

A Theory of Rule Development

Glenn Ellison and Richard Holden*

January 2012

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: School of Economics, Australian School of Business, University of New South Wales. E-mail: richard.holden@unsw.edu. This work was supported by MIT Sloan's Program on Innovation in Markets and Organizations and NSF grant SES-0550897. We are deeply indebted to Bob Gibbons for numerous discussions and suggestions. We also thank the editor Wouter Dessein, two anonymous referees, Oliver Hart, Bengt Holmstrom, Andrei Shleifer, Andy Skrzypacz, Oliver Williamson, and seminar participants at Columbia, Chicago GSB, Harvard, MIT, UC Berkeley and UCLA 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 rules will be inefficient is that the world is so complex that it would be impossible to list what is optimal in every contingency. In this paper, we develop a model to explore the economics of second-best rulemaking in such an environment. Our model does not have a traditional “agency problem” involving incentives and private benefits: we assume that agents follow any rules that are set out for them. Instead, the problem that our principal faces is one of communication: the language in which rules must be expressed is not sufficiently rich as to allow the first-best contingent plan to be described. The way in which we do this is the most novel aspect of our framework: we assume that the principal can convey analogy-based rules dictating that a given action be taken in states similar to previously-realized states. We examine the economics of second-best rulemaking and discuss the inefficiencies that arise and how they are related to various assumptions about the environment.

One application we have in mind for our model is rule-making in organizations. Here, our motivation is that many businesses have central decisionmakers who have developed a deep understanding of what is optimal for the firm. For example, an experienced venture capitalist knows a great deal about which ventures should be financed, the central management of a retail chain knows a great deal about how a store should operate, and an experienced human relations director has expertise in judging whether a given job applicant is likely to be a valuable employee. Central decisionmakers are also very busy and hence delegate tasks to agents: younger associates do background research on venture capital proposals; store managers are hired to run individual stores; and less experienced staff review many job applications. Some things that the central decisionmakers know can easily be written down in a manual, e.g. “always include in your report a complete list of other ventures that the entrepreneur has been involved with in the past ten years” and “always issue refunds to customers who have a receipt for a purchase within the past 30 days.” But there will be a great deal of residual expertise that cannot be so easily codified: one cannot describe the full set of proposals that might come in and what issues should

be investigated on each one; or the full set of people who might apply for a job and what questions you should ask each one; or all customer complaints that a store manager might encounter. Our primary focus is on how organizations develop less formal rules to capture the benefits of this residual expertise.

We feel that an important way in which nuanced rules are conveyed in the real-world is via analogy-based rules. For example, the VC analyst might be told later on that the firm might have made a better decision if his report on ENTREPRENEUR X's proposal had contained INFORMATION Y and he should include such information on any similar proposals; the store manager could be told that MAYOR X should have immediately been given a refund even though he didn't have a receipt and the same rule applies to anyone at all like the mayor; and an HR employee could be told that hiring PERSON X seems to have been a mistake and he should be wary of very similar candidates. Each new shared experience provides an opportunity to promulgate such rules. There is some scope to decide how broad or narrow to make each such rule. The efficiency of rule systems should improve over time as shared experience accumulates.

A second application we have in mind is to the interpretations of laws by common-law courts. Here, the principal would be the Supreme Court or some other higher-level court and the agents would be the lower-level courts that adjudicate each case. Much of what a law dictates is readily apparent from the text of the statute. But again there will be "residual expertise": there will be cases for which the interpretation of the statute is not as clear and the higher level court will have views about how each such case should be decided. The way in which this expertise gets reflected in lower court decisions is by legal precedent: the higher-level court agrees to hear some cases and its decisions in those cases become rules which lower-level courts will apply to similar cases in the future. These precedents are like the analogies mentioned above in that they exploit shared experiences and in that higher-level courts will have some scope to indicate whether each new precedent is broad or narrow. The accumulated set of precedents becomes more complete and lower court decisions come close to the higher court's preferred interpretation as the case law accumulates.

We describe our model formally at the start of Section 2. 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 of our communication constraints: 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 any rule that has been communicated to him, but these will typically be incomplete and/or suboptimal because of our 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. One other consideration that turns out to be important is the extent to which rules can be changed over time. Different assumptions may be appropriate for different applications, so we discuss three variants of the model. 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 basic economic observation is that decisions about the breadth of a rule will reflect the rule's average value in the marginal cases. Hence, except in special cases optimal rules will have "excess breadth" in that they are designed to 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 observation here is that in dynamic settings option-value considerations create a "starting small" incentive that counteracts the incentive noted above: narrow rules have the benefit of giving the principal more freedom to exploit future rule-making opportunities. In the no overwriting case, the option-value consideration may dominate to the extent that the principal declines to expand the breadth of a rule even though it would also be correct in all marginal cases.

Section 5 examines the infinite-horizon case. While rule books may grow ever more complicated, we note that they can be seen as converging if one thinks about them in

terms of the expected payoff they produce. These limiting payoffs depend on historical accident: the firm achieves a better long-run payoff if early realizations of the state are ones that lead to efficient analogy-based rules. We suggest that this might be a source of persistent performance differences among seemingly similar firms. Firms will always develop nearly optimal rules if the principal is very patient. Whether exactly optimal rules develop depends 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. Most closely related is the literature on languages and communication 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. Wernerfelt (2004) analyzes a team-theoretic problem in which agents may use different codes in communicating with a principal in order to optimally deploy an indivisible asset. Cremer et al. (2007) 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.

Also closely related is the literature on categorization and coarse decision making. Papers in this literature endogenize the coarseness of plans and/or examine implications of this coarseness. MacLeod (2002) is perhaps most similar to our paper in that its model involves organizations making state-contingent plans that will be implemented when a random state arises. It endogenizes incompleteness via per-state planning costs and examines learning curves in a dynamic model. Other notable papers endogenizing the coarseness of categorizations include Samuelson (2001), which examines a model in which players make analogies across games to economize on cognitive costs which are modeled using automata,

and Al-Najjar and Pai (2009), which develops a model in which coarse categorical models can be preferred when more refined models will over-fit limited data. Another set of interesting recent papers, including Mullainathan (2002), Fryer and Jackson (2007), and Mullainathan et al. (2008), focuses more on the types of problems that arise when agents use a limited number of categories. Relative to this literature, one can think of our paper as developing an alternate framework for studying endogenous coarseness – restricting agents to using distance based rules with breadth as a choice variable – and examining problems motivated by organizational decisionmaking.

More broadly, there is a long history of studying adaptation in organizations dating back to Knight (1921), Barnard (1938), and Hayek (1945). Williamson (2005) argues that “adaptation to disturbances [is] the central problem of economic organization”. The largest part of the formal literature examines the principal-agent relationship from incentive theory perspective¹. Our paper, in contrast, is more similar to the “team theory” approach pioneered by Marschak and Radner (1972) and including such subsequent contributions as Cremer (1980), Bolton and Dewatripont (1994), Wernerfelt (1997), Van Zandt (1999), and Vayanos (2003). We follow this literature in focusing on models in which there are no divergence in preferences between principals and agents. Instead, the fundamental obstacles to efficient organizations are the costs of information processing and communication. Our focus is somewhat different from the much of this literature: whereas many papers in this literature assume a reduced form “cost” of communication and focus on how organizations should be structured to minimize these costs, we spend most of our time characterizing the “costs” that result from the limitations we place on what can and cannot be communicated; and we only get to organizational implications in one section near the end.

Our paper also relates to other literatures that endogenize communication problems in organizations. A large literature following Crawford and Sobel (1982) examines limitations on communication using cheap-talk models. Dessein (2002) notes that such limitations provide a motivation to delegate decisions to agents. Two more recent papers focus more directly on endogenizing communication costs. Dessein (2006) analyze the tradeoffs between

¹Aghion et al. (1991) study a model of optimal learning by experimentation with a single decision maker who is uncertain about his own payoff function.

adaptation to changing environments and coordination across agents when communication is imperfect. It emphasizes that communication is difficult and includes an extension analyzing the quality of communication in a model in which communication can be made more accurate at a cost. Dewatripont and Tirole (2005) focus more directly on endogenizing the limited quality of communication and develop several models featuring differences in preferences, effort costs in the sending and receiving of messages, and various assumptions about the structure of actions. What we are doing is somewhat different from this literature in that we do not have preference differences and focus on the issue of communicating the desired contingent plan-of-action, whereas much of this literature abstracts from this part of the communication problem and simply assumes, for example, that the optimal plan is to set the action “equal” to the signal the agent receives.

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. A complete 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.

Finally, there is an experimental literature on communication in organizations. Most

closely related to our paper are Camerer and Weber (2003) and Selten and Warglien (2006). Camerer and Weber (2003) run an experiment involving 16 pictures depicting offices in which subjects can use natural language to try to tell each other which picture is relevant in this round. They are interested in the development of codes over time and in how two “firms” with different codes perform when merged. Selten and Warglien (2006) use abstract shapes instead of office pictures, and the available language consists of a few different letters, each with its own cost. The performance-improvements on these tasks are an example of organizational “learning by doing”.

We would also cite one piece of potentially relevant empirical evidence. Kellogg (2008) examines the experience of oil production companies and drilling contractors and finds evidence that producer-driller pairs become more efficient over the course of their relationship. A source of these economies may be that shared experience lets the producer describe measures that should be undertaken in particular circumstances to avoid problems that have occurred in the past.

2 The Model

In this section we describe our model formally and discuss motivations for several of the assumptions.

2.1 Statement of the Problem

Our model has two players: the Principal and the Agent. The Principal hires the Agent 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 $\omega_t \in \Omega$ and chooses an action $a_t \in \{-1, 1\}$. The principal then receives a payoff $\pi(a_t, \omega_t)$ that depends on A’s action and the state of nature. The principal’s payoff in the full game is the discounted sum of her per period payoffs: $V = \sum_{t=1}^T \delta^t \pi(a_t, \omega_t)$. We will sometimes write $V_t \equiv \pi(a_t, \omega_t)$ for the principal’s payoff in period t .

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,

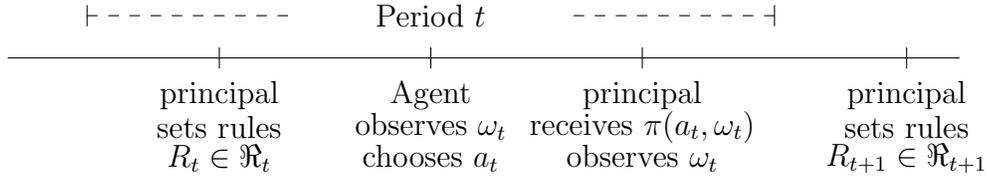


Figure 1: Model timing

however, will make this impossible.

A central element of our model is that we assume the existence of a commonly understood distance function $\|\cdot\|$ defined on Ω . This distance function makes it possible to define analogy-based rules. 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 assigns a priority to each rule: when multiple rules apply the higher priority rule is to be followed. We will sometimes write $\omega(r)$, $d(r)$, $a(r)$, and $p(r)$ for the functions 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 model the Agent as mechanically implementing whatever rule book he is given. If no rules apply the Agent is assumed to choose an action at random with both actions being equally likely. To express the assumptions about how agents act formally, write $D(R)$ for the set of states covered by at least one rule, $D(R) = \{\omega' \in \Omega \mid \|\omega' - \omega\| < d \text{ for some } (\omega, d, a, p) \in R\}$. Our assumption is that when rule book R_t has been communicated to the agent and state ω_t is then observed the agent's action choice is

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

We impose limits on communication both by restricting the principal to choose rule books as defined above and by imposing additional restrictions on the set \mathfrak{R}_t of feasible rule books.

First, we incorporate the notion that rules must be analogy-based by assuming throughout the paper that rules must be based on previous observations and that that at most one

rule may be based on any past observation:

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

Second, we incorporate the notion that changing rules can be difficult by imposing restrictions on how rule books are changed over time. We vary these in different sections of the paper. We consider three variants.

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 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.

Most of our analyses will discuss what happens in simple special cases of our model: we assume the state space is a circle and that states are independent draws from a uniform distribution on this set. Mathematically, this corresponds to setting $\Omega = [0, 1]$ and defining distances by $\|x - y\| = \min\{|x - y|, 1 - |x - y|\}$. We 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 which requires that $f(0) = f(1)$. We extend f to a periodic function on \mathbb{R} by setting $f(x) = f(x - \lfloor x \rfloor)$ 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. The Principal’s optimal strategy is unaffected if we change payoffs in any way that leaves $f(\omega)$ unaffected, so we will also normalize payoffs by assuming that $\pi(-1, \omega) + \pi(1, \omega) = 0$ for all ω . With this assumption, the Principal’s payoff is $V = \sum_{t=1}^T \delta^t a_t f(\omega_t)$.

2.2 Motivation

We intend for the model to capture applications for which Ω sufficiently complex so that communicating the first-best contingent plan is an impossibility. In the case of a venture capital firm giving instructions to a new associate, for example, the set Ω would be the set of all possible business proposals that the employee might get to review and the first-best contingent plan would be a list describing all possible proposals and instructing him as to what he should tell upper management in his analysis of each proposal. In the case of a firm giving instructions to a new human resources employee the set Ω would be the set of all possible employees who might interview for a job. In the legal application Ω would be the set of all cases that might come before a lower court.

What does happen in such settings? Obviously, some basic rules can easily be communicated to the Agent before any shared experience exists. For example, the human relations employee can easily be told that a candidate without a bachelor’s degree or one with an undergraduate GPA below 2.8 should not be hired. Our model is not about these rules. Instead, we are interested capturing that Principals in these examples (the experienced HR manager, the head of the venture capital firm, the higher court judges) know much more than can be stated so simply and would like to communicate this to their Agents so that they can more closely approximate the much more nuanced decisions that the Principal would take if she were personally performing the Agent’s job.

We focus on analogy-based rules because we feel that this is descriptively a good model of how knowledge is often communicated in real-world organizations. While it is hard to describe people in the abstract, it is easy to use a particular person that Principal and Agent both know as an example of someone who should or should not have been employed. Moreover, we feel that it is also descriptively accurate that people have an innate sense of “distance” that they can apply at least approximately even in complex spaces. For example,

one can say that John Kerry more similar to Mike Dukakis than is John Edwards. While distances are not conveyed numerically in real-world analogies, we feel that substantial information about desired distances is often conveyed by one of two methods: rules can be conveyed either by using distance words (“don’t nominate anyone who is *at all* like Mike Dukakis”) or by using lists of people who are separated by some distance in order to implicitly define the breadth of the analogy (“don’t let people invite themselves to give the theory seminar unless it’s someone like Paul Milgrom or Eric Maskin”). In the legal application, the case under consideration in the precedent-setting decision serves as the base case for the analogy and the language in the higher-level court decision helps convey how broadly the precedent is to apply.

Our additional restrictions on how rules change over time are designed to reflect difficulties inherent in promulgating and communicating rules. “Vast overwriting” might be appropriate for organizations in which a single principal and agent interact and in which the principal has the time to explain to the agent that are throwing out all of their old rules and starting from scratch with a new way of thinking about the job. “Incremental overwriting” prevents principals from revising existing rules and limits them to making one new rule per period. Retracting rules might be problematic in some applications. For example, if rules diffuse through a multi-agent organization via a word-of-mouth process then it may be hard for agents who hear different versions of the same rule to know which version of a rule is the current version. “No overwriting” additionally rules out priority systems. This may be appropriate in situations where agents are less sophisticated and could get confused if multiple rules were applicable, or if agents follow the first applicable rule they think of.

In the legal application there is a longstanding tradition that cases that have already been decided are “settled law” which should not be contradicted by new decisions except in very unusual circumstances. This could make the no overwriting model most appropriate.

2.3 An aside on one modeling choices: A fully rational variant

Note that in our model specification we made assumptions directly on the agent’s behavior – essentially making them robots who literally follow instructions – rather than giving them

utility functions and making them rational players in a game theoretic model. We have two motivations for this approach. First, we think that what we've done captures well how agents would behave in many applications. Second, what we've done is simpler.

Appropriately rationalizing the model is more complicated than one might guess at first because many rational models would have equilibria in which more communication than seems reasonable can occur in equilibrium. The “unreasonable” equilibria exploit a dubious feature of traditional solution concepts – agents behave as if they know the equilibrium strategies of the other players. For example, if we assume that the agent is rational with the same preferences as the principal, then the first-best can be obtained in a sequential equilibrium by having the principal effectively communicate the function f .

Proposition 1 *Suppose that the function f has a finite number of zero crossings. Suppose the agent is fully rational and has the same preferences as the principal but doesn't know f . Then, there exists a sequential equilibrium in which the optimal action is taken in every period after the first with probability one.*

To see how this can be done, suppose that after ω_1 is observed by the principal and agent the principal issues a second-period rule (ω_1, d, a, p) as follows. Rather than sending a rule that is to be interpreted literally, the principal uses the d part of the rule to encode several pieces of information. First, he sets the initial digits after the decimal point to be a binary string (zeroes and ones) which specifies the number of times f crosses zero. This is followed by a single digit of number 2, which is interpreted as a “stop code” indicating the next component. The second component is a single binary digit indicating which action to take on the first interval; it is followed by another 2. Finally, the location of each of the crossings is encoded as follows. The first digit of this component of the string is the first digit of the first crossing, the second is the first digit of the second crossing, and so on for all n crossings. Then the second digit of the first crossing follows, and so on. In this manner, a single real number encodes everything about the function f that the agent needs to know to take the optimal action. Under the standard assumption that players play best responses to each others' strategies, it will be an equilibrium for the principal to send this message in the first period and for the agent to take the optimal action from the

second period on.²

We have not stated the proposition in the greatest possible level of generality. For example, one can also communicate the locations of a countable number of zero crossings by arranging them in a triangular array, e.g. first giving the first digit of one crossing, then the second digit of the first crossing and the first digit of a second, then the third digit of the first, the second digit of the second, and the first digit of the third, and so on. One could try to refine away these equilibria by ruling out equilibria which involve the agent processing an infinite string of bits, but there are many other coding schemes one could use. For example, the Principal could communicate any rule that can be described in finite space in English-language text by encoding the text in ASCII just as is done when sending an e-mail.

We think of Proposition 1 not as a positive result, but rather as a cautionary note illustrating that the standard “rational” approach is unappealing for the problems we are trying to address. Specifically, the assumption of best-responding to equilibrium strategies embeds a degree of knowledge of the structure of the problem and of the meaning of language that does not fit with other aspects of our formulation. In the remainder of the paper we rule out schemes like the ones described above by assuming directly that agents interpret rules literally.

2.4 Another aside on our modeling choices: Asymmetric analogies and random choices

In this second aside we discuss the motivation for two other assumptions: our restriction to “symmetric” analogies and our assumption that agents choose randomly when no rule has been defined. It is another conscious part of our modeling strategy – we see it as the simplest way to create the kind of accuracy vs. breadth tradeoff that will arise in any complex state space – and are happy with it even though we could have avoided the obvious critiques of these assumptions via other means.

The symmetry restriction is that we have assumed the principal can only issue rules that apply when $\|\omega_t - \omega\| < d$ rather than allowing rules that apply when $\omega_t \in [\omega - d_1, \omega + d_2]$.

²This result relates to the literature on the necessary size of message spaces in implementation theory. See for instance Reichelstein and Reiter (1988) and Segal (2010) and the references therein.

Many readers will no doubt quickly think of a two-step critique. First, asymmetric rules are used in practice, e.g. an admissions committee member could be told that students like STUDENT X should only be accepted if their technical background is as strong or stronger. Second, allowing asymmetric rules would have a dramatic effect in some situations. For example, when f is as pictured in Figure 3 and $\omega_1 = \frac{3}{8}$, a Principal restricted to asymmetric rules will be forced to define a highly inefficient rule, whereas a Principal who could explain that action 1 should be taken if the action is within $\frac{3}{8}$ of ω_1 on the left or within $\frac{1}{8}$ on the right would be able to perfectly describe where action 1 should be taken.

We agree with both steps, but still feel that our symmetry assumption is a very good one. Recall that our motivation for considering analogy-based rules is that the world is sufficiently complex so as to make it impossible to describe the sets of states in which each action is optimal. In such environments it is natural that all feasible rules will only crudely approximate the first-best. For example, if the set of states in which action 1 is optimal were the irregular shaped region in \mathbb{R}^2 shown in Figure 2, then the best feasible rule would be inaccurate regardless of whether the agent could issue rules to apply on squares, or rectangles, and regardless of whether these had to be symmetrically centered on the point in question. We could have written this paper using a two-dimensional model without the symmetry restriction. Several of our propositions (including 2, 6, and 8) could be easily modified to apply in such an environment and examples could have been constructed to convey other points.

However, we feel that our simple one-dimensional model with the symmetry restriction captures the same breadth vs. accuracy tradeoffs and provides a neater way to convey our insights.

Similarly, some readers may wonder about our assumption that agents choose action randomly if a state occurs for which no rule has been defined. Why don't we let the Principal decree that a particular action should be the default if no other rule applies?³ Again, our answer is that yes, principals could often do this in practice, and yes it would allow the Principal of our model to do substantially better in some examples, but we don't

³This question mostly applies to the no overwriting variant. In the other variants the Principal can effectively issue default instructions by issuing a low-priority rule that applies to all states within distance one half of some previous state.

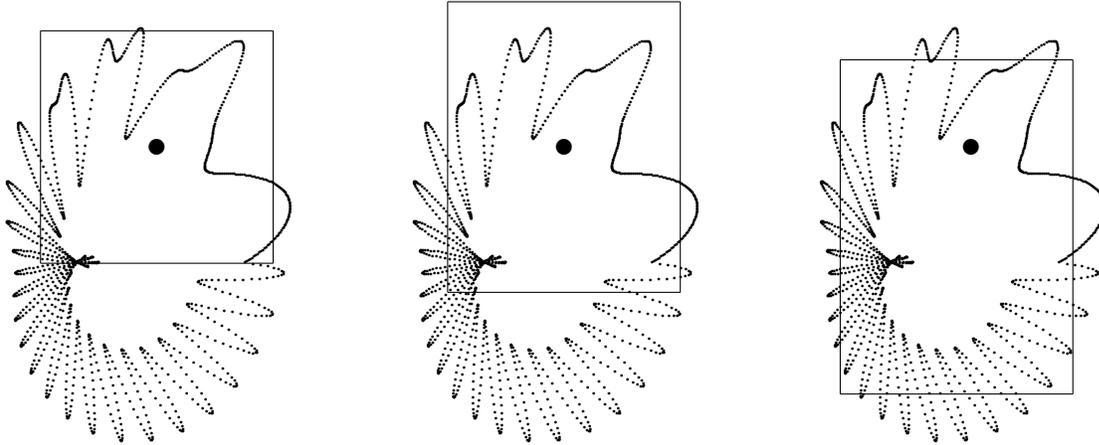


Figure 2: A higher-dimensional example

want to allow it because it's only real effect would be to force us to come up with more complicated examples to illustrate some of the things we want to illustrate.

3 Two Periods: A Simple Theory of Intentionally Inaccurate Rules

In this section we discuss the two-period version of our model. The two-period version is simple and brings out some basic economics about how the design of rules.

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) .

3.1 Endogenous inaccuracy

The next proposition contains two observations about optimal rules in this simple environment. The first is that rules do enable the Principal to increase her payoff except in one special case. The second is that rules are designed to have “excess breadth” in the sense that they are intentionally made to cover a broader domain than the domain on which they only produce correct decisions.

Proposition 2 *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 $d^*(\omega_1) > 0$ and the principal's expected second period payoff is positive, $E(V_2|\omega_1) > 0$.*
2. *Any interior optimal choice $d^*(\omega_1)$ is such that $f(\omega_1 - d^*(\omega_1)) = -f(\omega_1 + d^*(\omega_1))$.*

Proof

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

$$E(V_2|\omega_1) = a \int_{\omega_1-d}^{\omega_1+d} f(x)dx.$$

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 .

$E(V_2|\omega_1)$ 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 2 is in a sense an extreme statement about excess breadth – it says that rules are expanded until the average value of the rule in the marginal cases, $\omega_1 - d^*$ and $\omega_1 + d^*$, is exactly zero. Unless both points are cases of indifference, this implies that the rule must be leading to some 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 3 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 $\omega \approx \frac{1}{4}$. As noted in the proposition, the principal will, in fact, 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 $\omega_1 + d^* = \frac{5}{8}$.

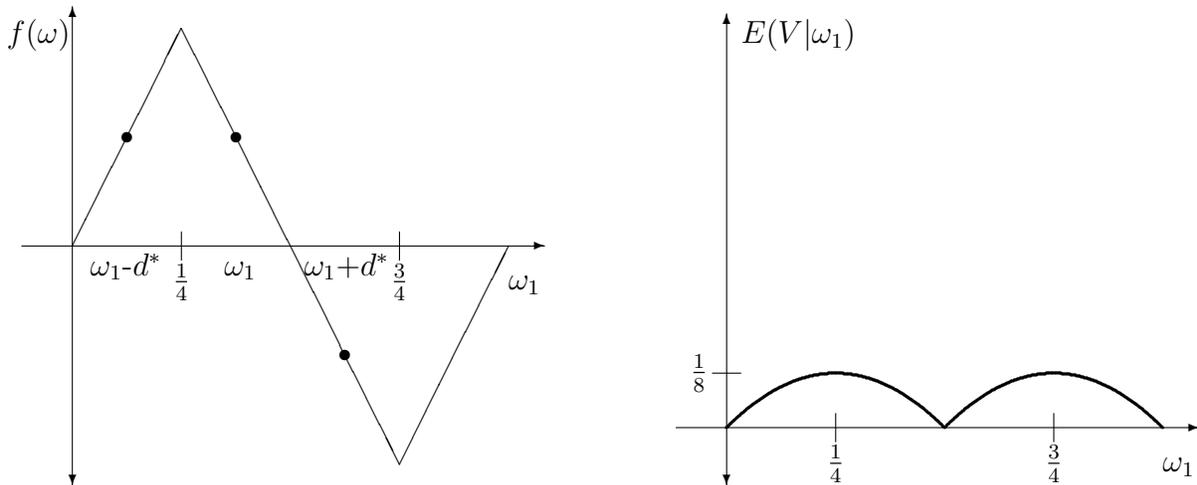


Figure 3: A two period example

3.2 Inefficient rules as a source of performance differences

There is a sizable 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)). In our model historical accidents can be a source of performance differences: “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.

The right panel for Figure 3 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 applicable. The payoff function turns out to be quadratic on the intervals around the two maxima. The quadratic function implies that there is no first-order loss in the efficiency of the rule-book when the initial state ω_1 is very close to $\frac{1}{4}$ or $\frac{3}{4}$. But the Principal does end up with a completely worthless rule book in the worst states.

3.3 Forms of inefficiency

When the initial state is not equal to one of the ideal points, it is obvious that the principal must reduce the breadth of the rule and/or its accuracy and therefore ends up with a lower payoff. Which of these types of inefficiencies occurs is less obvious. For the function in the example, it turns out that the principal entirely sacrifices accuracy: the breadth d^* that the Principal chooses 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 are 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 so the breadth/inaccuracy choice is indeterminate.

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 $f(\omega_1 - d^*) = f(\omega_1 + d^* - \frac{1}{2})$.⁴ 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 $d^* = \frac{1}{4}$. The latter

⁴For 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 $[0, 1]$.

implies that $\omega_1 = y + \frac{1}{4}$, a non-generic case.

3.4 Sophistication of the rule-making language

Throughout this paper we restrict the Principal to issue a single rule based on each experience and to issue rules that apply on symmetric neighborhoods around previously realized states. Relaxing these constraints illustrates that a Principal will be better off if she is able to communicate in a more sophisticated language. 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 . If the Principal also had the ability to communicate multiple rules and to communicate asymmetric rules she could achieve this payoff for any ω_1 : she would just again use a default rule of to apply for all states and then a higher priority rule on an asymmetric interval around ω_1 . Hence, the model does suggest that an ability to convey more sophisticated rules could be a source of performance differences depending on the historical events that arise.

In the particular environment of this example, the more sophisticated rules described above do not sound 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 seem less reasonable, e.g. the principal could dictate N different actions on concentric rings around w_1 by defining a first rule that applied 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 applied at distances less than d_3 , and so on. And it is only the unrealistic simplicity of the one-dimensional model with two actions that makes full efficiency achievable. Somewhat more sophisticated rules than we use in most of our analyses may be a source of interesting additional category of rules to study, but we felt that the three variants in the current paper were preferable as a simple framework in which to studying the second-best payoffs that analogy-based rules can provide in a complex environment.

4 Three Periods: Learning Curves and Building Routines

A number of literatures have focused on the dynamics of firm efficiency. There is a vast theoretical and empirical literature on learning curves. Another empirical regularity which has been documented is that firms often “start-small”. Gibbons (2006) refers to the possibility that communication may start out vague and become more precise over time as “building a routine”.⁵ In this section, we develop observations related to these topics by examining the simplest dynamic version of our model: the three period version. One observation related to the starting-small literature is that option value considerations can offset some of all of the incentive to issue excessively broad rules.

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 $\delta > 0$.⁶

4.1 No overwriting

Our first result on the no-overwriting model highlights that option value considerations lead the principal to “start small” and define narrower rules in the three-period model than would be used in the two-period model.

Proposition 3 *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

⁵See Chassang (2010) for a recent model.

⁶This is obviously a function of δ . We omit the dependence from the notation where it will not cause confusion.

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 expected payoff in the second period. In the third period, the expected payoff conditional on $\omega_3 \in [\omega_1 - d(r'_2), \omega_1 + d(r'_2)]$ is again greater under $(r_2^m, r'_3(\omega_2))$ by the optimality of r_2^m . The third period expected payoff conditional on $\omega_3 \notin [\omega_1 - d(r'_2), \omega_1 + d(r'_2)]$ is identical under the two rules.

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 and has excess breadth at the margin 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 4 provides an illustration. The left panel graphs the 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, $d(r_2^*(\omega_1))$ for same function f pictured in Figure 3.⁷ 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 3 the optimal second-period rule is narrower. The solid line in the figure gives this breadth. Note that sometimes the optimal rule has zero breadth. This is the case near $\omega_1 = 0$, $\omega_1 = \frac{1}{2}$ and $\omega_1 = 1$. For other values of ω_1 the optimal rule has breadth close to $\frac{1}{4}$.

The contrast between the static and dynamically optimal rules illustrates that our model provides a potential explanation for observations that firms sometimes intentionally “start small.” Watson (2002) argues that one 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 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. Early in its life a firm may decline opportunities that would be profitable in the short run because accepting the opportunity would establish a less-than-ideal precedent.

⁷The discount factor of $\delta = 1$ was used for these graphs.

The three-period model remains a model in which historical accidents can lead to persistent performance differences. The right panel graphs the expected payoff as a function of ω_1 . In the region where the optimal rule has positive breadth, the payoff function again declines smoothly away from maxima at $\omega_1 = \frac{1}{4}$ and $\omega_1 = \frac{3}{4}$.

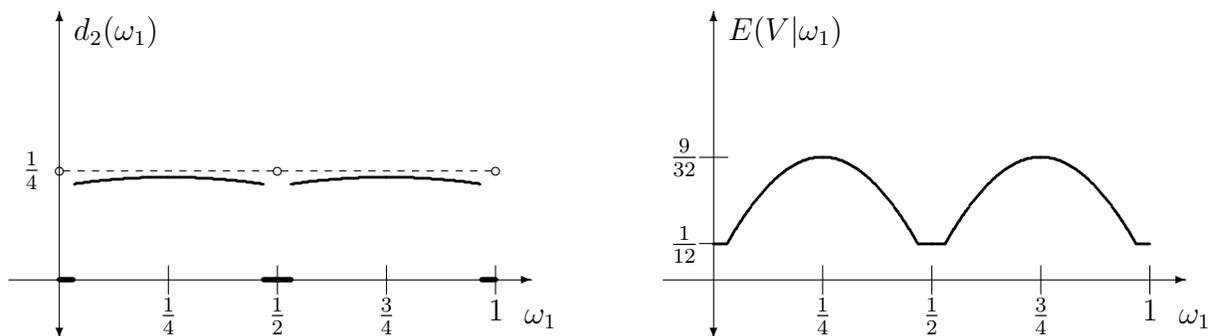


Figure 4: A three period example

We emphasized above that option-value considerations work against the incentive to promulgate excessively broad rules which we had noted in the two-period model. We would like to point out here that indeed the option value consideration can dominate in a fairly strong sense. We will say that the initial rule in the three-period model has *insufficient breadth at the margin* if $d_2^* < \frac{1}{2}$ and $f(\omega_1 - d_2^*)$ and $f(\omega_1 + d_2^*)$ are both of the same sign as a_2^* . In words, this says that the Principal chooses not to expand the domain of the rule even though it would be correct in all marginal cases. Our next proposition says that this will occur for some model specifications:

Proposition 4 *There exist payoff functions for the three-period no-overwriting model for which the optimal second-period rule has insufficient breadth at the margin for a positive measure of values of ω_1 .*

The left panel of Figure 4 illustrates that the payoff function given above provides the necessary example. One case in which it is intuitive that the optimal rule will have insufficient breadth is when ω_1 close to zero, one-half, and one. Here, the optimal breadth is zero, which will satisfy the insufficient breadth at the margin condition. 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 $[\frac{1}{8}, \frac{5}{8}]$.

A second noteworthy situation in which we get insufficient breadth at the margin is when $\omega_1 = \frac{1}{4}$. In this situation one's initial thought might be that the principal has received a fortunate draw and will choose $d = \frac{1}{4}$ to define the exactly optimal rule on half of the states. This exactly optimal rule, however, is still of zero value at the margin. Under the no overwriting assumption the marginal breadth has an opportunity cost. For example, if the second period state is $\omega_2 = \frac{5}{8}$ the principal will be able to dictate that action -1 be taken when $\omega_3 \in [\frac{1}{2}, \frac{3}{4}]$, but will be unable to promulgate a broader rule covering $[\frac{1}{2} - \epsilon, \frac{3}{4} + \epsilon]$.

We omit a formal proof of Proposition 4 because it is unexciting and the graph serves as a numerical proof.

4.2 Incremental overwriting

The above discussion has focused on the no-overwriting version of our model. In the incremental overwriting model, some of the constraints which we had mentioned in explaining why a principal might intentionally make a rule excessively narrow no longer exist. For example, the principal can overwrite the interval $[\frac{1}{2} - \sigma, \frac{1}{2}]$ 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 “start small” in the sense of choosing a narrower second-period rule than would be chosen by a myopic agent.

Proposition 5 *In the three period incremental overwriting model there exist payoff functions f and states ω_1 for which $d(r_2^*(\omega_1)) < d_2^m$.*

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.⁸ For $\omega_1 = \frac{1}{6}$ 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 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 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 $r_2^* = r_2^m$. In this section we regard the vast overwriting model primarily as reinforcing the observation that it is constraints on the Principal's ability to modify rules which create the incentive to start small.

⁸The example uses a discontinuous f to make the computations easy. A continuous example could be obtained by choosing a nearby continuous function.

5 The Infinite Horizon Model

In previous sections we have used two- and three-period models to bring out some of the primary economic considerations of rule development. In real-world organizations, of course, rules development is a longer and more gradual process. In this section we use the infinite horizon version of our model to develop some additional insights. The first subsection focuses on long-run performance differences. The second examines the natural idea that patient Principals will develop nearly optimal rules.

5.1 Long run performance differences

In the long run the Principal will have many opportunities to issue and/or refine rules. While the rule book may forever grow longer and longer as these opportunities keep arising, it is intuitive that these changes cannot continue to be significant: each firm's rule book should in some sense settle down and converge to some limit rule book. We should be able to talk about long-run efficiency differences across firms in terms of these limit rule books.

Our first result provides a way to formalize this intuition. The main idea is to focus on each rule book's expected-payoff rather than on the growing list of individual rules it might contain. Specifically, we define π_t to be the expected payoff in period t given the rule book which the Principal will have defined by that point in time. This π_t is a function of the random draws $\omega_1, \omega_2, \dots, \omega_{t-1}$ and thus can be thought of usefully both as a function of the random draws and as a random variable. A convergence theorem is:

Proposition 6 *Write $\pi_t(\omega_0, \dots, \omega_{t-1}) \equiv E(\pi(a_t, \omega_t, R_t(\omega_0, \dots, \omega_{t-1})))$ for the expected payoff in period t given the rule book which the principal has defined in equilibrium. In any of the overwriting cases there exists a limiting random variable π_∞ such that π_t converges to π_∞ in probability.*

Proof

In both the no overwriting and vast-overwriting versions of the model we have $E(\pi_t | \pi_{t-1}) \geq \pi_{t-1}$. In the vast overwriting case this is because it is always possible to define a weakly better rule book with a superset of previous realizations. In the no-overwriting case this is because the rule that the principal defined in period $t - 1$ must have been better than

defining a rule dictating that the opposite action be taken on the same domain. The process $\pi = (\pi_t : t \geq 0)$ is adapted and $E(|\pi_t|)$ is finite for all t , so π is a submartingale. Hence, Doob's Forward Convergence Theorem (Williams 1991, p.109) applies and guarantees the existence of a π_∞ such that π_t converges to π_∞ almost surely. This implies convergence in probability.

In the incremental overwriting case the argument is a little more complicated. Here, we define $v_t(\omega_0, \dots, \omega_{t-1}) \equiv (1 - \delta) \sum_{s=0}^{\infty} \delta^s E(\pi(a_{t+s}, \omega_{t+s}, R_t(\omega_0, \dots, \omega_{t+s-1})) | \omega_0, \dots, \omega_{t-1})$ to be the expectation of the average discounted future payoff at each point in time. Using the facts that $v_t = (1 - \delta)\pi_t + \delta E(v_{t+1} | \omega_0, \dots, \omega_{t-1})$ and $\pi_t(\omega_0, \dots, \omega_{t-1}) \leq v_t(\omega_0, \dots, \omega_{t-1})$ (which follows from the fact that the principal has the option of deciding to never change the rule book) we have $E(v_{t+1} | v_t) \geq v_t$ so again Doob's theorem guarantees the existence of a π_∞ such that v_t converges to π_∞ almost surely. To show that π_t converges in probability to π_∞ it then remains only to show that $v_t - \pi_t$ converges in probability to zero. To see this note that $v_t - \pi_t$ is a nonnegative valued random variable and $E(v_t - \pi_t) = \frac{\delta}{1-\delta} (E(v_{t+1}) - E(v_t))$ converges to zero because $E(v_t)$ and $E(v_{t+1})$ both converge to $E(\pi_\infty)$. For nonnegative random variables convergence of the expectation to zero implies convergence to zero in probability.

QED

One can think of the π_∞ defined in the above proposition as the efficiency of the rule book which the principle develops "in the long run". The statement that π_t converges to π_∞ is giving a formal sense in which we can say that the rule book that the principal develops over time always settles down and approaches some limit. The sense in which we have shown that it approaches this limit is a probabilistic statement about closeness in payoff space – the convergence in probability means that for t large there is a high probability that the efficiency π_t of the period t rule book is close to the limiting efficiency that the principal will obtain in the long run "after" all future modifications.

The distribution of π_∞ describes how often we end up with a limiting rule book that is very efficient/moderately efficient/highly inefficient in the long run. It can be thought of as describing the distribution of performance levels we would expect to see among ex ante identical firms purely because of historical accidents. The distribution of the random

variable π_t describes how frequently each efficiency level will be seen at time t . Convergence in probability implies convergence in distribution. Hence, Proposition 6 implies that the distribution of π_∞ can also be thought of as an approximation to the distribution of the efficiency of the period t rule book π_t for any large t .

The distribution of π_∞ will depend on all of the elements of our model: the underlying state-contingent payoff function; how patient the principal is; the overwriting restrictions the Principal faces; etc. In the no overwriting version of our model history has a long-run impact. For some realizations of $\omega_0, \omega_1, \dots, \dots$ the limiting rule book will be highly efficient and for others it may be very inefficient.

Figure 5 provides an illustration of the limiting distribution for one specification of the model: it plots the density of π_∞ for the f function of Figure 3 under the assumption that the principal acts myopically ($\delta = 0$) and is subject to the no overwriting constraint. One noteworthy feature of the distribution of the graph is that there is a significant amount of mass close to $1/4$. This reflects the simplicity of the optimal decision rule: as noted earlier there no first order loss in the second period payoff when the initial state differs slightly from $1/4$ or $3/4$; and whenever this happens the Principal's will eventually fill out the rule book to be nearly efficient in all other states. It is also noteworthy, however, that there is a substantial amount of mass a much lower efficiency levels. This highlights that our model predicts substantial path dependence. A series of early "bad" realizations—in the sense that they lead the firm to promulgate substantially inefficient rules—have a persistent effect. Extremely inefficient rule books arise only rarely in this specification – they only occur when the Principal receives multiple unlucky draws.

In the vast overwriting version of the model there will be no long run inefficiency: π_∞ would be equal to the first-best payoff with probability one. When the state space is a line and $f(x)$ crosses zero a finite number of times the argument is simple and yields a particularly strong conclusion. The set of states for which $a = 1$ is optimal (and for which $a = -1$ is optimal) is a union of intervals. Divide each interval into four equal subintervals. As soon as the observed history $\omega_1, \omega_2, \dots, \omega_T$ contains points both in the second and the third subintervals of each interval the first-best can be achieved by overlapping rules dictating that $a = 1$ (or -1 as appropriate) should be chosen in a neighborhood extending

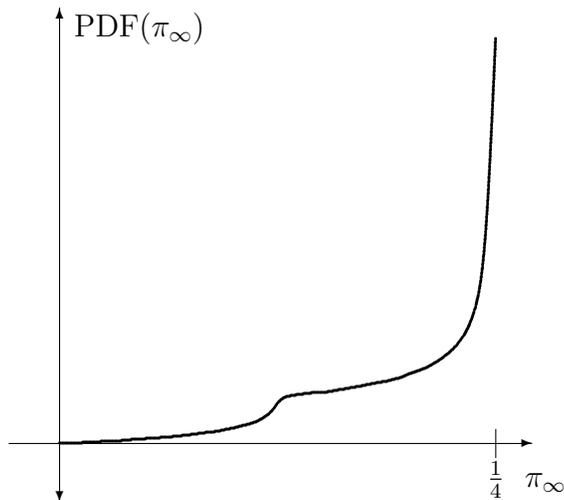


Figure 5: Distribution of long-run payoff π_∞ with myopic play

from those points exactly to the boundary of the region where $a = 1$ is optimal. This implies, in fact, that the first best rule book becomes possible in finite time (with probability one). In more general and complex state spaces like that pictured in Figure 2 the first-best will not be possible at any finite time t , but the result that there is no long run inefficiency will generalize under appropriate regularity conditions: when the set of observed ω_t is very rich an approximately optimal rule book can be achieved simply by defining the appropriate rule on a very small neighborhood around each observed ω_t . The lesson we would take away from this model is that with vast overwriting history matters but only in terms of affecting how long it takes and how much inefficiency is incurred before efficient rules to develop.

In the incremental overwriting model we conjecture that there is also no long run inefficiency under some set of regularity conditions. The intuition behind this conjecture is that early on inefficient rules like that illustrated in Figure 3 will be promulgated, but once the whole interval is covered opportunities will continually arise to overwrite previous mistakes including some which allow an interval of mistakes to be completely eliminated. Hence, we conjecture that the difference between the vast overwriting and incremental overwriting models is not in terms of what happens in the very long run, but in terms of how quickly or slowly efficiency is approached.

5.2 Limiting payoffs for patient principals

The figure in the previous subsection illustrated the efficiency of the limiting rule books developed by completely impatient Principals. It is intuitive that patient Principals should do better. For example, we noted in our analysis of the three-period model that patient Principals will “start small” and decline to issue any rule when the initial realization is sufficiently unfortunate. In this section we present some results on how (very) patient Principals do.

5.2.1 Vast overwriting

Limiting payoffs for patient principals are easiest to characterize in the vast-overwriting version of our model. The rules developed in this model are independent of the Principal’s discount factor because the Principal always rewrites rule book to be the best possible after each period. As a result a fully optimal rule book is not only achieved asymptotically, but develops in finite time if the function f is well behaved. 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. A result showing the Principal reaches this upper bound in finite time is:

Proposition 7 *Consider the infinite horizon vast-overwriting version of our model with any discount factor $\delta > 0$. 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 $t \geq T$.*

Proof

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 .⁹ Hence, R_t will result in action $a^{FB}(\omega_t)$ being chosen. Given that R_t achieves payoff V^{FB} for all $t > T$, any other optimal rule book must also lead the Principal to choose action $a^{FB}(\omega_t)$ with probability one.

QED

Remarks

1. An immediate corollary of the Proposition is that the expected per period payoff of the Principal converges to V^{FB} in the $\delta \rightarrow 1$ limit: the principal gets the best possible payoff from some time T on and when δ is close to one this is all that matters. Formally one can show this by defining T to be the time at which the first-best rule book is first achieved and writing the Principal's payoff as the expectation over T of the conditional expected payoff given T .
2. 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.
3. 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 in finite time.

⁹When two rules apply they give the same action so the priorities we specified were unimportant.

5.2.2 No overwriting and incremental overwriting with δ close to 1

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 8 *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 7 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 expected 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.

A final result of this section is that when no overwriting is allowed the first-best is not achieved in finite time. As in the situation described in the second remark after Proposition 7, this is a mechanical consequence of the principal's limited ability to define rules. This time we do give a formal statement and a proof.

Proposition 9 *Consider the no overwriting version of our model. Assume that there is no nontrivial subinterval (a, b) on which f is almost everywhere equal to zero. Then, $\text{Prob}\{E(V_{t+1}|\omega_1, \dots, \omega_t) = V^{FB}\} = 0$ for every t .*

Proof

Let Ω^∞ be the set of all sequences $\{\omega_t\}$. We show the result by partitioning Ω^∞ into two disjoint subsets $\Omega^\infty = \Omega_1 \cup \Omega_2$ and showing that the result holds for almost all $\{\omega_t\}$ in each Ω_i provided that Ω_i is not of measure zero.

The division is simple: we write d^0 for the width of the first nontrivial rule issued by the principal, i.e. $d^0(\{\omega_t\}) = d(r_s)$ where s is such that $d(r_s) > 0$ and $d(r_t) = 0$ for all $t < s$ and set Ω_1 to be the set of all $\{\omega_t\}$ with $d^0(\{\omega_t\}) \geq \frac{1}{2}$.

Case 1: $\{\omega_t\} \in \Omega_1$

In this case, the agent chooses at random in the first $s - 1$ periods and chooses a state-independent action in all future periods. Hence, we have $E(V_t|\omega_1, \dots, \omega_{t-1}) = 0$ for all $t < s$ and $E(V_t|\omega_1, \dots, \omega_{t-1}) \leq |\int f(x)dx| < \int |f(x)|dx = V^{FB}$ for all $t > s$, with the strict inequality following from our assumption that f is continuous and takes on positive and negative values.

Case 2: $\{\omega_t\} \in \Omega_2$

It suffices to show that with probability one there is a nonempty interval that is not covered by any rule. This follows easily by induction. In period s there is an interval of width $1 - 2d^0$ that is uncovered. If in any period t , some interval of width w with $0 < w < 1$ is uncovered, then a subinterval remains uncovered in period $t + 1$ unless ω_t lies exactly in the middle of the uncovered interval from period t . This occurs with probability one.

QED

The impossibility of obtaining the first best in finite time does not extend to the incremental overwriting model. Consider for example, the example shown in Figure 3. There is positive probability that first draw will be similar to that pictured and the second will be just to the right of three-fourths leaving the principal with a rule book with rules that are correct except on two intervals: one just to the right of zero and one just to the right of one-half. If the next two states lie in $[1/8, 1/4]$ and $[5/8, 3/4]$ the principal will issue rules (overlapping with previously defined rules) that completely fix the mistakes.

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, suppose there are two tasks $i = 1, 2$. The principal's payoff on task i is $\pi_i(a_i, \omega_i)$, where a_i is the action taken on task i and ω_i is a task-specific state of nature. As before, the principal's payoff in the full game is the discounted sum of her per period payoffs: $V = \sum_{t=1}^T \delta^t (\pi_1(a_{1t}, \omega_{1t}) + \pi_2(a_{2t}, \omega_{2t}))$. In each period, Principal has two options. First, he can communicate rules to both agents at the start of the period. They will then carry out these instructions as in our standard model. We call this "full delegation." Second, the Principal can instead carry out one of the actions herself. If she does this, we assume that she does not have sufficient time to communicate with the other agent so that agent carries out the most recently communicated rule book. We assume the Principal does observe both ω_1 and ω_2 and can base rules on these in the future. We call this "partial delegation."

In some situations it would be appropriate to assume that delegation structures are

time invariant. In this case we have:

Proposition 10 *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 7 and 8 we know that as $\delta \rightarrow 1$ the principal's average per-period payoff converges to V^{FB} . Partial delegation cannot achieve this since there is no communication on one of the tasks and the expected payoff on that task is zero. Even though the first best is obtained on the other task, the sum of the payoffs on the two tasks will be worse than a payoff that approaches the first-best on both as $\delta \rightarrow 1$.

QED

In other situations it may be reasonable to assume that firms can change their delegation structure as the firm matures. In this case, our model becomes a theory of organizational change. In the first period, the Principal must practice partial delegation. Full delegation will provide no benefits because there is no history on which to base rules, whereas with partial delegation she can achieve the first-best payoff on one of the two tasks. In the vast overwriting case, at least, the Principal must switch at some point to full delegation: as history accumulates it will eventually become possible to achieve a superior payoff by delegating both tasks. (Switching directly to partially delegating the other action is inferior because before any rule has been communicated it provides a payoff that is no better than the payoff that was achievable at the start of the game.) The long-run organization structure is indeterminate. If the first-best is achievable when the Principal first switches to full delegation then there is obviously no benefit from further switches in organizational form. But if the Principal can communicate the first-best rule on exactly one task at this time, then she strictly prefers to use partial delegation at least for some period of time after this.

Other variants of this model could be considered. For example, one could assume that the Principal and Agent do not accumulate shared experiences when the Principal is working on the other task. Principals in such models will face a tradeoff: do they take an

optimal action now or invest in developing shared experience. The dependence which the basic model exhibited can be magnified in such settings: there is both a direct effect of history on payoffs and an additional indirect effect whereby a fortunate history allows the Principal to adopt a delegation structure that is more efficient in the long run. We leave the exploration of such models to future work.

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 $\{\omega \mid \|\omega - \omega_t\| < d_t\}$ 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 categories¹⁰. 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 work-

¹⁰In the case of overlap a state can belong to multiple categories.

ers who may be black or white. Thus there are for 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.

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.

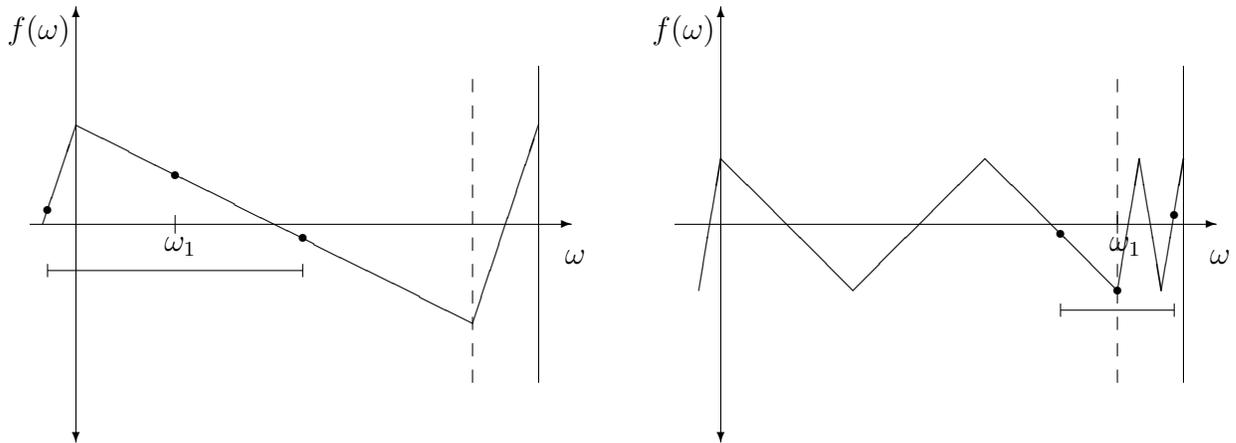


Figure 6: Categorization examples

In the second panel of figure 6 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

In this paper we have developed a simple model of second-best rule development. We model rules as instructions provided to agents who must react to idiosyncratic events in a timely manner. We assume that rules must take the form of analogies telling the agent how he should behave in situations similar to previously observed events. Accumulated history is what enables Principals to communicate complex contingent rules to agents. But Principals speed up the rule-development process by broadening rules in ways that make them somewhat inaccurate. We have discussed several considerations that come out of the basic economics of the model: the short-run incentive for excess breadth that comes from considering whether the rule has value in marginal cases, dynamic option-value considerations which provide an incentive to keep rules narrow, and the long run payoffs that will be achieved.

Our model is potentially relevant to a number of applied topics. One, which we mentioned in the introduction but have not emphasized since, is the development of refined interpretations of statutes by common law courts. Here, the Principal should be thought of as a higher-level court which has a complete understanding of how it would like the law to be applied to every possible case. The way in which it can communicate to lower courts, however, is by writing decisions explaining how the law should have been applied in particular cases. These prior cases are the basis of our analogies and the distance parameter can be interpreted as the writing around the decision suggesting that it should apply broadly or narrowly. Our incentive results describe tradeoffs higher-level courts face. The desire to change how lower courts act right away provides an incentive to write excessively broad

decisions. The desire to wait for a better test case is the counteracting force that may lead courts to decline to hear a case they would overcome. The understanding that precedents are not to be overturned except in rare circumstances is an “overwriting” constraint and our model illustrates that these may be important to the efficiency of the rules which develop. These ideas could certainly be extended in future work and other factors unique to this application (costs of hearing appeals, competing higher-level courts with different opinions, hierarchies of courts) could provide opportunities for many interesting extensions.

Our other main intended application, of course, is to the informal analogy-based rules that managers communicate to their subordinates. Here, our motivation was that the world is a complex place and baseline rules that can be communicated in general terms will be highly inefficient. Common experiences create the opportunity to improve on this baseline. Here, our model is potentially relevant to a number of existing literatures. One which we have stressed is the literature on persistent performance differences between seemingly similar firms. There are a number of existing explanations for such differences, e.g. they might correspond to different equilibria of a repeated game (Leibenstein (1987)). Our model provides an alternate explanation emphasizing the natural path dependence of rule development. 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.

Another related literature is that on learning curves. In our model firms performance will usually improve over time as each period’s experience provides an opportunity to improve upon existing rules. These improvements may account for the typical shape of a learning curve: at first big improvements can be made by defining broad rules but eventually most of the space will be covered and new rules will have only small incremental benefits. Rates will be different for different firms due to path-dependence and will also be related

to primitives like the rate at which experiences arrive, the patience of the Principal, and the scope for overwriting. It would be interesting to explore whether such explanations can illuminate existing empirical evidence on heterogeneity of rates of learning (Argote and Epple (1990)).

Similar comments can be made with regard to other topics within the literature on organizations. Our model highlights the effects of a number of parameters on performance difference. The contrasts between what we find in examining different variants of our model also brings out the potential relevance of the ability to “overwrite” as a characteristic that may distinguish how firms fare. For example, it may be that overwriting existing rules is easier in small organizations or in organizations where the Principal and Agent of our model directly interact, and overwriting is more difficult in large organizations where there is no direct communication and rules diffuse in a less controlled word of mouth process. Whether proxies for differences in the ability to “overwrite” might explain whether some firms in an industry end up more efficient than others, or whether industry-level proxies for differences in the ability to overwrite might explain why there is more performance heterogeneity in some industries than others are examples of questions that our model raises.

We conclude by highlighting a direction for further work. It would be desirable 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 think of our assumptions restricting firms to use analogy-based rules as a reduced form way of capturing what might occur in a model where state spaces were highly complex and some cost of communication made communicating via analogies preferable to trying to communicate by describing sets of states in other terms. It would be interesting to explore what kinds of environments, or what other cognitive limitations, give rise to the inability to perfectly communicate, and to see whether such models might suggest alternate rule-making models to the one we explore here.

References

- Philippe Aghion, Patrick Bolton, Christopher Harris, and Bruno Jullien. Optimal learning by experimentation. *Review of Economic Studies*, 58:621–654, 1991.
- Nabil I. Al-Najjar and Mallesh Pai. Coarse decision making. *Northwestern University Working Paper*, 2009.
- 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.
- Patrick Bolton and Mathias Dewatripont. The firm as a communication network. *Quarterly Journal of Economics*, 109:809–839, 1994.
- Colin F. Camerer and Roberto A. Weber. Cultural conflict and merger failure: An experimental approach. *Management Science*, 49:400–415, 2003.
- Bruce Chew, Kim Clark, and Timothy Bresnahan. *Measurement, coordination and learning in a multiplant network*, chapter 5. Harvard Business School Press, 1990.
- Vincent Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50:1431–1452, 1982.
- Jacques Cremer. A partial theory of the optimal organization of bureaucracy. *Bell Journal of Economics*, 11:683–683, 1980.

- Jacques Cremer, Luis Garicano, and Andrea Prat. Language and the theory of the firm. *Quarterly Journal of Economics*, CXXII:373–408, 2007.
- Wouter Dessein. Authority and communication in organizations. *Review of Economic Studies*, 69:811–838, 2002.
- Wouter Dessein. Adaptive organizations. *Journal of Political Economy*, 114:956–995, 2006.
- Mathias Dewatripont and Jean Tirole. Modes of communication. *Journal of Political Economy*, 113:1217–1238, 2005.
- 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.
- F.A. Hayek. The use of knowledge in society. *American Economic Review*, 35:519–530, 1945.
- Ryan Kellogg. Learning by drilling: Inter-firm learning and relationship persistence in the texas oilpatch,. *CSEM Working Paper*, 2008.
- Frank H. Knight. *Risk, Uncertainty, and Profit*. Hart, Schaffner & Marx; Houghton Mifflin, Boston, MA, 1921.
- Harvey Leibenstein. *Inside the Firm: The Inefficiencies of Hierarchy*. Harvard University Press, 1987.
- W. Bentley MacLeod. Planning costs and the theory of learning by doing. *USC Working Paper*, 2002.
- 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.
- Sendhil Mullainathan. Thinking through categories. *NBER working paper*, 2002.
- Sendhil Mullainathan, Joshua Schwartzstein, and Andrei Shleifer. Coarse thinking and persuasion. *Quarterly Journal of Economics*, 123(2):577–619, 2008.
- Stefan Reichelstein and Stanley Reiter. Game forms with minimal message spaces. *Econometrica*, 56:661–692, 1988.
- Larry Samuelson. Analogies, adaptation, and anomalies. *Journal of Economic Theory*, 97: 320–366, 2001.
- Ilya Segal. Nash implementation with little communication. *Theoretical Economics*, 5(1): 51–71, 2010.
- Reinhard Selten and Massimo Warglien. The emergence of simple languages in an experimental coordination game. *Mimeo*, 2006.
- Herbert Simon. A formal theory model of the employment relationship. *Econometrica*, 19: 293–305, 1951.
- Dimitri Vayanos. The decentralization of information processing in the presence of interactions. *Review of Economic Studies*, 70:667–695, 2003.
- Joel Watson. Starting small and commitment. *Games and Economic Behavior*, 38:176–199, 2002.
- Birger Wernerfelt. On the nature and scope of the firm: An adjustment-cost theory. *Journal of Business*, 70:489–514, 1997.
- Birger Wernerfelt. Organizational languages. *Journal of Economics and Management Strategy*, 13:461–472, 2004.

David Williams. *Probability with Martingales*. Cambridge University Press, 1991.

Oliver E. Williamson. The economics of governance. *American Economic Review Papers and Proceedings*, 2:1–18, 2005.

Timothy Van Zandt. Real-time decentralized information processing as a model of organizations with boundedly rational agents. *Review of Economic Studies*, 88:633–658, 1999.