with $f(0) = f(1)$. We extend f to a periodic function on \mathbb{R} by setting $f(x) = f(x) - f(\lfloor x \rfloor)$ to make some calculations easier. To rule out a trivial case where rule-making is extremely easy or irrelevant, we also assume that f takes on both positive and negative values. Nothing changes if the payoffs are shifted up or down for each ω , so we will also normalize payoffs by assuming that $\pi(-1, \omega) + \pi(1, \omega) = 0$ for all ω . With this assumption, we have $V = \delta^t \sum_{t=1}^T a_t f(\omega_t)$.

Before proceeding to the two-period model let us illustrate why our assumptions about behavior and describability are important to rule out schemes whereby the first-best can be approximated through the use of sophisticated strategies in the first period, thereby rendering the entire problem moot. If we simply assumed that the agent has the same preferences as the principal then the first-best can be arbitrarily closely approximated by the principal effectively communicating the function f . Suppose that in period 1 the principal sends a real number which consists of a binary string (zeroes and ones) which specifies the number of times f crosses zero, followed by the number 2 which is interpreted as a “stop code” indicating the next component. That component is a number which specifies how many digits the location of the crossings will have in what follows, this is followed by a stop code. Then, where each of the crossings is located is revealed by a real number, each added on to the string. Finally a single zero or one indicates whether in the first region action 1 or -1 is optimal. This is a single real number which encodes the entire function f , and it is a feasible strategy for the principal. The interpretation this is given in equilibrium is not a rule, but as revealing the function f . Given that the agent has the same preferences as the principal, the first best is achieved from the second period onward.

This highlights the strength entailed in common knowledge of equilibrium strategies, and we rule out such schemes (which are both unrealistic and render the problem uninteresting) by our behavior and describability assumptions.

3 Two Periods

In this section we discuss the two-period version of our model. The two-period version is simple and brings out some basic insights.

In the first period of our model, there is nothing for the principal to do and her expected

payoff is always zero. In the second period, the principal can take advantage of the common reference point ω_1 to define a rule R_2 that will result in a higher expected payoff. All three versions of our model are identical: the principal chooses a single rule (ω_1, d, a) .

Our first Proposition notes that the principal is able to increase her payoff except in one special case and that rules are designed to have “excess breadth” in the sense that they are intentionally chosen to produce some incorrect decisions.

Proposition 1 *Consider the two-period version of the models under any of the overwriting assumptions.*

1. *If f is not antisymmetric around ω_1 then all optimal rules have $d^* > 0$ and the principal’s expected second period payoff is positive, $E(V|\omega_1) > 0$.*
2. *Any interior optimal choice d^* is such that $f(\omega_1 - d^*) = -f(\omega_1 + d^*)$.*

Proof

The principal’s expected second period payoff with rule (ω_1, d, a) is

$$E(V|\omega_1) = a \int_{\omega_1-d}^{\omega_1+d} f(x)dx.$$

This is a continuous function of a and d and the parameters are chosen from a compact set so it achieves its maximum. The maximum is zero only if $\int_{\omega_1-d}^{\omega_1+d} f(x)dx = 0$ for all d , which implies that f is antisymmetric around ω_1 : $f(\omega_1 - d) = -f(\omega_1 + d)$ for all d .

The second period payoff function is differentiable, so any interior optimum has

$$\left. \frac{d}{dd} \int_{\omega_1-d}^{\omega_1+d} f(x)dx \right|_{d=d^*} = 0$$

which implies $f(x - d^*) = -f(x + d^*)$.

QED

We will say that a rule r has *excess breadth* if there exists a state x for which $|x - w(r)| < d(r)$ and $a(r)f(x) < 0$. The second part of Proposition 1 says that rules are expanded until the average value of the rule in the marginal cases, $\omega_1 - d^*$ and $\omega_1 + d^*$, is zero. Unless both points are cases of indifference, this implies that the rule must be leading to incorrect decisions.

Corollary 1 *If there does not exist a value d for which $f(\omega - d) = f(\omega + d) = 0$, then the optimal rule at ω has excess breadth.*

Figure 2 provides a simple illustration. The left panel presents a simple payoff function. Action 1 is optimal whenever $\omega \in (0, 1/2)$. If the first period state is $\omega_1 = 3/8$, then the principal will instruct the agent to choose action 1 if a nearby state arises in period 2. She will not, however, choose a rule that is always correct by picking $d = 1/8$. At this breadth, she benefits from broadening the rule because the loss from choosing action 1 when $\omega \approx \frac{1}{2}$ is much smaller than the gain from choosing action 1 when $\omega \approx \frac{1}{4}$. In the two-period model the principal will broaden the rule until it is completely worthless at the margin. In the illustrated case, the optimal choice is $d^* = \frac{1}{4}$. At this breadth, the gain from implementing action 1 at $\omega_1 - d^* = \frac{1}{8}$ matches the loss from implementing action 1 at $\omega_1 + d^* = \frac{5}{8}$.




Figure 2: A two period example

The right panel graphs the expected second-period payoffs as a function of the first period state. The principal is best off if $\omega_1 = \frac{1}{4}$ or $\omega_1 = \frac{3}{4}$. In these cases, the principal can issue a rule which is applicable half of the time and which always yields the correct decision when it is applicable. In intermediate cases, the principal cannot maintain both the breadth of this rule and its accuracy and ends up with a lower payoff. For the function in the example, it turns out that the principal entirely sacrifices accuracy: the optimal d^* is exactly one-quarter for almost all ω_1 . The only exceptions are the two states where the

principal is worst off: $\omega_1 = 0$ and $\omega_1 = \frac{1}{2}$. These states are particularly bad examples on which to base distance-based rules because they right on the boundary between the regions where the two actions are optimal. In these states the principal's second period expected payoff is zero for any choice of d and a .

Observations

1. The fact that the first-period draw of ω_1 primarily effects the accuracy of the chosen rule and not its breadth is somewhat more general than the above example. Suppose that f is uniquely maximized at y , uniquely minimized at $y + \frac{1}{2}$, that f is strictly monotonic between the minimum and maximum, and f is symmetric both in the sense that $f(y - z) = f(y + z)$ for all z and in the sense that $f(x + \frac{1}{2}) = -f(x)$ for all x . Then, the principal's optimal choice will be $d^* = \frac{1}{4}$ for almost all ω_1 . An easy argument for this is that the first-order conditions imply that $f(\omega_1 - d^*) = -f(\omega_1 + d^*)$. Symmetry gives $f(\omega_1 - d^*) = f(\omega_1 + d^* - \frac{1}{2})$.³ The function f takes on each value only twice. Hence we either have that the two arguments are the same, $\omega_1 - d^* = \omega_1 + d^* - \frac{1}{2}$, or that the arguments are equidistant from the peak on opposite sides, $y - (\omega_1 - d^*) = \omega_1 + d^* - \frac{1}{2} - y$. The former gives $d^* = \frac{1}{2}$. The latter implies that $\omega_1 = y + \frac{1}{4}$, a non-generic case.
2. The principal would typically be able to achieve higher payoffs if we allowed her to issue multiple rules based on the same state. For example, given the function f pictured in Figure 1 she could achieve the first-best payoff in the second period if the first period state was $\omega_1 = \frac{1}{4}$ by issuing two rules: a lower-priority rule telling the agent to take action -1 if w_2 is within distance infinity of ω_1 , and a higher priority rule telling the agent to take action 1 if w_2 is within distance $\frac{1}{4}$ of ω_1 . In this case, the rules are perhaps not so unreasonable: they are like telling the agent to do action 1 in some cases and action -1 in all others. In other cases, however, we felt that such overlapping rules seemed to be less reasonable, e.g. the principal could dictate N different actions on concentric rings around w_1 by defining a first rule that applied

³For some values of ω_1 and d the term on the left will need to be $f(\omega_1 - d^* + 1)$ and/or the term on the right-hand side will need to be $f(\omega_1 + d^* - \frac{3}{2})$ or $f(\omega_1 + d^* + \frac{1}{2})$ ensure that the values are in $[0, 1]$.

at distances less than d_1 , a second higher-priority rule that applied at distances less than $d_2 < d_1$, a third even higher priority rule that that applied at distances less than d_3 , and so on. Such rules may be interesting additional category of rules to study, but we decided to limit our current paper to just three variants.

4 Three Periods

In this section, we discuss the three period version of our model. A primary observation is that the option value of being able to define superior rules in the future reduces some of the excess breadth we saw in the two-period case.

Write $r_2^m(\omega_1)$ for the myopic optimal rule at $t = 2$. This can be defined either as the optimal choice in the two-period model described in the previous section or as the optimal choice in the three-period model for a principal with discount factor $\delta = 0$. Let $d_2^m = d(r_2^m(\omega_1))$ be the breadth of the myopic optimal rule. We say that this rule has *excess breadth at the margin* if $\min(f(\omega_1 - d_2^m), f(\omega_1 + d_2^m)) < 0$. Write $r_2^*(\omega_1)$ for the optimal rule at $t = 2$ for an agent with discount factor $\delta > 0$.⁴

4.1 No overwriting

Proposition 2 *Suppose no overwriting is allowed. Let f be differentiable and suppose that the myopic optimal rule $r_2^m(\omega_1)$ is a unique interior optimum and has excess breadth at the margin. Then, the optimal rule at $t = 2$ in the three-period model has $d(r_2^*(\omega_1)) < d_2^m$.*

Proof

We show that the expected payoff from any rule $r'_2, r'_3(\omega_2)$ with $d'_2 \equiv d(r'_2) \geq d_2^m$ is strictly lower than the expected payoff that is obtained from a strategy we'll specify that includes setting $r_2 = (\omega_1, d_2^m - \epsilon, a(r_2^m(\omega_1)))$ for a small $\epsilon > 0$.

First, consider any strategy $r'_2, r'_3(\omega_2)$ with $d(r'_2) > d_2^m$. We show that such strategies cannot be optimal by showing that the expected payoff is improved by switching to the strategy $r_2^m, r'_3(\omega_2)$. To see this, note first that r_2^m provides a strictly greater payoff in

⁴This is obviously a function of δ . We omit the dependence from the notation where it will not cause confusion.

the second period. In the third period, the payoff is again strictly greater if $\omega_3 \in [\omega_1 - d(r'_2), \omega_1 + d(r'_2)]$ by the optimality of r_2^m and is identical if $\omega_3 \notin [\omega_1 - d(r'_2), \omega_1 + d(r'_2)]$.

Next, consider any strategy $r'_2, r'_3(\omega_2)$ with $d(r'_2) = d_2^m$. We show that the expected payoff can be improved by switching to a strategy with a slightly smaller value of d_2 . To define this strategy, note first that the assumption that d_2^m is regular implies that $f(\omega_1 \pm d_2^m) \neq 0$. Assume WLOG that $f(\omega_1 + d_2^m) > 0$ and $f(\omega_1 - d_2^m) < 0$. Hence, we can choose $\eta > 0$ so that $f(\omega + x) > f(\omega + d_2^m)/2$ and $f(\omega - x) < f(\omega - d_2^m)/2$ for all $x \in [d_2^m, d_2^m + 3\eta]$. For $0 < \epsilon < \eta$ define $r_2(\epsilon) = (\omega_1, d_2^m - \epsilon, a(r_2^m(\omega_1)))$ and

$$r_3(\omega_2; \epsilon) = \begin{cases} r'_3(\omega_2) & \text{if } \omega_2 \notin [\omega_1 - d_2^m - \eta, \omega_1 + d_2^m + \eta] \\ (\omega_2, \|\omega_2 - (\omega_1 + d_2^m - \epsilon)\|, 1) & \text{if } \omega_2 \in [\omega_1 + d_2^m - \epsilon, \omega_1 + d_2^m + \eta] \\ (\omega_2, \|\omega_2 - (\omega_1 - d_2^m + \epsilon)\|, -1) & \text{if } \omega_2 \in [\omega_1 - d_2^m - \eta, \omega_1 - d_2^m + \epsilon] \end{cases}$$

In words, these strategies consist of narrowing the breadth of rule r_2 by ϵ and then taking advantage of the more narrow definition to choose an ϵ broader rule than would have been possible in these cases when the realization of ω_2 within η of the boundary of the initial rule. We show that $r_2(\epsilon), r_3(\omega_2; \epsilon)$ gives a higher payoff than $r'_2, r'_3(\omega_2)$ when ϵ is small.

To see this note first that the disadvantage of $r_2(\epsilon)$ in the second period is $O(\epsilon^2)$ because payoffs are lower if $\omega_2 \in I \equiv [\omega_1 - d_2^m, \omega_2 - d_2^m + \epsilon] \cup [\omega_1 + d_2^m - \epsilon, \omega_2 + d_2^m]$. In the third period payoffs differ in two cases: if $\omega_3 \in I$; and if ω_2 is in the interval where the third-period rule is different and $\omega_3 \notin I$. The former difference again gives a loss that is $O(\epsilon^2)$. The latter case gives a larger advantage. In particular, if $\omega_2 \in [\omega_1 - d_2^m - \eta, \omega_1 - d_2^m + \epsilon]$, then the correct decision is made whenever $\omega_3 \in [\omega_2 - z - \epsilon, \omega_2 - z]$ for $z = \|\omega_2 - (\omega_1 - d_2^m)\|$, whereas previously no rule was defined in this interval. The gain from this in expected value terms is at least $\epsilon f(\omega_1 - d_2^m)/2$. There is a similar $O(\epsilon)$ gain from getting the correct decision when $\omega_2 \in [\omega_1 + d_2^m - \epsilon, \omega_1 + d_2^m + \eta]$ and ω_3 is in $[\omega_2 + z, \omega_2 + z + \epsilon]$ for $z = \|\omega_2 - (\omega_1 + d_2^m)\|$. Hence, for sufficiently small ϵ the strategy $r_2(\epsilon), r_3(\omega_2; \epsilon)$ gives a higher payoff.

QED

Figure 3 provides an illustration. The left panel graphs the same function f pictured in Figure 2. The right panel graphs the breadth of the myopic optimal rule, $d(r_2^m(\omega_1))$, and the breadth of the optimal second-period rule in the three-period model, $d(r_2^*(\omega_1))$.⁵

⁵The discount factor of $\delta = 1$ was used for these graphs.

Recall that the myopic optimal rule had a distance of $\frac{1}{4}$ for almost all ω_1 . This is shown as a dotted line in the graph. In the three period model, this breadth is never optimal. By Proposition 2 the optimal second-period rule is narrower. The solid line in the figure gives this breadth.




Figure 3: A three period example

An interesting observation from the graph is that the optimal second-period rule in the three-period no overwriting model can have *insufficient breadth at the margin*: $f(\omega_1 - d_2^*)$ and $f(\omega_1 + d_2^*)$ are both of the same sign as a_2^* .

Proposition 3 *In the three-period no overwriting model with any f , there exist states ω_1 for which the optimal second-period rule has insufficient breadth at the margin.*

The right panel of Figure 3 illustrates that for that payoff function the the optimal breadth is exactly zero when ω_1 close to zero, one-half, and one. A rule with zero breadth will satisfy the insufficient breadth at the margin condition. In the example, the optimal breadth is zero because the option value to keeping rules undefined outweighs the potential short-run benefit to making a rule. For example, suppose $\omega_1 = \frac{1}{2} + \epsilon$. In this case, a second-period gain of approximately ϵ could be obtained by making a rule that action -1 should be chosen if $\omega \in [\frac{1}{4} + \epsilon, \frac{3}{4} + \epsilon]$, or a smaller gain of $\frac{1}{2}\epsilon^2$ could be obtained by defining a much narrower rule that action -1 should be chosen on $[\frac{1}{2}, \frac{1}{2} + 2\epsilon]$. However, in each case,

such a rule can prevent a much better rule from being defined in the third period. For example, if $\omega_2 = \frac{3}{8}$, the former would prevent us from defining any rule at all, and the latter would force us to limit the domain of the rule to $[\frac{1}{4}, \frac{1}{2}]$ instead of $[\frac{1}{8}, \frac{5}{8}]$.

The statement in the Proposition that a similar result holds for any f means for any f satisfying the assumptions in Section 2 we maintain throughout the paper, including that f is continuous and takes on both positive and negative values. Given these assumptions, one can always find ω_1 where the optimal breadth is similarly zero even though f is positive by considering points very close to points where f crosses zero. We omit a formal demonstration.

4.2 Incremental overwriting

The above discussion has focused on the no-overwriting version of our model. In the incremental overwriting model, the constraints that created the advantage to making the rule somewhat narrower in the above example do not exist: one can overwrite the interval $[\frac{1}{2}, \frac{5}{8}]$ at $t = 3$ if the draw of ω_2 makes this attractive. The incremental overwriting model does have some constraints, however, and these can still provide an incentive to choose a narrower second-period rule than would be chosen by a myopic agent.

Proposition 4 *In the three period incremental overwriting model there exist payoff functions f and states ω_1 for which $d(r_2^*(\omega_1)) < d(r_2^m(\omega_1))$.*

Proof

We prove this by providing an example. Suppose $f(x) = 1$ if $x \in [\epsilon, \frac{1}{3} - \epsilon]$ and $f(x) = -1$ otherwise.⁶ For $\omega_1 = \frac{1}{6}$ there are the myopically optimal rule is to take action -1 everywhere: this gives a payoff of $\frac{1}{3} + 4\epsilon$, whereas the more obvious narrow rule dictating that action 1 be taken on $[\epsilon, \frac{1}{3} - \epsilon]$ gives a payoff of $\frac{1}{3} - 2\epsilon$.

In the three-period incremental overwriting model, however, the narrower rule provides a strictly greater expected payoff if ϵ is small and δ is close to one. To see this, note that the narrower rule can be improved in the third period whenever $\omega_2 \notin [\epsilon, \frac{1}{3} - \epsilon]$. When

⁶The example uses a discontinuous f to make the computations easy. A continuous counterexample could be obtained by choosing a nearby continuous function.

such ω_2 arise, the optimal incremental rule is to prescribe action -1 on the largest interval around ω_2 contained in $[\frac{1}{3} - \epsilon, 1 + \epsilon]$. On average this interval has width $\frac{1}{3} + \epsilon$, so the expected gain is $(\frac{2}{3} + 2\epsilon)(\frac{1}{3} + \epsilon) \approx \frac{2}{9}$. The broad rule can be improved only if $\omega_2 \in [\epsilon, \frac{1}{3} - \epsilon]$. Again, the best way to improve the rule is to define r_3 to extend to the nearest edge of the interval $[\epsilon, \frac{1}{3} - \epsilon]$. There is a payoff gain of two whenever ω_3 is such that the correction is effective, but the expected gain is still just $2(\frac{1}{3} - 2\epsilon)(\frac{1}{6} - \epsilon) \approx \frac{1}{9}$. The nontrivial difference in third period expected payoffs will easily outweigh the ϵ order second-period differences if δ is close to one and ϵ is not too large.

QED

An intuition for why the narrower rule is preferable in the example in the above proof is that a constraint imposed by the incremental overwriting model is that one cannot go back at $t = 3$ and redefine the rule involving ω_1 . Hence, there is a benefit to immediately employing ω_1 in its best long-run role. In this case, the natural use for $\omega_1 = \frac{1}{6}$ is as the leading example of the set of situations where action 1 is optimal.

4.3 Vast overwriting

In the vast overwriting model there is no incentive to leave rules undefined. The long-run optimal strategy is always to define each period's rulebook to maximize the payoff in that period. Hence $r_2^* = r_2^m$.

5 Infinite Horizon

In a model with more periods a patient principal will be able to exploit the large number of available examples that arise to achieve higher payoffs. In this section we present a couple of results illustrating the sense in which optimal rules will be developed by a patient principal.

The first-best rulebook would specify that $a_t = a^{FB}(\omega_t) \equiv \text{sign}(f(\omega_t))$. Write $V^{FB} = \int_0^1 |f(x)| dx$ for the per-period payoff that this rule would give. This is an upper bound on the expected payoff that a rational player can achieve in any period.

5.1 Vast overwriting

Developing rules is easiest in the vast-overwriting version of our model. Here, it is fairly easy to see that a fully optimal rule book is developed in finite time.

Proposition 5 *Consider the vast-overwriting version of our model. Suppose that $f(x)$ crosses zero a finite number of times. Then, with probability one there exists a T such that the action $a^{FB}(\omega_t)$ is chosen in period t for all $t \geq T$. As $\delta \rightarrow 1$ the principal's average per-period payoff converges to V^{FB} .*

Proof

The second conclusion follows directly from the first. To see the first, suppose that $f(x)$ crosses zero $n - 1$ times and let $0 = x_0 < x_1 < \dots < x_n = 1$ be such that $f(x)$ is always nonnegative or always nonpositive on each interval $S_j = [x_{j-1}, x_j]$. Write a^j for the optimal action on S_j . Let $S_{j1} = [x_{j-1} + \frac{1}{4}(x_j - x_{j-1}), x_{j-1} + \frac{1}{2}(x_j - x_{j-1})]$ and $S_{j2} = [x_j - \frac{1}{2}(x_j - x_{j-1}), x_j - \frac{1}{4}(x_j - x_{j-1})]$. With probability one there exists a T such that $\{\omega_1, \omega_2, \dots, \omega_T\} \cap S_{jk} \neq \emptyset$ for all $j = 1, 2, \dots, n$ and for $k = 1, 2$. For any $t > T$ we can define $\tau(j, k)$ so that $\omega_{\tau(j, k)} \in S_{jk}$. Define a rulebook for period t by

$$R_t = \bigcup_{j=1}^n \{(\omega_{\tau(j,1)}, \omega_{\tau(j,1)} - x_{j-1}, a^j, j), (\omega_{\tau(j,2)}, x_j - \omega_{\tau(j,2)}, a^j, -j)\}$$

Note that one or more of the rules applies to each $\omega \in [0, 1]$ and that for any $\omega \in S_j$ any rule that applies to ω specifies action a^j .⁷ Hence, R_t will result in action $a^{FB}(\omega_t)$ being chosen.

QED

Remarks

1. The assumption that $f(x)$ crosses zero a finite number of times is necessary for the first part of the result. If $f(x)$ has an infinite number of crossings then no finite rule can specify the optimal action in between each of the crossings. Hence, the upper bound V^{FB} cannot be achieved in any period.

⁷When two rules apply they give the same action so the priorities we specified were unimportant.

2. The rulebook used in the proof employs rules with overlapping domains that specify the same action on the overlap. Some such overlapping is necessary to achieve the first-best in finite time. If compatible rules are not allowed to overlap, then on each interval S_j we will be able to define rules that extend exactly to the right endpoint or exactly to the left endpoint, but not both (unless the ω 's satisfy a relation that is satisfied with probability zero). We need rules to end exactly at $n + 1$ endpoints to achieve the first best. Hence, the first-best is not achievable.

5.2 No overwriting and incremental overwriting

The no overwriting and incremental overwriting versions of our model impose additional constraints on the principal and will reduce her payoffs. Nonetheless, it is fairly easy to show that payoffs still must approximate the first-best when the principal is patient.

Proposition 6 *Suppose that $f(x)$ crosses zero a finite number of times. In the no overwriting and incremental overwriting models the principal's average per-period payoff converges to V^{FB} in the limit as $\delta \rightarrow 1$.*

Proof

A suboptimal strategy for the principal which is feasible in the incremental overwriting model would be define no rule except in the $2n$ periods in which the current state falls into one of the subintervals S_{jk} defined in the proof of Proposition 5 for the first time, and in those to define the rule exactly as we do in that proof. If the principal follows this strategy she will achieve a payoff V^{FB} in all periods from some period T on. Hence, her average payoff will be something that approaches V^{FB} as $\delta \rightarrow 1$. Her optimal strategy yields a higher payoff, and hence must also approach V^{FB} .

The above strategy is not feasible in the no overwriting model because the principal is restricted from issuing overlapping rules even when they disagree. More complicated strategies can be used, however, to obtain a payoff that approximates the first best. One way to do this is to define no rule unless the state is within ϵ of the center of one of the S_j . There may be a long wait until such states occur, but once they have occurred in each S_j the expected payoff in all future periods is at least $1 - 2n\epsilon$. This similarly puts a lower bound on the payoff that a rational player can receive.

QED

Remark

1. The assumption that $f(x)$ has a finite number of crossings is convenient for this proof, but is not necessary. If $f(x)$ crosses zero a countable number of times, we can implement a similar strategy in each subinterval. Not all subintervals will be covered, but each half of a subinterval of width w is covered with probability $(1 - \frac{w}{2})^{t-1}$ in period t . Hence, the expected payoff again converges to the first best.

6 Delegation Structures

So far our environment has been sufficiently simple that there has been a single possible delegation structure for the authority relationship/firm: a single principal, a single agent and a single task. We now consider a richer environment where the principal has a single unit of time, but may do two things: perform a task herself, or communicate about it to an agent. We assume that performing a task is more time consuming than communicating about it. In particular, we assume that performing a task takes one unit of time, but communicating to an agent takes only $1/2$ a unit of time. The principal has two agents available.

More concretely, there are two tasks $i = 1, 2$. On each task the principal's payoff on task i is $\pi_i(a_i, \omega_i)$ that depends on the action taken and the state of nature. Suppose that the principal's benefit function for the two tasks is given by $B(\pi_1(a_1, \omega_1), \pi_2(a_2, \omega_2)) = \min(\pi_1(a_1, \omega_1), \pi_2(a_2, \omega_2)) + (1 - \gamma)\max(\pi_1(a_1, \omega_1), \pi_2(a_2, \omega_2))$ and the "f functions" are f_1 and f_2 respectively. Note that γ parameterizes the degree of complementarity of the tasks, with $\gamma = 1$ being strict complementarity. As before, the principal's payoff in the full game is the discounted sum of his per period payoffs and is given by $V = \sum_{t=1}^T \delta^t B(\pi_1(a_1, \omega_1), \pi_2(a_2, \omega_2))$. There are three possible delegation structures in this setting: (i) the principal communicates with each agent after observing ω_t and then task 1 is performed by agent 1 and task 2 by agent 2 in period $t + 1$, (ii) task 1 is performed by an agent and task 2 is performed by the principal, but without communication from

the principal, and (iii) task 2 is performed by an agent and task 1 is performed by the principal, but without communication from the principal. We call these *full delegation*, *type 1 delegation* and *type 2 delegation* respectively.

We can now offer a number of results about the relative attractiveness of different delegation structures.

Proposition 7 *Consider any of the overwriting assumptions, suppose that $f(x)$ crosses zero a finite number of times, and that delegation strategies are time invariant. Then if δ is sufficiently large then full delegation is optimal.*

Proof

From Propositions 5 and 6 we know that as $\delta \rightarrow 1$ the principal's average per-period payoff converges to V^{FB} . Neither type 1 or type 2 delegation can achieve this since there is no communication on one of the tasks and the expected payoff on that task is zero which is bounded away from V^{FB} , and the payoff on the other task is V^{FB} .

QED

The next result shows that if the principal starts out with full delegation then he will switch to type 1 or type 2 delegation if he gets enough “lucky” realizations of the state which allow him to establish very effective rules on one task.

Proposition 8 *Consider the vast overwriting version of the model and suppose that $f(x)$ crosses zero a finite number of times. Then with probability one there exists a T such that type i delegation is optimal in all future periods $t > T$.*

Proof

From Proposition 5 we know that with probability one there exists a T such that for all $t \geq T$, $a^{FB}(\omega_t)$ is chosen. If this happens on task 1 (respectively task 2) then type 1 delegation (respectively type 2 delegation) is optimal since $a^{FB}(\omega_t)$ is chosen on both tasks.

QED

Remarks

1. The proposition says that type i delegation is weakly optimal, but it will be strictly optimal except in the special case where the first-best rulebook is established on both tasks in exactly the same period.
2. It seems intuitive that if a highly effective rulebook is developed on one task relative to that on the other task then it would be optimal for the principal to delegate the task with the effective rulebook and perform the other himself. However, it is unclear to us how to formalize this intuition at present.

While Proposition 7 showed that full delegation can be optimal if the principal is sufficiently patient, we now show that if she is sufficiently impatient but the tasks are strongly complementary then full delegation is optimal for some amount of time.

Proposition 9 *Consider any of the overwriting regimes. Then if δ is sufficiently small and γ is sufficiently close to one then there exists $T \geq 1$ such that full delegation is optimal for all $t = 1, \dots, T$*

Proof

At $t = 1$ the expected payoff from under full delegation is zero and no other governance structure improves on this since the expected payoff on the delegated task is zero and γ is close to 1 by assumption. At $t = 2$ the expected payoff at $t = 1$ on each task under full delegation is positive, but under the other governance structures one task still has expected $t = 2$ payoff of zero. Therefore for γ sufficiently close to 1 full delegation is optimal.

QED

A final observation of this section is that the path dependence which the basic model exhibited can be magnified in this richer setting. There there was simply a direct effect of history on payoffs. Here there is an additional indirect effect whereby a change in delegation structure can further improve the principal's payoff if he is able to switch to a more effective delegation structure.

7 Endogenous Categories

As we mentioned in the introduction, our model suggests a mechanism via which categories are formed in models of category based thinking (such as Mullainathan (2002) or Fryer and Jackson (2007)). Roughly speaking, these models depart from the Bayesian framework by grouping together multiple states into a single category. In Fryer and Jackson the decision maker allocates experiences to a finite and limited number of categories—and they focus on the optimal way to categorize (given the number categories available). They show that experiences which occur less frequently are (optimally) placed in categories which are more heterogeneous. Mullainathan (2002) considers decision makers who have coarser beliefs than pure Bayesians, since they have a set of categories which are a partition of the space of posteriors and choose just one category given data which they observe. They then forecast using the probability distribution for the category.

An important question which these models leave open is where exactly these categories come from. In Mullainathan (2002) they are completely exogenous. In Fryer and Jackson (2007) they are formed optimally, although the number of categories is exogenous. Our model provides one way to think about how categories are formed. In our setting we think of the categories as being the sets $\{\omega s.t. \|\omega - \omega_t\| < d_t\}$ on which individual rules are defined. Thus in an N period model if action 1 is taken on several disjoint set we think of them as being different categories⁸. The analog of perfect categorization is that all states are covered by a rule (with highest precedence) which maximizes the principal's payoff. Imperfect categorization comes from one of two sources: (i) rules which do not maximize the principal's payoff in a given state or, (ii) states which are not covered by a rule.

Fryer and Jackson (2007) offer a labor market example with high and low ability workers who may be black or white. Thus there are four types: high-black, low-black, high-white, low-white. They assume that the respective population shares of the types are: 5%, 5%, 45%, 45%. In the context of their model they show that if there are only three categories then high and low ability blacks are categorized together, whereas high-white and low-white each have their own category.

⁸In the case of overlap a state can belong to multiple categories.

In the context of our model, suppose that action 1 is the decision to hire a worker and action -1 is the decision not to hire. The left panel of figure 4 depicts an arrangement of workers which (from left to right) is: high-white, low-white, low-black, high-black. The dotted vertical line represents the division between whites and blacks. Now suppose that there are two periods and the first period realization of the state is ω_1 . The optimal rule categorizes all workers in the band between the first and third dots together and the rule is that any worker in that category should be hired. In a sense this rule is quite “fair”, though not perfectly so. Most high ability blacks are categorized as to-be-hired, though not all, and a few low ability whites are also put in that category.




Figure 4: Categorization examples

In the second panel of figure 4 outcomes can be much less fair. The difference between the two panels is the frequency of the function. This captures the idea that high and low ability types are harder to distinguish than they are in the left panel—they are less naturally clumped together. Not only can this lead to a less effective rule being developed, it can lead to a less fair one. If the realization of the state is ω_1 then the optimal rule categorizes all workers between the first and third dots as not-to-be-hired. More than half the high-ability blacks are so categorized, and this comes from the principal’s desire to expand the rule on the left of ω_1 so as to prevent many low-whites from being hired. The key driver of the unfairness is that whites are relatively more populous than blacks. This is a similar underlying cause as in Fryer and Jackson, though the models are obviously different. In a

version of our model with more than two periods both the number of categories and the nature of the categorization arise endogenously.

8 Discussion and Conclusion

The model we have analyzed sheds light on a number of phenomena. The first is path dependence, whereby there may be differences in performance based purely on the history of events. In the model, a series of “lucky” early realizations of the state can lead to a very effective rule book being established, delivering highly efficient outcomes. Conversely, bad draws early on can have persistent effects which cannot be overcome. There is a sizeable empirical literature which documents what Gibbons (2006) calls “persistent performance differences among seemingly similar organizations”: that is, substantial differences in the performance of organizations which are difficult to account for (Mairesse and Griliches (1990), McGahan (1999), Chew et al. (1990)).

One potential theoretical explanation for this is that differences in performance correspond to different equilibria of a repeated game (Leibenstein (1987)). As Gibbons (2006) points out, such an explanation provides a basis for the existence of different equilibria, but is silent about the dynamic path of differences, and how one might switch equilibria. In contrast, our model exhibits path dependence. It can be the case that a firm has more efficient communication than another purely because of chance: they got some good draws early on which allowed them to develop an effective rule book.

Furthermore, the model with multiple activities highlights that there can be a complicated evolution of governance structures because of communication needs. Various activities may or may not be delegated, and this can change over time as more effective rules are developed. At a minimum this suggests that bounded communication is a potential explanation for the heterogeneity of observed organizational forms.

The model also provides a potential explanation for *learning curves*, whereby a firm’s cost declines with cumulative output/experience. The fact that learning can take place at different rates depending on history or different abilities to overwrite is consistent with the evidence on heterogeneity of rates of learning (Argote and Epple (1990)).

Another empirical regularity which has been documented is that firms often “start-

small” in the sense that they grow less quickly than they otherwise would. On its face this is puzzling, as profits are being foregone in the process. Watson (2002) argues that a rationale for starting small is to build cooperation in a low stakes environment and then enjoy the benefits of the cooperative equilibrium in a high stakes environment. Our model offers a different (and potentially complementary) notion of starting small. There is option value in developing rules and consequently rules can be under-inclusive, particularly early on (recall Proposition 2). One way to interpret the actions in our model that would make it connect with the starting-small idea would be to regard the zero expected payoff that obtains in states for which no rule has been defined as coming from a blanket instruction that the firm should decline to decline any business opportunities that might arise which do not fall into any of the categories for which rules have been defined. With this interpretation, our model would be one in which the size of the business grows over time as new situations arise and enable the firm to promulgate rules about how it should exploit such opportunities.

The possibility that communication may start out vague and become more precise over time has the flavor of what Gibbons (2006) refers to as “building a routine”. The model also highlights that the nature and extent of inefficiency from bounded communication—and the degree to which rules are over or under inclusive—can depend importantly on the ability of the principal to amend previously developed rules. Indeed we saw that in the vast overwriting regime there was no option value from leaving rules undefined as they could always be amended.

There are two obvious directions for further work. The first is to provide a more concrete foundation for the inability of the principal to describe the state to the agent. We remarked early on that we had in mind state spaces which were complex in some sense, though we conducted our analysis on the unit interval. It would be useful to understand what kinds of environments, or what other cognitive limitations, give rise to the inability to perfectly communicate.

Finally, the second direction would be to more explicitly model the tradeoff between the price mechanism and the authority relationship. This would allow one to address the question of when the price mechanism is superior at adapting to changing circumstances than the authority relationship: Barnard versus Hayek. What we have done is to highlight

one particular cost of the authority relationship, and analyzed how different governance structures perform in that context. We remarked earlier that a complete understanding of the tradeoff between the efficiency of markets and firms in changing circumstances would seem to require a general model of price formation. That is, a model in which one can analyze how frictions affect the incorporation of information into the price mechanism and hence the circumstances in which the price mechanism works well and those where it does not. Although the literature on rational expectations equilibrium and securities markets provided some steps in this direction⁹, we are not aware of a suitable model for studying this question. For that model would not only need to address the question of how information is incorporated into prices, but also be able to rank economies according to the allocative efficiency effects of imperfect prices. Even the first question is daunting enough. Any model of price formation with rational expectations is subject to the Milgrom-Stokey No-Trade Theorem (Milgrom and Stokey (1982)), raising the awkward question of how information comes to be incorporated into prices if those who collect the information cannot benefit from it. Nonetheless, a complete model of price formation seems essential to satisfactorily answer the larger question of when markets or firms adapt better to changing circumstances.

⁹For example: Grossman (1976), Allen (1981), Glosten and Milgrom (1985), Kyle (1985)

References

- Beth Allen. Generic existence of completely revealing equilibria for economies with uncertainty when prices convey information. *Econometrica*, 49:1173–1199, 1981.
- Linda Argote and Dennis Epple. Learning curves in manufacturing. *Science*, 247:920–924, 1990.
- Kenneth J. Arrow. *The Limits of Organization*. Norton, New York, NY, 1974.
- Kenneth J. Arrow. On the agenda of organizations. In *Collected Papers of Kenneth J. Arrow*, volume 4, chapter 13, pages 167–184. Belknap Press, 1984.
- George Baker, Robert Gibbons, and Kevin J. Murphy. Governing adaptation: Decision rights, payoff rights, and relationships in firms, contracts, and other governance structures. *Mimeo*, 2006.
- Chester Barnard. *The Functions of the Executive*. Harvard University Press, Cambridge, MA, 1938.
- Bruce Chew, Kim Clark, and Timothy Bresnahan. *Measurement, coordination and learning in a multiplant network*, chapter 5. Harvard Business School Press, 1990.
- Ronald Coase. The nature of the firm. *Economica*, 4:396–405, 1937.
- Jacques Cremer, Luis Garicano, and Andrea Prat. Language and the theory of the firm. *Quarterly Journal of Economics*, CXXII:373–408, 2007.
- Roland G. Fryer and Matthew O. Jackson. A categorical model of cognition and biased decision-making. *Contributions in Theoretical Economics*, B.E. Press, forthcoming, 2007.
- Robert Gibbons. What the folk theorem doesn't tell us. *Industrial and Corporate Change*, 15:381–386, 2006.
- Lawrence R. Glosten and Paul R. Milgrom. Bid, ask and transactions prices in a specialist model with heterogeneously informed traders. *Journal of Financial Economics*, 14:71–100, 1985.

- Sanford Grossman and Oliver Hart. The costs and benefits of ownership: A theory of vertical and lateral ownership. *Journal of Political Economy*, 94:691–719, 1986.
- Sanford J Grossman. On the efficiency of competitive stock markets where traders have diverse information. *Journal of Finance*, pages 573–585, 1976.
- Oliver Hart and John Moore. Property rights and the nature of the firm. *Journal of Political Economy*, 98:1119–1158, 1990.
- F.A. Hayek. The use of knowledge in society. *American Economic Review*, 35:519–530, 1945.
- Benjamin Klein, Robert Crawford, and Armen Alchian. Vertical integration, appropriable rents and the competitive contracting process. *Journal of Law and Economics*, 21:297–326, 1978.
- Frank H. Knight. *Risk, Uncertainty, and Profit*. Hart, Schaffner & Marx; Houghton Mifflin, Boston, MA, 1921.
- Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53:1315–1335, 1985.
- Harvey Leibenstein. *Inside the Firm: The Inefficiencies of Hierarchy*. Harvard University Press, 1987.
- Jacques Mairesse and Zvi Griliches. Heterogeneity in panel data: are there stable production functions. In P. Champsaur et al., editor, *Essays in Honor of Edmond Malinvaud*, volume 3: Empirical Economics. MIT Press, Cambridge MA, 1990.
- Jacob Marschak and Roy Radner. *Economic Theory of Teams*. Cowles Foundation and Yale University, 1972.
- Anita M. McGahan. The performance of us corporations: 1981-1994. *Journal of Industrial Economics*, XLVII:373–398, 1999.

- Paul Milgrom and Nancy Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 26:17–27, 1982.
- Sendhil Mullainathan. Thinking through categories. *NBER working paper*, 2002.
- Sendhil Mullainathan, Joshua Schwartzstein, and Andrei Shleifer. Coarse thinking and persuasion. *Quarterly Journal of Economics*, forthcoming, 2008.
- Herbert Simon. A formal theory model of the employment relationship. *Econometrica*, 19: 293–305, 1951.
- Joel Watson. Starting small and commitment. *Games and Economic Behavior*, 38:176–199, 2002.
- Oliver Williamson. Transaction cost economics: The governance of contractual relationships. *Journal of Law and Economics*, 22:233–261, 1979.
- Oliver E. Williamson. The vertical integration of production: Market failure considerations. *American Economic Review*, 61:112–123, 1971.
- Oliver E. Williamson. The economics of governance. *American Economic Review Papers and Proceedings*, 2:1–18, 2005.