

Loopholes and the Evolution of Contract Form^{*}

Philippe Jehiel[†] and Andrew F. Newman[‡]

March 2014

Abstract

We explore a simple overlapping-generations model of the evolution via “loopholes” of incentive systems such as contracts, organizations, or regulatory regimes. Young principals, who are uncertain about what actions will be feasible for their agents, get limited observations of earlier agents’ behavior, and rationally update their beliefs about feasible actions. When old, they offer contracts that deter only the harmful actions they deem sufficiently likely to exist; their agents then cheat if it is both feasible and undeterred by the contract, i.e., if there is a loophole.

In a loophole equilibrium, principals who observe cheating close loopholes when they offer contracts. But loophole-free contracts deter all cheating, thereby conveying little information about feasible actions to the next generation of principals; they rationally update in favor of the view that cheating is unlikely or infeasible and may therefore offer prove to be contracts with loopholes. The result is cycling of contract types that alternately deter and encourage undesired behavior, yielding heterogeneity across principals. There are also “bureaucratic” equilibria in which contracts deter behavior that is actually infeasible.

Depending on whether principals sample concurrently or historically, population dynamics in a loophole equilibrium may display aggregate cycling or convergence to a stationary, nondegenerate distribution of contracts or rules.

Keywords: Bayesian learning, contract theory, aggregate dynamics; organizational change, organizational heterogeneity, multi-task principal agent model; gaming incentive systems.

1 Introduction

The State of California implements a statute requiring motorcyclists to “wear helmets” when riding on public roads. The riders soon comply by strapping the helmets to their

^{*}We are grateful for discussion with Patrick Bolton, Wouter Dessein, Bob Gibbons, Bengt Holmstrom, Meg Meyer, Nicola Persico, Julie Wulf, Jeff Zweibel, and seminar participants at BU, ESSET Gerzensee, NYU, Columbia, ECARES, Amsterdam, Erasmus University, the Harvard/MIT Organizations Workshop, PSE, Namur, Oxford, UCL, Cambridge, the NBER Organizational Economics Working Group Meeting, the Barcelona ThRed Meeting, Harvard/MIT Applied Theory Workshop, the NGO Workshop at LSE, Princeton, and HEC Montreal.

[†]PSE, UCL, and CEPR; jehiel@pse.ens.fr

[‡]Boston University and CEPR; anewman@bu.edu

arms.

An electrical equipment manufacturer equips its office typewriters with counters that register the number of times the keys are operated, and pays its stenographers by the keystroke. It is soon noticed that one typist is earning much more than any of the others: she takes lunch at her desk, “using one hand for eating and the other for punching the most convenient key on the typewriter as fast as she can.”

Division managers at a large food products company receive bonuses only when earnings increase from the prior year. They respond by manipulating the timing of shipments, falsifying dates on sales invoices, and accruing expenses in the wrong period.

These are instances of *loopholes*: a principal offers an incentive scheme to an agent, who then “complies” with the letter of the scheme in ways that harm the principal, but cannot be deterred by punishments since they are not violations of the contract. Given the information that this kind of cheating is possible, the principal would redesign the contract.

Indeed, the helmet law was rewritten to specify (and more important, enforce) that helmets be worn on heads.¹ The stenographers’ piece rate scheme was withdrawn by their company, Lincoln Electric, and replaced by fixed wages and close supervision (Berg and Fast, 1975). And the food products company, H.J. Heinz, implemented a costly set of accounting and supervisory procedures and modified the company’s reporting structure to deal with the problems with its division managers (Goodpaster and Post, 1981).

Examples like these, in which firms and other organizations attempt to boost profits or other goals by introducing incentive schemes that appear to backfire, pepper the management literature.² Economists favor them as pedagogical tools for illustrating the pitfalls in designing of incentive systems, or simply to prove that people really do respond to incentives. There is no shortage of economic analysis showing why rational agents would respond the way they have (Holmström-Milgrom, 1991; Baker, 1992), and plenty of empirical evidence illustrating just how real people manipulate the schemes that govern their workplaces. Piece rate schemes such as that for the stenographers have well known quality problems;³ timing of sales and other performance quotas such as those at Heinz have been noted in a variety of settings and widely studied (Chevalier and Ellison, 1997; Oyer, 1998; Courty and Maschke, 2004, 2008; Frank and Obloj, 2009; healthcare providers have received particular scrutiny recently: see for example Bevan and Hood, 2004, 2006; Propper et al., 2008; Gravelle, Sutton, and Ma, 2010). And motorcyclists around the world continue to comply with their local ordinances by strapping helmets to their arms. In other words, the same mistakes seem to be made many times over.

¹For instance the California Vehicle Code now specifies that “... ‘wear a safety helmet’ or ‘wearing a safety helmet’ means having a safety helmet...on the person’s head that is fastened with the helmet straps and that is of a size that fits the wearing person’s head securely without excessive lateral or vertical movement.” (Division 12, Chapter 5, Article 7, Section 27803.

²The classic reference on this is Kerr (1975), “On the Folly of Rewarding A while Hoping for B,” which also emphasizes that loopholes are abundant in many walks of life besides firms. Although our analysis is confined to formal contracts in simple production settings, much of what we say can be applied, with appropriate modification, to “relational” contracts, as well as to incentive systems more broadly such as to legal or regulatory regimes.

³See for instance Stiglitz (1975), though the problems had been documented in the sociology literature and much earlier, e.g. Whyte et al. (1955) and were apparent to Frederick Taylor and his followers.

Of course, firms and other principals respond. Once the harmful action is discovered, the firm will attempt to deter it in future by revising its rules or contracts.⁴ Sometimes though, employees find new ways to harm the principal, leading to new contracts. And so on.

These phenomena do raise the question why so many well-intentioned and motivated principals (such as profit-maximizing firms) would provide perverse incentive schemes in the first place. Indeed, they sit somewhat uneasily with the standard contracting setting, in which the principal designs a contract to induce her agent to behave in desirable ways. There, in equilibrium, nothing “unexpected” happens: the agent chooses the action prescribed by the contract, and every verifiable outcome of the relationship is an anticipated one, in the support of the distribution of outcomes corresponding to the prescribed action. Looking across similar economic environments, contracts chosen are similar, and unless some feature of the environment changes, there is no reason for the form of the contract to change over time. From this perspective, the undesired and apparently unexpected behavior by agents that sometimes happens in the real world, together with subsequent revision of the incentive schemes that govern them, is something of an anomaly.

A simple resolution to the paradox is to acknowledge that real-world, principals, rational as they may be, must sometimes design incentive schemes with only incomplete information about their environments. This certainly would account for the occasional loophole. But subsequent agent behavior would reveal the oversight to the principal, who could then close it, thereby deterring the undesired behavior in the future. In other words, principals can *learn* about the environment by observing agent behavior.

One might think such a process would converge to a state of affairs in which principals know pretty much everything about the environment, or at least to one in which contracts they offer deter harmful behavior. But a key feature of learning about agency environments, which differs from the “technological” environments that have received the bulk of attention, is that *the optimal incentive scheme inhibits learning*. Because it deters all undesired actions, an observer will never see that such behavior is feasible, or the damage it can do. So he will rationally reduce his belief in its feasibility or harmfulness. As a result, to reduce the costs of incentive schemes that deter behavior which in his view is unlikely to happen anyway, he may offer a scheme that turns out to have loopholes. Unfortunately for him, his agent will cheat, but now others will learn from his mistake and close the loophole. The process begins anew.

Thus loopholes, though damaging to principals who offer them, generate a *positive learning externality*, which in the simplest settings may lead to cycling through a limited number of possible incentive schemes. More generally, they may be viewed as an engine of the “constancy of change” in organizations.

The purpose of this paper is to explore some of the implications of these observations in a simple model of contracting and learning. The basic setting we consider in this paper is an overlapping-generations model in which at any date there is a large number of ex-ante identical principals and agents. A typical agent can potentially take different types of actions, only one of which benefits her principal; though he is aware of all the

⁴Thus we distinguish the loophole case from the one in which occasional cheating is anticipated and tolerated simply because it is too costly to deter.

logically possible actions an agent might take, he is initially uncertain as to which actions are possible.

Before offering a contract to an agent, the principal spends an apprenticeship observing a random selection of other principal-agent relationships. We consider a number of assumptions about what is observable during these encounters, but the simplest is to imagine that the agents can be observed either to be “cheating” or “not cheating,” but not (all of) the constraints – in particular, the details of the contracts – that govern their behavior.

If cheating is observed, then it is obviously known to be possible and the principal will offer a contract that deters it. If it is not observed, then the principal is not sure whether that is because it is impossible or because he happened to observe only agents who were governed by contracts that deter cheating. Indeed, he will update his beliefs in favor of the scenario that cheating is impossible, conditional on this observation. He then designs a contract optimally given those beliefs. This setup captures what appears to be an important feature of different managerial “styles,” namely that they are largely conditioned by early career experiences (Schoar, 2007).

Other than the limitations on what is observed and how many earlier relationships principals may observe, all assumptions are standard: principals and agents are Bayesian rational, there is a common prior about the true state and strategies are assumed to be known in equilibrium.⁵

We obtain the following findings:

- A fully informed contract cannot emerge in equilibrium: that is, sometimes principals don’t choose the optimal contract given the actual set of feasible actions.
- There are two basic types of equilibria: (1) In *bureaucratic* equilibria, principals offer contracts that deter behaviors that cannot actually happen; (2) In *loopholey* equilibria, some of the contracts offered have “loopholes,” and agents will sometimes “cheat,” i.e., take actions not recommended by the contract.
- In *loopholey* equilibria, when cheating is feasible, contracts offered by a lineage of principals change over time without converging — since the number of possible optimal contracts is finite, this means that there is always “cycling” (usually with stochastic frequency) at the lineage level. There will also be heterogeneity in the contract offers across otherwise identical lineages.
- In these equilibria, as the sample size grows without bound, a vanishing fraction of the population offer *loopholey* contracts, while the rest offer “tight” ones that deter cheating, so that the outcome approximates the “standard” case in which principals are fully informed about the environment.
- From the principals’ point of view, *loopholey* equilibria yield higher average payoffs than *bureaucratic* equilibria whenever the two kinds of equilibria coexist.

⁵For the simplest model, for principals to act they need only know the strategies and not the priors. So the common prior assumption can be relaxed.

- Cycling may also occur at the population level – most (in the simplest case, *all*) principals may simultaneously offer loopholey contracts, followed by most offering tight ones, followed by mostly loopholes, etc. The nature of these cycles depends on the temporal structure of observation: if principals look only concurrently at other relationships, rather than observing the outcomes of historical relationships, the aggregate cycles are stable; if they look historically, they tend to converge to a steady state in which a constant fraction of the population offers loopholey contracts.
- In an extension of the basic model, principals may pick new contracts that are “more complicated” (and more costly) than earlier ones, in that they explicitly prohibit more types of actions. Sufficiently costly contracts may then be replaced by simple ones: a period of incremental improvements is then followed by a major reform, and the cycle begins anew.
- In another extension, principals can choose whether to look concurrently or historically. Variation in the fraction of the population offering different contracts over time favors historical sampling, but if the world changes at all, even slowly, concurrent sampling gives more relevant information for next period’s decision. In general there will be a mix of concurrent and historic samplers, with aggregate limit cycles that move faster than the underlying environment.

1.1 Literature

This paper is related to several theoretical literatures in economics. As we have indicated, the setting we are interested in is one in which agents may engage in multiple “tasks,” with varying benefits for the principal, but theoretical attention there (Holmstrom-Milgrom, 1991; Baker, 1992) has mainly been devoted to studying the optimal design of incentive schemes in such environments, in particular the deterrence of gaming (Ederer, Holden, Meyer, 2010), rather than to explain why gaming occurs in the first place. Most close in spirit is a series of recent papers studying aspects of contracting under unforeseen contingencies and/or cognitive constraints (e.g. Tirole, 2009; Bolton and Faure-Grimaud, 2005). Part of the distinction here is how we interpret “cognitive limitations”: we take it as a fact of life that the signal structure we are able to process is coarse relative to the world we are trying to understand (and acquiring a finer one would be costly), but this need not be inconsistent with Bayesian rationality; the other papers interpret the unwillingness of agents to get full information about the state of the world as bounded rationality. More significant is the difference in focus: we are interested in how contracts evolve as a result of observational limitations rather than on the static properties of contracts that are chosen by cognitively limited agents. Ellison and Holden (2009) seeks to understand how rules are developed, but it uses a different framework, one in which principals are perfectly informed but have difficulty communicating to agents who always obediently follow instructions. It does admit increasing complexity or refinement of rules but the tendency is always toward greater efficiency over time, with eventual convergence to a stable set of rules. There is also a literature on strategic incompleteness (Spier 1992, Bernheim and Whinston, 1998; Chung and Fortnow, 2007), but there principals deliberately write incomplete contracts

with perfect anticipation of the behavior that will result and again there is no evolution of contract forms.

There is also a very large "social learning" literature in which individuals observe one aspect or another of other individuals' behavior and use those observations to update their beliefs about the optimal course of action. Examples include papers on social multi-armed bandits (Bolton and Harris, 1999) or game theoretic or economic models in which signals are of limited dimension relative to other individuals' information (Ellison and Fudenberg, 1993; Banerjee and Fudenberg, 2005; Piketty 1995). The difficulty of making correct inferences from coarse signals is also at the heart of the information cascades and herding literature (Banerjee, 1992; Bikchandani, Hirshleifer and Welch, 1992); Scharfstein and Stein (1990) uses a somewhat different logic to derive herding behavior in a setting in which agency rather than pure inference problems are crucial. None of these papers has the feature that choosing the "right" decision for the state (in our case, the optimal contract that deters the entire set of harmful behaviors that exist at least cost) hides information from others, and thus while they will often have convergence to the wrong action, they do not display the constancy of change (e.g. cycling) that our model displays.

Our model is also related to the literature on learning in games and in particular the idea of self-confirming equilibrium (Fudenberg and Levine, 1998) in which the observation made by subjects does not allow them to pin down a unique interpretation of others' strategies in a dynamic learning process (in particular, about what strategies would be off the equilibrium path). Our model differs from that literature in that principals' beliefs derive from Bayesian updating and a common prior assumption, and are not just assumed to be subjectively consistent with the observation.⁶

Our model bears some connection with papers on imperfect recall and/or coarse information processing (Piccione and Rubinstein, 1997; Mullainathan, 2000; Wilson, 2002; Jehiel, 2005; Baliga and Ely 2009) though formally our agents don't forget anything they observe, they simply have no incentive to remind anyone else. Moreover, that literature doesn't look at implications for contracting or organizational evolution.

Finally, there is a law and economics literature (Gennaioli and Shleifer, 2007a, 2007b) on the evolution of the common law that bears some connection to what we are doing. We briefly discuss that in the Conclusion.

2 Model

2.1 Technology and Preferences

The economy lasts forever, with time measured in discrete periods. In each period, there is a continuum with unit measure of principals and agents who enter into one-period production relationships. All individuals live for two periods, youth and adulthood. Principals only

⁶Moreover, in our setting, even on the equilibrium path, some principals may propose contracts that trigger harmful behavior, and this is because we assume there is not a public record of what happens in past interactions. At the same time, the analysis reveals the value of such a record: of course, it is likely to be costly to maintain.

offer contracts in adulthood; in their youth they serve an “apprenticeship,” in which they watch and learn. Agents are idle in youth and work in adulthood.

At any time t , all (risk-neutral) principals own identical projects. In the simplest case, there are two logically possible dimensions of action that an agent might take. Following Kerr, we refer to them as B and A (this can be generalized to B plus A_1, \dots, A_{n-1} ; we will refer to n -dimensional worlds and states i corresponding to the 2^{n-1} sets of A -dimensions that exist).

Now, either all projects are in fact one-dimensional or they are two-dimensional. Principals don’t know what state they are in. They do know that the world is either in one state or the other for all time, and they have an (exogenous) prior μ^0 , the probability of the “good state,” in which only the B action is available. This belief will be updated based on observations of a sample of previous principal-agent relationships.⁷ A project yields $R > 0$ to the principal with probability e_B , 0 with probability $1 - e_B$, where e_B , which is not verifiable, is the effort exerted in the B dimension.

All agents are identical: they are risk neutral in income, have limited liability, and an outside option equal to zero. Agents preferences are given by (y is income)

$$y + \alpha e_A - \frac{1}{2}(e_B + e_A)^2,$$

where $\alpha \geq 0$ is an agent’s private benefit of engaging in action A . Of course in the state of the world in which A -actions exist don’t exist, we necessarily have $e_A \equiv 0$, while in the “bad” state in which they do, $e_A \in [0, 1]$.

We assume that agents don’t learn the state until after contracting, when the production period begins; indeed we may assume they learn only if they happen to succeed at taking action A . The parameters satisfy

$$1 > R > \alpha > R/2.$$

This assumption ensures that it is both efficient and costly to deter the A action. A first-best allocation in which effort is verifiable would have $e_B = R$ and $e_A = 0$.

2.2 Full-Information Second-best Contracts

In general, principals will be uncertain about the dimensionality of the world they are operating in. One might imagine then that the optimal contract would be some sort of

⁷Formally we treat this prior as common among all young principals and exogenous. However, the reader may want to keep in mind the following derivation, which will be pertinent in our discussion in Section 4.0.8. Suppose that over time, the economy may be in either state. However transition times between states are very long, so for all intents and purposes a principal can be quite certain that he spends his whole life in one state. Specifically, if it is in state 0, the economy transitions out of it into state 1 with probability ϵ_0 and from 1 to 0 with probability ϵ_1 . The economy therefore spends $\mu^0 = \frac{\epsilon_1}{\epsilon_0 + \epsilon_1}$ of the time in state 0 with average duration of a spell equal to $1/\epsilon_0$. We think of the ϵ_i as very small so that principals (and we) may neglect them in their updating calculations. The generalization to any finite number of states is straightforward.

The common prior assumption can be substantially relaxed without changing many of our results; indeed in some cases they are “easier” to obtain. We leave discussion of that case to a future draft.

“compromise” that would perform well on average, where the average is computed with respect to the principal’s belief about the world he is in.

In our model, however, the optimal contract with uncertainty about the state of the world is always equal to one of the certainty optima. This serves to highlight the effects we are interested in. In more general settings, contracts and efforts taken by agents might adjust incrementally to principals’ beliefs, but it would still be true that principals with strong beliefs in the good state would find agents exerting higher than expected levels of A effort, and principals with strong beliefs in the bad state would offer contracts that lead to low levels of A effort; insofar as observing A effort is more likely when there is more of it, our results will stand.

We compute the full-information optimal contracts for our model here.

In the 1-dimensional world (state 0), the optimal second best (effort not verifiable) contract in the world pays w_0 if R , 0 otherwise, where w_0 solves

$$\begin{aligned} & \max_w e_B(R - w) \\ \text{s.t. } e_B &= \arg \max_w e_B w - \frac{1}{2}(e_B)^2; \end{aligned}$$

since the agent sets $e_B = w$, maximizing $w(R - w)$ yields

$$w_0 = R/2$$

In a 2-dimensional world (state 1), there are two ways to encourage agents to choose the productive effort. The first is simply to pay a sufficiently high incentive wage w_1 if R and 0 otherwise, where w_1 solves

$$\begin{aligned} & \max_w e_B(R - w) \\ \text{s.t. } e_B &= \arg \max_w e_B w + e_A \alpha - \frac{1}{2}(e_B + e_A)^2. \end{aligned}$$

Since the agent sets $e_B = 0$, $e_A = \alpha$ if $w < \alpha$ and $e_B = w$, $e_A = 0$ if $w \geq \alpha$, the principal gets 0 if $w < \alpha$; $w(R - w)$ is decreasing for $w \geq \alpha$ (since $\alpha > R/2$), so

$$w_1 = \alpha > w_0$$

Hence deterring the A effort is costly, and it may be preferable to use a second method of deterrence, namely to add a clause to the contract that prohibits the A action. Such a clause will be costly to enforce, because it requires monitoring to make it verifiable (in some cases the cost may assume the form of foregone production benefits – employees socializing in the coffee room may occasionally communicate information that might enhance productivity and benefit the principal – in short the two actions may not be perfect production substitutes). We model this by supposing that the principal incurs a cost c for each enforceable clause he adds to the contract. Then a clause contract consisting of the incentive wage w_0 and a clause would pay the principal $\frac{R^2}{4} - c$.⁸

⁸For simplicity, we rule out random monitoring, wherein the principal incurs the cost pc for detecting

There is thus an optimal contract C_0 for state 0 in which cheating is not possible and a different optimal contract C_1 for state 1 in which cheating is possible. Denote by π_i the optimal profit in state i ; thus $\pi_0 = R^2/4$ and $\pi_1 = \max\{\frac{R^2}{4} - c, \alpha(R - \alpha)\}$.

2.3 Contracting with Principal Uncertainty about Environment

Let μ_0 be the Principal's (posterior) belief that the world is 1-dimensional; he gets $\mu_0\pi_0$ by offering the C_0 contract and π_1 with C_1 . It is easy to see that he does not gain by hedging with some sort of compromise contract: any optimal contract will either induce the agent to be productive (choose B effort) in both states or it will induce him to be productive only in state 0 (it cannot be productive in state 1 without also being productive in state 0). If the contract is productive in both states, it is dominated by C_1 . If it is productive only in state 0, it is dominated by C_0 : offering less than w_0 yields him less than $R^2/4$ in state 0 and zero in state 1, since it is not incentive compatible there; and offering more than w_1 will always yield him less than offering w_1 , which in turn is (weakly) dominated by contract C_1 . Finally, offering $w \in (w_0, w_1)$ is weakly dominated by offering w_0 . Note the argument does not depend on whether C_1 is the wage or clause contract.⁹

We thus have a simple characterization of how the contract offered will depend on the Principal's belief about the dimensionality of the agent's action set:

Lemma 1 *The Principal offers C_0 if $\mu_0 > \mu_0^* = \pi_1/\pi_0$ and C_1 otherwise.*

Note that μ_0^* is small if c (or α , in case the high wage contract is optimal in state 1) is "large" – close to $R^2/4$ (α is close to R). In other words, if deterrence is costly, the principal does not need a strong belief in state 0 in order to offer the contract suited for it. By the same token, when c (or α) is small, the principal might as well deter, even if he is fairly certain that it is not necessary.

We should emphasize that our basic results do not depend on the discreteness of the principal's optimal action set (though the simplicity of computations surely does). Whatever contract that would be offered would either deter action A or it would not, and it is the action rather than the contract that we are taking to be observed by young principals.

3 Learning

Of course, the immediate question is where does μ_0 come from? Beliefs are endogenous in our model: we suppose a principal learns (imperfectly, in general) about the state of the world based on observations only of what other *agents* do. In particular, principals

A effort with probability p while paying a wage of $(1 - p)\alpha$ – little substantive would change if we allowed for it.

⁹In fact, if one thinks of the state as the agent's preference that he learns after contracting, then one can use Myerson's (1982) framework to show that these remain the optimal contracts even if agents's were allowed to report their types once they learn them. Separating agent types profitably is difficult because both "good" ($\alpha = 0$) and "bad" ($\alpha > R/2$) agents like high wages, and good agents mind clauses less than bad agents. So in a separating menu, the good types would get more clauses, which is counterproductive. Details are omitted from this draft.

do not base their beliefs on “stories” told by other principals (though we could allow for some relaxation of this assumption without changing things very much). The main point is that learning will be based on a limited (compared to the complexity of the world) set of observations. This is the extent of “bounded rationality” or coarse information processing in our model.

The sequence of activity during a principal’s life time is as follows.

- **Principals have a prior μ^0** that the world is one dimensional, i.e. that there is no way to cheat other than to withhold B -effort. This prior is exogenous, perhaps, as we mentioned, a long historical average.
- **Each apprentice (young principal) gets k draws about past interactions** that (almost certainly) took place in the same state of the world that he is in.
- **While an apprentice, principal observes which actions are taken by agents governed by the earlier k relationships she sees** (as well as success/failure of the projects, i.e. she observes only whether $e_B > 0$ and $e_A > 0$, but not their magnitudes).
- **Principal forms a posterior** based on observations, knowing what fraction of the population are offering each type of contract in each state of the world (the equilibrium assumption).
- **Principal offers contract to his agent, who then takes an action,** possibly while another apprentice watches.

For now we assume that the apprentice principals observe very little. They don’t observe the size of wage paid or the contract itself, or the magnitude of e_B (which would reveal the wage). In fact it would make no difference if they observe the magnitude of e_A ; even observing e_B would not matter if one extends the model slightly – see the section on observing contracts).

Our calculations assume that all relationships the principal sees are in the same state as her own. The results we get will not depend on this provided transition probabilities between states are small.

We are otherwise making conventional assumptions about rationality and equilibrium play: principals update beliefs using Bayes’ rule, and know the strategies other principals and agents are playing. In particular they know the map from states to distributions of contract offers (but not, of course, the actual distribution prevailing at the time).

The equilibria we study are all symmetric and stationary: all agents use the same strategies, and those strategies do not vary over time.

3.0.1 Impossibility of Fully-Informed Equilibria

The first observation is that the standard principal-agent setting, in which the principal is always fully informed about the set of action the agent is capable of taking and designs his incentive scheme accordingly, cannot emerge in equilibrium. The reason is very simple: a principal’s strategy must be a map from what he observes to a contract offer, and in the

standard setting, the principal would always observe the same thing, no matter how large his (finite) sample. Indeed, if fully informed, he would offer C_0 in the good state and C_1 in the bad, as would everyone else, and action A would never be taken by any agent. But since a principal’s sample would then look the same regardless of the state he is in, he must offer the *same* contract in both states (be it C_0 or C_1), which is a contradiction.

What then *can* happen in equilibrium? Broadly speaking, the possible outcomes will assume one of two forms, “bureaucratic” and “loopholey.”

3.0.2 Stationary Equilibria

We consider first steady-state equilibria of our model, that is situations in which the distribution of contracts that are offered within a state of the world is constant over the time that state persists (this distribution is simply the fraction x offering C_0 ; the rest offer C_1). We shall refer to principal beliefs in terms of the probability of the one-dimensional world, which is the one they would prefer; (given the optimal contracts, agents have the opposite preference).

Assume that each principal views k relationships, where k is exogenous.¹⁰ Since we are focused on stationary equilibria, it does not matter whether the observations the principal gets are from one period (say, his youth) or from relationships that occurred deeper in the past.

Proposition 2 (*Bureaucratic Equilibrium with Low Priors*): *Assume that $\mu^0 < \mu_0^*$. For any k , the following is a stationary equilibrium: all principals offer C_1 ; the sample always reveals that only task B was performed.*

That this is an equilibrium should be clear: with everyone offering C_1 in both states of the world, no agent will ever be observed cheating in either state, no matter what the sample size. Thus, what is observed cannot reveal any information about the state, and the posterior equals the prior; since this is less than μ_0^* , it is optimal to offer C_1 . Notice that in this stationary equilibrium, when the world is one dimensional, contracts are *bureaucratic* in the sense of providing the principals with more protection than they need, and this is costly. Of course no principal *knows* her contract has redundant protection, he just believes it is necessary. These equilibria bear a close resemblance to many “incomplete learning” equilibria in the literature, where agent’s beliefs eventually settle down to a point where they do not wish to experiment, even though doing so would reveal that they were behaving suboptimally.

This is not the only stationary equilibrium compatible with a prior that is less than μ_0^* . There are others in which the C_0 contract is offered, at least some of the time, even in state 1.

To see this, suppose $k = 1$. We have:

Proposition 3 (*Loophole Equilibrium with High Priors*): *Assume that $\mu^0 > \frac{\mu_0^*}{2 - \mu_0^*}$ and each apprentice observes one prior relationship ($k = 1$). The following is a steady state*

¹⁰We have in mind that each observation is costly; one could model a principal’s choice of k explicitly, but that case is best left for a future paper. See the discussion in Section 4.0.8.

equilibrium: Upon observing no effort on task A, a principal offers C_0 ; upon observing effort on task A, a principal offers C_1 ; in the 0-state, only C_0 is offered; in the 1-state, half of the contracts are C_0 and half are C_1 .

If the entire population use the strategy described in the proposition, then in state 0, no one ever sees A, and immediately everyone offers C_1 . In state 1, if half offer C_0 in state 1 and half don't, a randomly chosen relationship will have been governed by a C_0 contract with probability $\frac{1}{2}$ and those agents will exert A effort. The observer will offer C_1 when it is his turn. The other half of the young principals will observe no A, and will offer C_0 ; thus once again half the contracts are C_0 . To verify that this strategy is optimal, note that those who don't see A but know that in state 1 half the population offer C_1 forms the posterior μ_0 , where

$$\begin{aligned} \mu_0 &= \frac{\Pr(\text{Observe no A} \mid 0)\Pr(0)}{\Pr(\text{Observe no A} \mid 0)\Pr(0) + \Pr(\text{Observe no A} \mid 1)\Pr(1)} \\ &= \frac{1 \cdot \mu^0}{1 \cdot \mu^0 + \frac{1}{2} \cdot (1 - \mu^0)} > \mu_0^* \iff \mu^0 > \frac{\mu_0^*}{2 - \mu_0^*}, \end{aligned}$$

so that offering C_0 is optimal. The equilibrium is a “loophole equilibrium” in that there are states of the world in which some of the contracts being offered are suboptimal for that state and do not deter cheating behavior.

As simple as this version of a loophole equilibrium is, it contains several features of interest in the study of organization, particularly if we focus on the “bad” state. First, there is *heterogeneity of fundamentally similar organizations*: the contracts in use differ only because the principals who offer them have had different “experiences” (c.f. Schoar, 2007) in terms of what agent behavior they observed. Moreover, there is a *constancy of organizational change*: within a lineage, the contract offered in one period will differ (perhaps stochastically) from the one offered in the next.¹¹

3.0.3 Welfare

Loopholes, while inconvenient for the principal who offers them, are not all bad. They do provide a *positive externality*: apprentice principals are able to learn from other principals’ mistakes. It is worth contrasting the loophole equilibrium, which in state 1 is suboptimal, with the loophole-free equilibrium in Proposition 2. In that case, everyone always writes a tight contract, but because of that there is no possibility of observing that cheating is infeasible, and so it is suboptimal in state 0. In fact, when both equilibria exist (as they do when $\mu_0^* > \mu^0 > \frac{\mu_0^*}{2 - \mu_0^*}$), the loophole is better on average for principals than the stationary, loophole-free equilibrium:

Proposition 4 *Assume that $\mu_0^* > \mu^0 > \frac{\mu_0^*}{2 - \mu_0^*}$. The stationary loophole equilibrium always*

¹¹If a principal observes what happened in his predecessor’s firm, he will do the opposite, his successor will do as his predecessor did, etc. If principals draw randomly from the whole population, the evolution occurs stochastically at the lineage level. Apart from this detail, the stationary equilibria do not depend on whether the sample is supposed to within or across lineages.

delivers a higher expected payoff to the principals than the stationary loophole-free equilibrium

The condition for the welfare superiority of the loophole equilibrium is precisely the one that allows it to exist: without loopholes the payoff is always π_1 ; with loopholes, the expected payoff is $\mu^0 \pi_0 + \frac{1}{2} (1 - \mu^0) \pi_1$, which exceeds π_1 whenever $\mu^0 > \frac{\mu_0^*}{2 - \mu_0^*}$. The idea is that it may be (socially) better to make mistakes sometimes (even half the time, as happens in state 1 of the loophole equilibrium) and learn than to persistently do the wrong thing (state 0 of the bureaucratic equilibrium).

Of course, depending on what alternatives one imagines a social planner would have at his disposal, it is not straightforward to assess whether the loophole equilibrium is optimal: ideally one would have a very small fraction of the population “experiment” by offering the contract that is not ex-ante optimal, and convey the results to the rest. But the costs of this information transmission are possibly quite high (after all, this is implicit in the assumption that k is finite), not to mention the costs imposed on the principals who conduct the experiment. Indeed, with the same instruments, the planner could improve upon the bureaucratic equilibrium as well: forcing a principal to offer a C_0 contract would reveal the state. The point here is simply that loopholes allow agents to convey useful information, while agent behavior under tightly written contracts tells the rest of us very little about the world.

Note that the principals are better off in the loophole equilibrium than in the bureaucratic equilibrium regardless of whether the optimal contract for state 1 is the clause contract or the high wage. Agents, however, will feel differently about the two equilibria depending on which contract is optimal in state 1. If the high wage is optimal, agents prefer getting that all the time (bureaucratic case) to getting it only (half the time) in state 1. But if the clause contract is optimal, they prefer the loophole equilibrium, because they receive a low wage no matter the state in the bureaucratic equilibrium, while benefiting from occasional cheating in the loophole equilibrium (the wage is the same in both equilibria). Thus, the loophole equilibria may actually ex-ante Pareto dominate.

3.0.4 Larger samples

Principals are unlikely to cast their lot entirely with one mentor and will instead get their information from several sources. In terms of our model, the sample size could be larger than 1. Indeed a natural question to ask is whether the outcome converges to the standard fully-informed case (in which the principal always offers the contract optimal for the state) as the sample gets large.

We have already seen that this need not be true if the prior is low (Proposition 2), since for any k there is an equilibrium in which only C_1 is offered every period in both states, which as we remarked, is what one might expect from the standard incomplete-learning results familiar in the social learning and multi armed bandit literatures.

However, it is also true that for any prior, there is a sufficiently large sample size for which there is a stationary loophole equilibrium (i.e. one in which a positive measure of principals are offering C_0 in state 1). Stationarity will require that the set of unfortunates who offer loophole-laden contracts in state 1 is very small, because with large sample sizes,

they are still very likely to be observed by future principals, for whom the state is therefore revealed. This means that nearly everyone will be offering C_1 in state 1, while everyone offers C_0 in state 0. Thus the equilibrium approximates the standard full-information case.

Let \bar{x}_k be the unique solution in $[0, 1]$ to

$$x = (1 - x)^k$$

We note that \bar{x}_k decreases with k and tends to 0 as k tends to infinity. If \bar{x}_k is the proportion of contracts in the 1-state leading to positive effort on task A (i.e., C_0 contracts) the probability that a k sample includes no observation of effort A is $(1 - \bar{x}_k)^k = \bar{x}_k$, and the posterior after observing no A effort in the k -sample that the state is B would thus be $\frac{\mu^0}{\mu^0 + \bar{x}_k(1 - \mu^0)}$. For someone to be willing to offer a such a contract, we need this posterior to exceed μ_0^* . But for k large enough (therefore \bar{x}_k small enough) we can make this expression as close to unity as we like, even if μ^0 is small.

To summarize, we have :

Proposition 5 *For all $\mu^0 > 0$ there exists a \bar{k} such that $k \geq \bar{k}$ implies $\frac{\mu^0}{\mu^0 + \bar{x}_k(1 - \mu^0)} > \mu_0^*$. For such k sample size the following is a steady state equilibrium: Upon observing no effort on task A , a principal chooses C_0 ; upon observing effort on task A , a principal offers C_1 . As k tends to infinity, this equilibrium approaches the first-best situation in which the state is perfectly observed by principals.*

This proposition underscores once again the positive nature of the loophole externality: an arbitrarily small fraction of the population can perform the social experiment (offering C_0) on behalf of everyone else, provided only that everyone else is sufficiently likely to observe them. It also underscores that any notion of optimality of equilibrium will depend sensitively on what the sampling structure (size of k) is. In particular, if very large values of k are costly to sustain, it may be (constrained) optimal for significant portions of the population to offer loophole-laden contracts. We shall discuss this further below in the section on population dynamics.

3.0.5 Observing Contracts

So far we have been assuming that the information available to young principals is limited to observation of the dimension of agents' effort choices. In particular the contracts they are operating under are not observed. This raises the question of what happens if they can observe more.¹²

Suppose then that principals can observe previous contracts as well as effort types. Possible observations are then $(C_0, \sim A)$ and $(C_1, \sim A)$ in state 0 and (C_0, A) and $(C_1, \sim A)$ in state 1. If A is part of any observation, the state is revealed. It cannot be part of a (symmetric) equilibrium for principals to offer C_0 if $(C_0, \sim A)$ or $(C_1, \sim A)$ is observed and to offer C_1 otherwise: for the only way an earlier principal would offer C_1 according to this

¹²While there are good reasons in general why a principal might want to prevent a young one from observing all or part his agent's contract (e.g., Jehiel, 2010), it is worth asking what happens when we relax the assumption.

strategy is in response to having seen (C_0, A) , which reveals the state is 1. But observing C_1 would then tell the current apprentice that the state is 1, and he would offer C_1 and not C_0 , a contradiction. In fact, the only equilibrium strategy involving possibly different choices of contracts will be to offer C_0 if $(C_0, \sim A)$ is observed and C_1 otherwise. But then, everyone will offer the optimal contract for the state. If the prior is sufficiently low ($\mu^0 < \mu_0^*$), there is also the bureaucratic equilibrium of Proposition 2, since under that scenario no information is transmitted by the observation. Whatever the equilibrium, in each state there is a single contract being offered, which would seem to undermine the main insights about the persistence of heterogeneous contracts in state 1.

Of course, if we suppose that individuals are capable of processing more information than we have allowed so far, we must also recognize that the world is also far more complex than what we have discussed. In keeping with the precept that what people can observe and process must surely be small compared to the complexity of the world they are trying to understand, we allow principals to observe more, but we also give them a more complicated inference to make. Suppose instead that C_1 contracts might be offered for several reasons besides the possibility of cheating. For instance a clause prohibiting congregating in the coffee room might be there because there is asbestos in the ceiling as much as because it encourages nonproductive chit-chatting among workers. Contracts list provisions without explaining their purpose. So an apprentice will not be able to infer the payoff relevant state from observing the contract.

To see this formally, consider the case in which C_1 is the clause contract and suppose that β of exceptional principals are “born bureaucrats” with $c \leq 0$, while the rest are regular ($c > 0$) (alternatively, these exceptional principals are vulnerable to cheating independent of the state of the world that pertains to the rest of the population). The exceptional principals optimally use a clause regardless of the state, so nothing can be learned from observing them or the contracts they offer. As long as the apprentice isn’t sure whether the principal he observes is exceptional or not, he may still offer a C_0 contract even in state 1, since observing $(C_1, \sim A)$ doesn’t tell him for sure that he is in state 1 – he may simply have an exceptional principal. Formally, we have:

Proposition 6: Assume that $\frac{\beta\mu^0}{\beta\mu^0 + \frac{1-\mu^0}{2-\beta}} > \mu_0^*$. The following is a steady state with a $k = 1$ sample: Upon observing $(C_1, \sim A)$ or $(C_0, \sim A)$, a regular principal offers C_0 ; upon observing (C_0, A) she offers C_1 ; exceptional principals always offer C_1

The steady-state fraction of C_0 contracts will be $\frac{1-\beta}{2-\beta}$, close to the $\frac{1}{2}$ we had when contracts were not observed when β is small. Of course, regular principals need to be more “optimistic” to sustain loophole equilibria in this case, so in that sense, bureaucratic equilibrium is more likely.

A similar exercise can be done for the case in which C_1 is the high wage contract by supposing that exceptional principals offer higher wages because they are more profitable (of course, observing agents’ wages be less likely for an apprentice than observing clauses or work rules). Or we could depart from the common prior assumption and suppose they are principals whose priors are lower than those of the rest of the population. Either way, apprentices who see C_1 contracts will reason that the preventive measures they see being taken need not apply to them, and rationally they offer C_0 in order to avoid the expense

of deterrence.¹³

3.0.6 Summary

So far we have studied stationary equilibria in a very simple contracting environment. The main finding is that stationary loophole equilibria often exist in which lineages of principals switch (stochastically) from offering the optimal contract for state 1 to the “wrong” contract, and then back again. Looking across the population, there will be both types of contracts offered at each period. (In such equilibria, in state 0 everyone offers the optimal contract for that state.) This captures important features of the phenomena we are interested in: heterogeneity, misguided incentive schemes, and responses by patching with schemes that work, only for those schemes to contribute to their own undoing.

The fact that firms that offer tight contracts effectively release little information has interesting implications of the persistence of success and failure within organizations and their relation to managerial turnover. In the bad state, individual lineages might be successful for a while – if their principals happen to observe loophole firms in their apprenticeships and therefore offer the optimal contract for the state. If we assume that every apprentice observes his “father” (consistent with our assumptions, but not necessary), then every lineage that offers a loophole contract will correct this in the ensuing period. But if observations are entirely of other principals, the possibility of persistent loopholes within lineages arises. Interpreting lineages as organizations à la Crémer (1986), with the overlapping generations as turnover of managers, this suggests that organizations that recruit managers from outside, particularly from successful rivals, may be in for disappointment: those managers may be “naive” in the sense that they do not believe sufficiently in the possibility of cheating to offer tight contracts. From the organizational point of view, it may be better to recruit from within, or to even seek outside managers from failures.¹⁴

3.0.7 Increasing the number of states

As the dimension of the world increases, the state space rapidly grows large. In fact, since with n possible cheating dimensions, any subset (including the empty one) of those n dimensions might be feasible, the state space has 2^n elements. Rather than attempt an exhaustive analysis, we will merely explore some of the possibilities.

It will be particularly illuminating to consider a special case in which the state space is a restricted to contain the following $n + 1$ elements: \emptyset (i.e., cheating is not feasible), $\{1\}$, $\{1,2\}$, $\{1,2,3\}$, ..., $\{1,2,3,\dots,n\}$. The agents’ marginal utility for A_i is α_i and their cost function is $\frac{1}{2}(e_B + \sum_i^{n-1} e_{A_i})^2$. Each of the α_i is in $(R/2, R)$, so that it is both efficient and

¹³If vulnerability to cheating is assumed to be specific to the principal rather than the organization, then this reasoning can also apply *within* an organization: even a principal who observes her predecessor’s contract to be a C_1 may rationally infer that “times have changed,” that he, unlike herself, was vulnerable, and she should therefore be willing to offer a (possibly loophole) C_0 instead of another C_1 .

Of course, if the state really does change (perhaps slowly) over time, this kind of reasoning leading to loopholes would be even easier to sustain.

¹⁴This is meant more as empirical implication than policy advice, since there are other reasons why one might not want to recruit exclusively from failed firms.

costly to deter each type of cheating. We also assume that α_i are decreasing in i . Thus, agents will prefer to cheat first in dimension 1, then in dimension 2, etc. All of this is assumed common knowledge; only the actual state is unknown to the principals.

Principals may deter A_i (and $A_j, j > i$) by paying a success wage $w \geq \alpha_i$; the minimum wage that will deter all cheating is then α_1 . There would be no need for further analysis if this was only means of controlling cheating, since states 1 through n could be conjoined and treated as our previous state 1. Things become more interesting with clause contracts. One can now deter all cheating by using n clauses or by paying α_1 and using no clauses, but in general there will be an $m < n$ for which the state optimal contract for state $j < m$ is a j -clause contract with wage $R/2$ while the optimal contract for $j \geq m$ is a $(m - 1)$ -clause contract with wage α_m : at some point it's better to pay the agent a larger rent than to try to deter all conceivable misbehaviors with rules. This keeps the effective dimension of the problem down to m . Here we take $m = 2$.

There will now be three state optimal contracts, C_0 for the state in which no cheating is possible, C_1 in which only the first type of cheating is feasible, and finally C_2 when both types can happen, with corresponding profits $\pi_0 > \pi_1 > \pi_2$. Denote the weight the posterior beliefs place on states 0 and 1 by μ_0 and μ_1 ; the belief space is now cut into three regions in which posteriors must sit in order for each of the contracts to be offered.

As before, there will be several types of equilibria depending on parameters. Here we focus on a particular loophole equilibrium in which all three contracts are offered in state 2. In state 0 apprentices never observe A_1 or A_2 . In state 1, they may observe A_1 (if they draw a C_0 shop). Finally, in state 2 they may observe neither A_1 nor A_2 , A_1 but no A_2 , or A_2 . Letting x_i denote the frequency of C_0 state i , and y_i the frequency of C_1 in state i , we have the following

Proposition 6 *There exist priors (μ^0, μ^1) such that the following is a steady state: strategies are*

Offer C_0 if never observed A_1 or A_2

Offer C_1 if observed A_1 but never A_2

Offer C_2 if ever observed A_2 ,

and the outcome is: In state 0, only C_0 is offered; in state 1, C_0 and C_1 offered by different principals; in state 2, all three contracts are offered by some part of the population.

The proof is straightforward, and simply requires calculating steady-state values of (x_i, y_i) resulting from the strategies and the corresponding posteriors corresponding to the values of (x_i, y_i) to ensure optimality of the strategies. Details are omitted from this draft.

Within a lineage, there will continue to be (stochastic) “cycling” from one period to the next. Thus there will continue to be heterogeneity across organizations in states 1 and 2. Moreover, in state 2, if one assumes that each apprentice samples from his predecessor organization, there will be a progression of contract forms within each organization. For instance if $k = 1$, then C_0 will be followed by C_1 followed by C_2 : the contracts get increasingly “complex” (at least if they are all clause contracts); but after that they return to the “simple” C_0 form.

More interestingly, consider what happens when apprentices also sample from other organizations (so that $k > 1$). Then an apprentice who grew up in a C_0 shop but who

stumbles into only C_2 shops on his travels will close the loophole (in which A_1 was the agent’s response) by offering C_1 , thereby revealing another loophole. But an apprentice who happened to enter a C_1 will know about A_2 , and will offer a C_2 . Thus not only is there heterogeneity across fundamentally similar organizations, but *there will be heterogeneous responses by similar organizations facing similar problems*. This is certainly consistent with evidence that heterogeneity in “managerial fixed effects” plays a significant role in the differential performance of organizations (Bertrand and Schoar, 2003): in this case, the differences in the early experiences of principals lead to differences in beliefs and therefore in the decisions they make.

If principals don’t observe their predecessors, but only take over organizations from outside, then there is the possibility that the same organization will have persistent loopholes, since each period the new principals may not believe sufficiently in state 2 based on what they have observed. By contrast, those who come from within will always offer (possibly unsuccessful) solutions.¹⁵

Finally, more complexity in the state space allows for the possibility that observing *part* of a contract but not all of it. That is one might see a clause deterring A_1 but be unsure of the wage being paid, which would then make it impossible to distinguish a C_1 from

4 Population Dynamics

This still leaves open the question of whether these steady state equilibria are representative of the overall behavior of an economy in which loopholes are possible, something we would want to know both for descriptive and welfare analysis. In this section, we investigate whether the steady state is stable, or whether the aggregate behavior follows some other temporal pattern.

It turns out that the answer depends on what assumptions we make not only about sample size but more importantly on the *lag structure* of observations. We shall consider two distinct lag structures. In the first, which we call “Ford,” each principal takes all $k > 1$ observations from the period in which he is young. In the second, “Santayana,” case, the principals take one observation from each of the k periods prior to the one in which they offer a contract.¹⁶ In the latter case, the aggregate behavior of loophole equilibria eventually settle down to the steady states we described earlier, with (stochastic) cycling at the lineage level but not in aggregate. In the former, though, the stochastic cycling of individual lineages becomes coordinated into (deterministic) cycles of the aggregate.

Let x_t^s = fraction of the population offering C_0 at t in state s . We assume as before that μ^0 is constant over time and common for the population. We maintain the standard

¹⁵If the prior puts most weight on states 0 and 2, this progressive equilibrium disappears, but another, “over-reactive” equilibrium emerges. In it, principals who see no cheating offer C_0 , but principals who see A_2 or A_1 offer C_2 . In state 0, only C_0 is offered, while in states 1 and 2, some C_0 as well as some C_2 are offered. In particular, no C_1 is ever offered. Thus in state 1, the equilibrium is over reactive: when no cheating is observed, there is a roll back to no deterrence, but as soon as cheating is discovered, agents are overregulated by contracts that deter forms of cheating that are feasible, but also some forms that are not.

¹⁶George Santayana is the originator of the aphorism, “those who cannot remember the past are condemned to repeat it,” while Henry Ford was reported to have said, “history is more or less bunk.”

equilibrium assumptions, though it is worth emphasizing that away from a steady state this means in particular that individuals are aware of t and x_t^s .¹⁷ We begin with a

Lemma 7 *Consider the dynamical system*

$$x_{t+1} = F(x_t) \equiv (1 - x_t)^k, k > 1$$

There is a unique steady state $\bar{x}_k < 1/2$, which is unstable, and a limit cycle with basin of attraction $[0, 1] \setminus \{\bar{x}_k\}$.

The uniqueness of the steady state, along with the fact that it is less than $1/2$, has already been established. The presence of a period-2 cycle is easy: if $x_t = 1$, then $x_{t+1} = 0$, $x_{t+2} = 1$, etc. That this cycle is an attractor is proven in the Appendix.

Ford Apprentices born at t draw all k of their observations from principals active at t . Consider the following strategy. Those who observe A in their samples offer C_1 (they know the state is 1). Those who never observe A offer contract C_0 . This strategy is optimal whenever $\mu^0 / [\mu^0 + (1 - \mu^0)(1 - x_t^1)^k] \geq \mu_0^*$, i.e., when x_t^1 exceeds some minimum value \underline{x}_k (which is 0 if $\mu^0 \geq \mu_0^*$ and positive if $\mu^0 < \mu_0^*$). In this case, we have $x_{t+1}^0 = 1$ (where it remains forever after), while $x_{t+1}^1 = (1 - x_t^1)^k$. If instead $x_t^1 < \underline{x}$, then it is optimal to offer C_1 in period $t + 1$ regardless of whether A is observed.

We thus have two sorts of long run behavior, depending on the relation between μ^0 and μ_0^* .

In the case $\mu^0 < \mu_0^*$, then in state 1, if $x_t^1 < \underline{x}$ for some t , then $x_{t+1}^1 = 0$ from which it follows that $x_{t'}^1 = 0$ for all $t' > t$. And if the economy starts with $x_t^1 \geq \underline{x}$, then from Lemma 7 the sequence x_{t+i}^1 , $i = 1, 2, 3, \dots$ must eventually (in finite time) fall below \underline{x} as it approaches the 0-1 limit cycle; thus eventually, everyone offers C_1 . In state 0, the knowledge that x_t^1 eventually falls below \underline{x} and remains there forever ensures that eventually only C_1 contracts are offered ($x_t^0 = 0$ for all t after some finite date). In other words, with $\mu^0 < \mu_0^*$, the economy eventually settles into the bureaucratic equilibrium described above. (The only exception is if it starts out at the steady state loophole equilibrium – which exists if and only if $\bar{x}_k \geq \underline{x}$ – in which case it remains there, but that is unstable, as we have seen.)

In the case $\mu^0 \geq \mu_0^*$, then $\underline{x} = 0$. In state 0, everyone offers C_0 . In state 1, anyone who does not observe A will offer C_0 , and we have $x_{t+1}^1 = (1 - x_t^1)^k$; as long as $x_t^1 \neq \bar{x}_k$, the economy converges to the limit cycle. This is a loophole equilibrium of sorts, though rather than having a constant fraction of the population offering loophole contracts every period (as in the steady state, which exists but is unstable), *the whole population is eventually synchronized*, alternately offering C_0 then C_1 .

The reason this state-1 cycle is stable is that if x is initially small, it is unlikely that anyone will observe the C_0 firms and see cheating; thus most principals will be convinced to offer C_0 contracts; but the next generation is now *likely* to see cheating, making the ensuing fraction of C_0 shops small, and so on. The fact that a successful incentive scheme

¹⁷In fact, for the analysis we conduct, it will become apparent that individuals need not know this much in order to form their strategies.

by definition masks undesirable behavior makes it all the more likely that the possibility of that behavior will be “forgotten,” not just by individual lineages, but by the population as a whole, contributing to its widespread re-emergence in the future.

To summarize, we have the following

Proposition 8 *If $k > 1$, then under the “Ford” (concurrent-sampling) assumption:*

- (1) *The loophole steady state, if it exists, is unstable in state 1*
- (2) *If $\mu^0 < \mu_0^*$ and the economy does not start out at the loophole steady state, it reaches the bureaucratic steady state in finite time in either state*
- (3) *If $\mu^0 \geq \mu_0^*$ and the economy does not start out at the loophole steady state, then it converges to a period-2 loophole cycle, in which the entire population offers C_0 in State 0 and alternates between offering C_0 and offering C_1 in State 1.*

Santayana Suppose that instead of taking their observations simply from concurrent experience, the apprentices take a more historical view, drawing one observation from each of the last $k - 1$ periods as well as their own. It is easy to see that the 2-period cycle in state 1 cannot persist as an equilibrium as long as $k > 1$. For if everyone offers C_0 in one period and C_1 in the next, every apprentice is sure to get one observation from a C_0 shop and therefore learn of the feasibility of cheating, so that it can never be a best response to offer a C_0 contract.

In a steady state on the other hand, the population from which one samples is by definition the same no matter what period it is in, so that historical sampling and concurrent sampling will always yield the same information. Thus the steady state loophole equilibrium exists under the same conditions as were established in Proposition 5 (indeed that result is proved without reference to the temporal sampling structure). We now establish a stronger result, namely that the loophole steady state is stable starting from any point in the interior of the unit square. Consider first the case $\mu^0 \geq \mu_0^*$.

Formally, the probability that an apprentice observes no A in state 1 is $\prod_{i=1}^k (1 - x_{t-k+i}^1)$; his posterior after such an observation is $\frac{\mu^0}{\mu^0 + (1 - \mu^0) \prod_{i=1}^k (1 - x_{t-k+i}^1)} \geq \mu^0 \geq \mu_0^*$. So $\prod_{i=1}^k (1 - x_{t-k+i}^1)$ of the population offer C_0 at $t + 1$, while the rest offer C_1 . In state 0, of course, no one observes A and so everyone offers C_0 .

Note first that in state 1, there is a cycle of period $k + 1$: the economy spends k periods with no C_0 contracts offered ($x = 0$); in the ensuing period, everyone will have observed no A , conclude they are likely in state 0 and offer C_0 , so that $x = 1$. This of course reveals the state is 1 to everyone in the next k periods, so that $x = 0$ then, and the cycle continues.

This cycle serves to illustrate once again the way information is revealed in this economy. After a sufficient time in which the correct contracts are being offered, the weight of evidence, which precisely because of the correct design of incentive schemes conceals information about feasible actions, serves to convince the population that the expense of protecting against those actions is not worthwhile, and thus loopholes are offered.

The basin of attraction for this cycle turns out to be very small, however.¹⁸ Rather, the long run behavior of the Santayana economy is better described by the steady state

¹⁸Though larger than the orbit itself, the basin is a strict subset of the boundary of the unit hypercube.

equilibrium, with individual but not aggregate cycling, that was discussed in the previous section.

The dynamics are given by $x_{t+1} = S(x_t, x_{t-1}, \dots, x_{t-k+1}) \equiv \prod_{i=1}^k (1 - x_{t-k+i}^1)$. Clearly, \bar{x}_k is (the unique) steady state of $S(\cdot)$. It is straightforward to see (at least for $k = 2$) that \bar{x}_2 is locally stable.¹⁹ This is already a departure from the Ford case, where \bar{x}_2 is unstable. In fact we can show much more: for any $k > 1$, the steady state is an attractor for the entire interior of the unit square.

Proposition 9 *When $\mu^0 \geq \mu_0^*$ and principals are all Santayanas, any sequence of population frequencies x_t of contract C_0 with initial conditions $(x_0, x_1, \dots, x_{k-1})$ in $(0, 1)^k$ converges to the steady state \bar{x}_k .*

The proof is somewhat lengthy and is deferred to the Appendix.²⁰

Why are cycles unstable under historical sampling? In contrast to the Ford case in which sampling is concurrent, a Santayana's learning is not particularly sensitive to the time in which he is born. A Ford who is born in state 1 when many C_0 contracts are being offered will be very likely to learn the state is 1, since there will be many instances of cheating agents. In response, many C_1 contracts will be offered next period. That in turn makes observing A very unlikely for the next generation, leading to a large number of C_0 offers, and so on. Thus concurrent sampling will tend toward aggregate volatility.

Such volatility does not survive when everyone is a Santayana. Compare one who is born when few C_0 contracts to one who is born subsequently when many are offered. They both look back at a history that is broadly similar: the latter day apprentice still samples from the (previous) low- x period, as does his predecessor, for whom it is contemporaneous. Before that, history looks nearly the same to both (it is one period shorter for the latter day apprentice). Thus despite the difference in x at the times of their births, they will be likely to behave similarly, offering C_0 contracts with similar probabilities. Thus differences in x across periods do not persist. Historical sampling has a dampening effect on aggregate fluctuations in organizational form.

Lastly, in the case $\mu^0 < \mu_0^*$, the bureaucratic equilibrium exists as before. But if a loophole steady state exists with $\bar{x}_k > \underline{x}_k$, the bureaucratic equilibrium will have a smaller

¹⁹Denoting $x_t - \bar{x}_k$ by z_t , the linearization of $S(\cdot)$ about \bar{x}_k is

$$z_{t+1} = -(1 - \bar{x}_k)^{k-1} \sum_{i=0}^{k-1} z_{t-i}.$$

For $k = 2$, rewrite this as $\mathbf{z}_{t+1} = \mathbf{A}\mathbf{z}_t$, where $\mathbf{z}_{t+1} = (z_{t+1}, z_t)'$ and $\mathbf{A} = \begin{pmatrix} -(1 - \bar{x}_2) & -(1 - \bar{x}_2) \\ 1 & 0 \end{pmatrix}$. The eigenvalues of \mathbf{A} are $\frac{1}{2} \left[(1 - \bar{x}_2) \pm i\sqrt{4(1 - \bar{x}_2) - (1 - \bar{x}_2)^2} \right]$, which have norm $1 - \bar{x}_2 < 1$, so the steady state \bar{x}_2 is locally stable.

²⁰Banerjee and Fudenberg (2005) study a social learning model in which agents sample from a *random* set of past periods; they establish global stability of a steady state by finding a Lyapounov function, a technique that does not appear to be tractable for Santayana sampling. Their steady states are very different from ours, moreover: eventually all individuals behave the same way (whether correctly or incorrectly), so the heterogeneity and lineage-level cycling that characterize our loophole steady states are absent. As discussed below, a random-sampling strategy is unlikely to be individually optimal in our setting.

basin of attraction than in the Ford case. For since the steady state is stable starting from $(0, 1)^k$ under historical sampling, there is an open set of initial conditions from which the trajectories will never fall below \underline{x}_k and will instead converge to \bar{x}_k .

Welfare Comparing the expected payoff of Ford and Santayana principals over the long period, we see that when $\mu^0 \geq \mu_0^*$, the Santayanas are better off if $k \geq 2$: since they converge to \bar{x}_k in state 1, they offer the loophole contract only $\bar{x}_k < 1/2$ of the time, whereas the Fords in their 2-cycle do so half the time (both offer C_0 in state 0). In this sense, the volatility that results from ignoring history is welfare reducing.²¹ Moreover, not only does historical sampling lead to higher welfare of loophole equilibria than concurrent sampling, it also is less likely to result in the (lower welfare) bureaucratic equilibrium when $\mu^0 < \mu_0^*$.

4.0.8 Extensions

Allowing for Agents to take a while to discover how to cheat The assumption that agents always discover the way to cheat may seem somewhat stylized, so it is worth asking what happens if they do so only sometimes.²² We offer a brief description of what happens here.

Suppose in state 1 the agent discovers A only with probability p ; discovery is independent across agents and periods. To avoid the trivial case in which the principal always offers C_0 , tolerating occasional cheating, assume $\mu_0^* > 1 - p$. With x as the fraction of C_0 contracts offered in state 1, a sampled firm will now reveal A with probability $1 - px$ instead of $1 - x$. The analysis is similar to before except that the Bayesian update for $\Pr(0|\text{Observe no } A)$ is just $\frac{\mu^0}{\mu^0 + (1 - \mu^0)(1 - px)^k}$ and the new critical prior μ_p^* is just $\frac{\mu^* - 1 + p}{p} < \mu^*$. This changes little in terms of the possibility of loophole equilibria, though principals now need to be less optimistic in order to offer loophole contracts – they won't be victimized as often. A loophole steady state still exists.

However it does change the dynamics somewhat. First of all, a principal may offer a loophole contract and be lucky enough to avoid having his agent cheat. An ensuing principal is therefore less likely to observe information that reveals the state is 1. In other words, a contract appears to work for a while, but then begins to break down. So the rate at which contracts change is, not surprisingly, lower. Since $(1 - px)^k$ is decreasing in p , the (unique) steady state value of x is larger, the smaller is p .

²¹The Santayana cycle yields even higher welfare: in it, the incorrect contracts are offered only $1/(k + 1)$ of the time, which is less than x_k .

To see that $\bar{x}_k > \frac{1}{1+k}$, note that for the map $F(\cdot)$ defined in Lemma 7, $F^2(x) = [1 - (1 - x)^k]^k$ has derivative $k^2[1 - (1 - x)^k]^{k-1}(1 - x)^{k-1}$; at \bar{x}_k , this is just $[k(1 - \bar{x}_k)^{k-1}]^2$. As mentioned in the proof of that Lemma, F^2 has three fixed points, 0, \bar{x}_k , and 1; its derivative is 0 at 0 and 1 and therefore $[k(1 - \bar{x}_k)^{k-1}]^2 > 1$ or $k(1 - \bar{x}_k)^{k-1} > 1$. Since $(1 - \bar{x}_k)^{k-1} = \bar{x}_k/(1 - \bar{x}_k)$, it follows that $\bar{x}_k > \frac{1}{1+k}$.

²²We conjecture that similar results obtain in a dynamic extension of the environment leading to Proposition 6, wherein there is residual uncertainty about principal types or only parts of contracts are observed.

At the aggregate level, things depend somewhat delicately on p . If p is close to 1, then under the Ford assumption, the limit cycle still exists, but will have a smaller amplitude. As p decreases, eventually the limit cycle disappears, and the steady state becomes globally stable. The reason is similar to the one giving stability of the steady state under Santayana that we had before: when agents are less likely to cheat, changes in the number of C_0 contracts has less impact on what apprentices see, so their behavior is less responsive to those changes. Under the Santayana assumption, the steady state is stable as before.

So aggregate cycling appears to be most likely when agents are significantly better informed than principals (p close to 1) and when principals ignore history. Otherwise, the aggregate behavior is better described by our steady state, in which lineages but not the population as a whole change behavior over time.

Endogenous Sampling Structures So far we have assumed that k and the temporal structure of samples are exogenous. But of course, principals might *choose* both the size of the sample and whether or not to ignore history.

Sample Size

Consider the two-state case with Ford sampling. Let k be an individual's chosen sample size, \hat{k} the population average. Suppose that x of the population choose C_0 in the bad state. Let the prior that the world is in the good state be $\mu > \mu^*$.

The individual gets

$$\mu\pi_0 + (1 - \mu)(1 - (1 - x)^k)\pi_1$$

if he samples k times. The reason is that with probability $1 - \mu$, it's the bad state, and he will offer a C_1 and earn π_1 only if he gets at least one observation of A , and that happens with probability $1 - (1 - x)^k$; with probability μ it's the good state, and the individual will then surely offer C_0 and earn π_0 .

Treating k as a continuous variable, the marginal value of information is

$$-(1 - \mu)\pi_1(1 - x)^k \ln(1 - x) > 0.$$

and equating this to $c'(k)$, where $c(\cdot)$ is some increasing cost function, characterizes the optimal choice. The marginal value of k is decreasing in k given x . It is nonmonotonic in x , however (it is single-peaked).

Clearly since the value of sampling depends on x , a general characterization is not immediate if we want to take full account of dynamics, where x is changing over time (the case $\mu < \mu^*$ may admit similar complications). This must await future research.

For now, let's ask whether multiple steady state equilibria can happen. Restrict attention further to symmetric equilibria. Then $x(k)$ is really a function of \hat{k} , the choice made by everyone else, and equilibrium requires $k = \hat{k}$. Then we are asking whether the equation

$$-(1 - \mu)\pi_1(1 - x(k))^k \ln(1 - x(k)) = c'(k)$$

can have multiple solutions.

In principle it might, because when \hat{k} is small, $x(\hat{k})$ is large, so people are likely to see

at least one instance of cheating even if they do not have a large sample, so they might as well pick a small sample. If \hat{k} is large, the chance of seeing cheating without a large sample is small, so this gravitates toward choosing a larger sample.

But since the marginal value is not monotonic, any such “strategic complementarity” is limited in scope. Indeed, since $x(k) = (1 - x(k))^k$, the steady-state equilibrium equation reduces to

$$-(1 - \mu)\pi_1 x(k) \ln(1 - x(k)) = c'(k).$$

Now $x(k)$ is decreasing in k , so the left hand side is decreasing (since $-x \ln(1 - x)$ is increasing); thus as long as $c(\cdot)$ is weakly convex and positive, there is a unique steady state with a positive fraction of the population offering loophole contracts in the bad state.

Ford vs. Santayana

Things are more interesting when one considers endogenizing the choice of temporal structure. Here we provide only a sketch of the argument.

Suppose as suggested above that the economy makes a transition state 1 to state 0 with probability ϵ_1 and from 0 to 1 with probability ϵ_0 , that principals are aware of this, and that we take $\mu^0 = \epsilon_0/(\epsilon_0 + \epsilon_1)$, the stationary distribution of this Markov chain. The sample size k is fixed. Principals can only decide whether to sample entirely from the last period or once from each of the previous k periods.

Optimal (pure) strategies in terms of contract offers here would resemble those we studied before, namely to offer C_0 if and only if the posterior exceeded some threshold (which now would differ from μ_0^*).

Now the informativeness of a sample depends both on the contract composition of the period(s) from which it was drawn and on how recent it is. For instance if it is known that in the bad state nearly everyone is offering C_1 , then not seeing A is not very informative in the sense that it does little to convince someone that she is in the good state, but if it is known that few are offering C_1 , then the same observation is very convincing of being in the good state. And if it is known that the world is changing, then more recent observations are more informative than old ones.

If everyone were a Ford, then within an era in which the world is in the bad state, the economy would be cycling and it would be strictly better to take an historical sample, at least provided the transition probability is not too large. On the other hand, if everyone were a Santayana, then within a state, the economy is converging to a steady state. Near the steady state, the compositional informativeness of every period is (nearly) identical, and therefore it is strictly better to be a Ford, which then leads to cycling. So it cannot be an equilibrium for either everyone to be a Ford or everyone to be a Santayana, and instead there must be a mix (which may vary over time). Moreover, there must be at least some aggregate cycling,²³ else everyone would wish to be a Ford. Thus with endogenous sampling, we should expect at least some aggregate cycling during bad state eras, typically of smaller amplitude than in the pure Ford case, but of a frequency much higher than the average fundamental transition rate.

²³More precisely, there must be some aggregate nonstationary behavior: it is possible that the dynamics are complex enough to admit behavior other than steady-state convergence or periodic cycles. Further research is warranted on this question.

5 Discussion

This paper explores the possibility that incomplete knowledge of the contracting environment that is mediated through social learning may serve as an engine of contractual, organizational and institutional change and development. The key observation is that a “tight” contract hides information about the environment, inhibiting social learning and opening the door for future “loopholey” contracts. Very large sample sizes apart, the learning process cannot converge to everyone getting things right. On the contrary, getting things right sows the seeds of its own destruction, and the model generates the possibility of “cycling” and organizational heterogeneity.

A number of further analytical extensions are possible. We might want to consider more general contracting environments – risk aversion, multiple agents, learning across agents and “contagion” effects, the possibility that agents’ outside options bind, variation in the benefits of cheating. The last two raise the interesting possibilities for studying how rates of cheating, learning, and organizational cycling and/or heterogeneity will tend to ebb and flow with industry life cycles or business cycles.

We also want to consider larger state spaces. Some empirical forms of “organization” (e.g. regulatory settings – regulate or not, executive compensation–stock options or not) do seem to conform to our low period cycles. In other cases, (human resource management strategies) what is more striking is the rich set of schemes that have been tried over time, and how those tend to evolve without obvious repetition (piece rates being a possible exception). Large state spaces allow for the latter possibility (period length longer than the age of the universe) and more important for understanding what factors determine whether cycles or short or long.

Although we have focused on the design of incentive schemes in a simple multi-task environment, we believe this type of analysis can be extended to include a broader variety of organizational design issues, such as the choice of ownership structure. For instance, Mattel began its Asian operations in the 1980s by outsourcing production of inputs. In response to significant quality control problems, it turned to an integration strategy which seemed to solve the problem. Yet by the mid 2000s they had returned to outsourcing, which has been attributed cause of the widely publicized and corporate embarrassment of lead toys (Lee, Tseng, and Hoyt, 2008). And conceptually, if not quite formally, the model pertains to the design of relational contracts. Learning all aspects of an implicit agreement between a manager and her employees is likely at least as difficult as learning formal contracts or workplace rules, and no less important for firm performance. Or regulation: witness the adoption tight financial regulation the 1930s following the 1929 crash, the deregulation of the 1990s, the call for new regulation following the 2008 financial crisis.

Our discussion of the effects of different types of sampling begs the question of what institutions, public or private, might serve to foster large samples or historical rather than concurrent sampling. One institution is the common law court system in its resolution of contractual disputes and torts. Broadly this has three features. It is invoked only when there are disputes, which suggests that some sort of “mistake” has occurred between parties. Secondly, it makes its decisions as well as the reasoning leading to them a matter of public record, which in terms of our model serves effectively to increase k as well as what

aspects of contracts are observable. Third, its reliance on precedent may serve to mimic the backward sampling of our Santayanas (a second, related, effect is to lengthen memory). Of course this institution is costly to set up and maintain.²⁴

Moreover, the scope of the court system is limited. Rarely does it get involved internal governance of firms, and even less so in the design of compensation schemes and other organizational design elements. What other institution could serve the role of broadcasting to a wide audience who will be responsible for these crucial decisions, credibly instructing them in the effects of improper organizational design, while drawing its cases not just from current management fashion, but also from the history of management practice? This would seem to be the natural role of the business schools.

To conclude, as a number of economists have remarked, one takeaway from the burgeoning literature on incentives, contracts and organizations, including its formal counterpart, mechanism design, is that getting incentive systems right is “hard,” as much so, it seems, as getting a technology such as an air frame or a microchip right. This represents a considerable (and not yet universally accepted) shift in perspective for economics, which traditionally views technology as exogenous (read infinitely hard), while incentive systems are supposed to adapt quickly and easily from one environment to the next. Yet it doubtless contains a large grain of truth: a casual look at economic history would suggest that the rate of progress in the design of incentive systems has been rather slow compared to that of technology, at least over the past few hundred years.

Now while most economists from the traditional point of view will grant that many situations that economic agents find themselves in are difficult to understand and that mistakes will therefore be made, they would also argue that eventually agents will learn to get it right, so that our models of instant understanding are good enough approximations. What this paper suggests is that a possible caveat to that view, which helps lend credence to the newer perspective, is that compared to technological design, incentive design is not so much harder to *do* as it is harder to *learn*.

6 Appendix

Proof of Lemma 7

Let $F(x_t) = (1 - x_t)^k$. Then $x_{t+1} = F(x_t)$ and $x_{t+2} = F^2(x_t) \equiv F(F(x_t)) = [1 - (1 - x_t)^k]^k$. These two functions are plotted in the figure. It is straightforward to verify that $F(0) = 1 = 1 - F(1)$, F is decreasing, and has a unique fixed point \bar{x}_k . Meanwhile F^2 is increasing with fixed points $0, \bar{x}_k$, and 1 . Simple calculations show there is a unique \hat{x}_k such that F^2 is strictly convex on $(0, \hat{x}_k)$ and strictly concave on $(\hat{x}_k, 1)$. Thus F^2 has no other fixed

²⁴In particular, it is subject to its own agency problems and performance depends delicately on design. In an interesting pair of papers, Gennaioli and Shleifer (2007a,2007b) show that if judges are committed only to making fine distinctions between cases, that convergence of the law to (in general, second-best) efficiency can result, even if judges are self-interested and have policy biases. This parallels our convergence to the loophole steady state under aggregate Santayana dynamics. However, if (self-interested) judges with policy preferences have the power to *overturn* precedent, there may be a tendency toward cycling. The parallel with our Ford case is perhaps less clear because overturning in effect puts the judge in a role similar to our agents.

points in the unit interval, and $F^2(x) < x$ in $(0, \bar{x}_k)$ and $F^2(x) > x$ in $(\bar{x}_k, 1)$, as depicted in the figure. (In fact, the inflection point $\hat{x}_k = 1 - (\frac{1}{k+1})^{1/k}$ exceeds \bar{x}_k , which follows from the argument in footnote 21.) Observe that there is a cycle of period 2: if $x_t = 1$, then $x_{t+1} = 0$, $x_{t+2} = 1$, etc.




Figure 1: One-step (black) and two-step (blue) Ford dynamics

We show that all trajectories starting away from the point \bar{x}_k converge to the 2-cycle $(0, 1, 0, 1, \dots)$. Let $x_0 \in (0, \bar{x}_k)$. Then $x_1 > \bar{x}_k$ (since $F(x) > \bar{x}_k$ on $(0, \bar{x}_k)$). But the “even iterates” x_2, x_4, \dots given by $F^2(x_0), F^4(x_0), \dots$ converge monotonically to 0. Meanwhile the odd iterates $x_1, x_3, x_5, \dots = F(x_0), F^2(F(x_0)), F^4(F(x_0)), \dots$ converge monotonically to 1. Thus, the sequence $\{x_t\}$ converges to the 2-cycle. A similar argument holds for $x_0 \in (\bar{x}_k, 1)$; thus any solution originating in $(0, 1) \setminus \{\bar{x}_k\}$ converges to the 2-cycle.

Santayana Dynamics and Proof of Proposition 9

We study the equation

$$x_{t+1} = \prod_{i=1}^k (1 - x_{t+1-i}).$$

To save some notation, it will be convenient to study the dynamics of $y = 1 - x$; our equation becomes

$$y_{t+1} = 1 - \prod_{i=1}^k y_{t+1-i}. \quad (1)$$

Note that $y = 1 - y^k$ has a unique solution \bar{y} in $[0, 1]$ that is the (unique) steady state of (1), and that \bar{y} is increasing in k ; we suppress this dependence to save notation. As pointed out in the text, there are periodic solutions to this equation with $y_t = 0, y_{t+i} = 1, i = 1, \dots, k$. However for solutions originating within the interior of the unit cube, all trajectories converge to \bar{y} , as we will now show.

Denote by \bar{C} the unit cube $[0, 1]^k$ and C its interior. Write $\overline{Co}(X)$ for the convex hull of an arbitrary $X \subset \bar{C}$ and $Co(X)$ for its interior. Let initial conditions $\mathbf{y}_0 = (y_{-k+1}, y_{-k+2}, \dots, y_0)$ be given; a solution $\mathbf{y}(\mathbf{y}_0) = \{y_{-k+1}, \dots, y_0, y_1, y_2, \dots\}$ with $\mathbf{y}_0 \in Y \subseteq \bar{C}$ is said to *originate* in Y . Call a finite sequence of consecutive y values $\{y_t, y_{t+1}, \dots, y_{t+n}\}$ that are part of a solution an *n-string*. An *n-string* $\{y_t, y_{t+1}, \dots, y_{t+n}\}$ in which all elements weakly exceed (are less than or equal to) \bar{y} , with at least one inequality strict, and for which $y_{t-1}, y_{t+n+1} < \bar{y}$ ($y_{t-1}, y_{t+n+1} > \bar{y}$) is called a positive (negative) *semicycle*. We provide a series of lemmas that help to characterize any solution originating in C .

Lemma 10 $\mathbf{y}_0 \in C$ implies $y_t \in C$ for all $t \geq 1$.

Lemma 11 No semicycle exceeds k in length.

Proof. Note that $y > \bar{y}$ if and only if $1 - y^k < \bar{y}$. Thus for a k -string $\{y_{t+1}, \dots, y_{t+k}\}$ that is part of a positive semicycle with $y_{t+i} > \bar{y}$ for at least one $i \in \{1, \dots, k\}$, we have $y_{t+k+1} = 1 - \prod_{i=1}^k y_{t+i} < 1 - \bar{y}^k = \bar{y}$; the negative semicycle case is similar. ■

Lemma 12 Let $\mathbf{y}_t^{k+1} = \{y_t, y_{t+1}, \dots, y_{t+k}\}$ be a $k+1$ -string. Then $\bar{y} \in Co(\mathbf{y}_t^{k+1}) \subset (0, 1)$.

Proof. If $\bar{y} \in \mathbf{y}_t^{k+1}$, there is nothing to prove. In the other case, from the previous lemma, \mathbf{y}_t cannot contain elements that all exceed \bar{y} or are all less than \bar{y} , else there would be a semicycle of length greater than k . So there are $y_{t+i}, y_{t+j} \in \mathbf{y}_t^{k+1}$ with $y_{t+i} < \bar{y} < y_{t+j}$, $i, j \in \{0, \dots, k\}$. Thus $\bar{y} \in (y_{t+i}, y_{t+j}) \subseteq Co(\mathbf{y}_t) \subset (0, 1)$, where the second inclusion follows from lemma 10. ■

Lemma 13 Let $\{y_t, y_{t+1}, \dots, y_{t+k}\}$ be a $k+1$ -string, and let M and m be its maximal and minimal elements. Then

- (a) $y_{t+k+1} \in [m, M]$
- (b) $\min\{y_t, y_{t+k}\} \leq y_{t+k+1} \leq \max\{y_t, y_{t+k}\}$.

Proof. We first show that $y_{t+k+1} \leq M$. Assume instead that $y_{t+k+1} > y_{t+i}$, $i = 0, 1, \dots, k$. Since $y_{t+k} = 1 - y_t y_{t+1} \dots y_{t+k-1}$ and $y_{t+k+1} = 1 - y_{t+k} y_{t+1} \dots y_{t+k-1}$, the case $i = k$ is equivalent to $y_{t+k} < y_t$ or

$$1 - \pi_t < y_t \quad (2)$$

where $\pi_t = y_t y_{t+1} \dots y_{t+k-1}$.

For the case $i = 0$, note that $y_{t+k+1} > y_t$ can be written $1 - \frac{(1-\pi_t)\pi_t}{y_t} > y_t$, or

$$(1 - \pi_t)\pi_t < y_t(1 - y_t) \quad (3)$$

We now show that 2-3 are inconsistent. First, we cannot have $\pi_t = 1/2$, since then $y_t(1 - y_t) \leq (1 - \pi_t)\pi_t$, violating (3). If $\pi_t < 1/2$, then (3) implies $1 - \pi_t > 1 - y_t$ and $1 - \pi_t > y_t$, and the latter violates (2). And if $\pi_t > 1/2$, then (3) implies that $\pi_t > y_t$, which contradicts lemma 10.

Similarly, we can show that $y_{t+k+1} \geq m$. If instead $y_{t+k+1} < y_{t+i}$, $i = 0, 1, \dots, k$, then we must have from $y_{t+k} > y_{t+k+1}$

$$y_t < 1 - \pi_t \quad (4)$$

and from $y_t > y_{t+k+1}$

$$(1 - y_t)y_t < \pi_t(1 - \pi_t) \quad (5)$$

By an argument similar to the one used to show $y_{t+k+1} \leq M$, (4) and (5) are inconsistent and we conclude $y_{t+k+1} \geq m$. This establishes (a). ■

Since only the inequalities involving y_t and y_{t+k} have been used, the argument also establishes (b).

Lemma 14 *If $y_{t+k+1} = y_t$ then $y_{t+k} = y_t$; if $y_{t+k+1} = y_{t+k}$ then $y_{t+k+1} = y_t$*

Proof. $y_{t+k+1} = y_t$ is equivalent to $y_t(1 - y_t) = \pi_t(1 - \pi_t)$ using the previous notation. From lemma 10 we have $\pi_t < y_t$. Thus, we must have $\pi_t = 1 - y_t$. Since $y_{t+k} = 1 - \pi_t$, we establish the first claim. The second hypothesis is equivalent to $1 - \frac{\pi_t(1-\pi_t)}{y_t} = y_{t+k} = 1 - \pi_t$, whence $\frac{1-\pi_t}{y_t} = 1$, therefore $y_t = y_{t+k} = y_{t+k+1}$, as desired. ■

Assume now that our solution \mathbf{y} originates in C . Define $H_t = \overline{Co}(\mathbf{y}_t^{k+1})$ for every $k + 1$ -string \mathbf{y}_t^{k+1} that is part of \mathbf{y} . By lemma 13, $H_{t+1} \subseteq H_t$. Denote $H = \bigcap_{t=1}^{\infty} H_t$, which is closed, convex, and, by lemma 12, contains \bar{y} . Assume that $H = [l, L]$, where $l \leq \bar{y} \leq L$. Since the H_t are nonincreasing in the set inclusion order, we must have $l > 0$ and $L < 1$. We will now show that $L = \bar{y}$; combined with a similar argument that shows $l = \bar{y}$, we will thereby establish

Proposition 9. *When $\mu^0 \geq \mu_0^*$ and principals are all Santayanas, any path of population frequencies of C_1 that originates in C converges to the steady state \bar{y} .*

The proof reasons on the accumulation points of the sequence of $k + 1$ strings of \mathbf{y} , considered as points in $(0, 1)^{k+1}$. Denote a typical element of the sequence by $\sigma^t = (y_t, y_{t+1}, \dots, y_{t+k})$, where by definition $y_{t+k} = \hat{S}(y_t, y_{t+1}, \dots, y_{t+k-1}) = 1 - y_t y_{t+1} \dots y_{t+k-1}$. All accumulation points of the sequence $\{\sigma^t\}$ (which exist since $\{\sigma^t\}$ is contained in the compact set $[0, 1]^{k+1}$) are contained in H^{k+1} . Since every σ^t has a component that is a maximal element M_t of the corresponding $k + 1$ -string, and the maxima $M_t \rightarrow L$, the accumulation point $\bar{\sigma}$ has at least one component equal to L . Since $H = [l, L]$, we must have $\bar{\sigma}_i \leq L$, $i = 0, 1, \dots, k$.

We now show

Lemma 15 For any accumulation point $\bar{\sigma} = (\bar{\sigma}_0, \bar{\sigma}_1, \dots, \bar{\sigma}_k)$ of $\{\sigma^t\}$, we have $\bar{\sigma}_k = L$.

Proof. Suppose not. Let $\{\sigma^{t_n}\} = \{(y_{t_n}, y_{t_n+1}, \dots, y_{t_n+k})\}$ be a subsequence converging to $\bar{\sigma}$. Begin by assuming $\bar{\sigma}$ has exactly one component $\bar{\sigma}_j = L$, $j < k$. By continuity of $\hat{S}(\cdot)$, $y_{t_n+k+1} = \hat{S}(y_{t_n+1}, \dots, y_{t_n+k})$ converges to a limit $\bar{\sigma}_{k+1} = \hat{S}(\bar{\sigma}_1, \dots, \bar{\sigma}_k)$ and thus $\bar{\sigma}' = (\bar{\sigma}_1, \dots, \bar{\sigma}_k, \bar{\sigma}_{k+1})$ is also an accumulation point of $\{\sigma^t\}$, the limit of the shifted subsequence $\{\sigma^{t_n+1}\} = \{(y_{t_n+1}, \dots, y_{t_n+k+1})\}$. Since $\min\{y_{t_n}, y_{t_n+k}\} \leq y_{t_n+k+1} \leq \max\{y_{t_n}, y_{t_n+k}\}$ by lemma 13(b), we have $\min\{\bar{\sigma}_0, \bar{\sigma}_k\} \leq \bar{\sigma}_{k+1} \leq \max\{\bar{\sigma}_0, \bar{\sigma}_k\}$. Since $\bar{\sigma}_k < L$, we cannot have $\bar{\sigma}_{k+1} = L$ unless $\bar{\sigma}_0 = L$; but then by lemma 14, we have $\bar{\sigma}_k = L$, a contradiction; thus $\bar{\sigma}_{k+1} < L$. Since $\bar{\sigma}$ has at most one component equal to L , so does $\bar{\sigma}'$.

Continuing in this way through the shifted subsequences $\{\sigma^{t_n+i}\}$, we arrive at $\{\sigma^{t_n+j}\}$ whose limit is $(\bar{\sigma}_j, \bar{\sigma}_{j+1}, \dots, \bar{\sigma}_{j+k})$, with $\bar{\sigma}_{j+i} < L$, $i = 1, \dots, k$. Then $\bar{\sigma}_{j+k+1} < L$, else by lemma 14 $\bar{\sigma}_{j+k} = L$, a contradiction. But then $(\bar{\sigma}_{j+1}, \bar{\sigma}_{j+1}, \dots, \bar{\sigma}_{j+k+1})$ is an accumulation point of $\{\sigma^t\}$ with all components less than L , which contradicts that $H = [l, L]$. We conclude that we cannot have an accumulation point with $\bar{\sigma}_k < L$ that has only one component equal to L .

Suppose now that if r components of $\bar{\sigma}$ are equal to L with $\bar{\sigma}_k < L$, then there is an accumulation point $\bar{\sigma}'$ with $r - 1$ components equal to L with $\bar{\sigma}'_k < L$, where $2 \leq r \leq k$. Since this implies by induction that there is an accumulation point $\bar{\sigma}''$ with one component equal to L and $\bar{\sigma}''_k < L$, we have a contradiction, and we conclude that $\bar{\sigma}_k = L$.

For the inductive step, starting from $\bar{\sigma}^0 = (\bar{\sigma}_0, \dots, \bar{\sigma}_k)$, construct accumulation points $\bar{\sigma}^j = (\bar{\sigma}_j, \bar{\sigma}_{j+1}, \dots, \bar{\sigma}_{j+k})$, $j = 1, \dots, k$ as above. Let \underline{j} be the lowest j such that $\bar{\sigma}_j = L$. For $j < \underline{j}$, $\bar{\sigma}^j$ has r components equal to L and the k^{th} component less than L . And $\bar{\sigma}_{\underline{j}+k+1} < L$, else by lemma 13(b) $\bar{\sigma}_{\underline{j}+k} = L$, so that $\bar{\sigma}^{\underline{j}}$ has $r + 1$ components equal to L , a contradiction. But then $\bar{\sigma}^{\underline{j}+1} = (\bar{\sigma}_{\underline{j}}, \bar{\sigma}_{\underline{j}+1}, \dots, \bar{\sigma}_{\underline{j}+k})$ is an accumulation point with $r - 1$ components equal to L and its k^{th} component less than L , as claimed. This completes the proof of the lemma. ■

Now we show that in fact there is an accumulation point $\bar{\sigma}^* = (L, L, \dots, L)$. Starting with $\bar{\sigma}^0$, which by lemma 15 is of the form $(\bar{\sigma}_0, \bar{\sigma}_1, \dots, L)$, construct the accumulation point $\bar{\sigma}^1 = (\bar{\sigma}_1, \dots, L, \bar{\sigma}_{k+1})$; also by lemma 15 we must have $\bar{\sigma}_{k+1} = L$. After repeating this for k steps, we obtain an accumulation point $\bar{\sigma}^* = (L, L, \dots, L)$ as desired. Since $\bar{\sigma}_k^* = \hat{S}(L, L, \dots, L) = 1 - L^k = L$, we conclude that $L = \bar{y}$.

A similar argument can be applied to l to conclude that $l = \bar{y}$. Then $H = \{\bar{y}\}$, and the proof that $y_t \rightarrow \bar{y}$ is complete.

References

- [1] Baker, G. P. (1992), "Incentive Contracts and Performance Measurement," *Journal of Political Economy* 100(3): 598-614.
- [2] Baliga, S., and J. Ely (2009) "Mnemonomics: The Sunk Cost Fallacy as a Memory Kludge," mimeo

- [3] Banerjee, A. (1992) "A Simple Model of Herd Behavior," *Quarterly Journal of Economics* 107(3): 797-817.
- [4] Banerjee, A. and D. Fudenberg (2005), "Word of Mouth Learning," *Games and Economic Behavior*, 46 (1): 1-22.
- [5] Berg, N., and N. Fast (1975), "The Lincoln Electric Company," Harvard Business School Case no. 9-376-028.
- [6] Bernheim, B.D. and M. Whinston (1998), "Incomplete Contracts and Strategic Ambiguity," *American Economic Review* 88(4): 902-932.
- [7] Bertrand, M. and A. Schoar (2003), "Managing with Style: The Effect of Managers on Corporate Policy," *Quarterly Journal of Economics*, 118 (4): 1169-1208.
- [8] Bevan and Hood (2004), "Targets, Inspections, and Transparency," *British Medical Journal*, 328: 598 (11 March 2004).
- [9] ----- (2006), "What's Measured is What Matters: Targets and Gaming in the English Public Health Care System", *Public Administration*, 84(3): 517-538.
- [10] Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992), "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100(5): 992-1026.
- [11] Bolton, P. and C. Harris (1999), "Strategic Experimentation," *Econometrica*.
- [12] Bolton, P. and A. Faure-Grimaud (2005) "Thinking Ahead: The Decision Problem", NBER Working Paper No. W11867
- [13] Chung, K.S. and Fortnow (2007), "Loopholes," mimeo
- [14] Courty, P., and G. Maschke (2004), "An Empirical Investigation of Gaming Responses to Explicit Performance Incentives," *Journal of Labor Economics* 22(1): 23-56.
- [15] ----- (2008), "A General Test for Distortions in Performance Measures," *Review of Economics and Statistics*, 3: 428-441.
- [16] Crémer, J. (1986), "Cooperation in Ongoing Organizations", *Quarterly Journal of Economics*, 101(1): 33-49.
- [17] Chevalier, J. and G. Ellison (1997), "Risk Taking by Mutual Funds as a Response to Incentives," *Journal of Political Economy* 105(6): 1167-1200.
- [18] Ederer, F., R. Holden, and M. Meyer (2010), "Gaming and Strategic Ambiguity in Incentive Provision," mimeo.
- [19] Ellison, G. and D. Fudenberg (1993), "Rules of Thumb for Social Learning," *Journal of Political Economy* 101(4): 612-643.

- [20] Ellison, G. and R. Holden (2009), "A Theory of Rule Development," <http://faculty.chicagobooth.edu/richard.holden/papers/rules.pdf>
- [21] Frank, Douglas H., Jr. and Tomasz Obloj (2009), "Ability, Adverse Learning and Agency Costs: Evidence from Retail Banking," http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1374784
- [22] Fudenberg, D. and D. Levine (1998): *The Theory of Learning in Games*, MIT Press
- [23] Gennaioli, N. and A. Shleifer (2007a), "The Evolution of Common Law," *Journal of Political Economy* 115: 43-68.
- [24] ----- (2007b), "Overruling and the Instability of Law," *Journal of Comparative Economics* 35: 309-328.
- [25] Goodpaster, K.E. and R.J. Post (1981), "H.J. Heinz: The Administration of Policy," Harvard Business School Case no. 382-036.
- [26] Gravelle, H., M. Sutton, M., and A. Ma (2010), "Doctor Behaviour under a Pay for Performance Contract: Treating, Cheating, and Case Finding," *Economic Journal*, 120(542): F129-F156.
- [27] Holmström, B. and P. Milgrom (1991), "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics and Organization* 7: 24-52.
- [28] Jehiel, P. (2005): "Analogy-based Expectation Equilibrium," *Journal of Economic Theory* 123: 81-104.
- [29] ----- (2010), "On Transparency in Organization," mimeo.
- [30] Kerr, S. (1975) "On the Folly of Rewarding A while Hoping for B," *Academy of Management Journal*, 18: 769-783.
- [31] Lee, H. L., M. M. Tseng, and D. Hoyt (2008). "Unsafe For Children: Mattel's Toy Recalls and Supply Chain Management," Stanford University, Graduate School of Business, case study GS63.
- [32] Mullainathan, S. (2002) "A Memory-based Model of Bounded Rationality," *Quarterly Journal of Economics* 117: 735-774.
- [33] Myerson, R.B. (1982) "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics* 10: 67-81.
- [34] Oyer, P. (1998), "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality," *Quarterly Journal of Economics*, 113(1): 149-185.
- [35] Piccione, M. and A. Rubinstein (1997), "On the Interpretation of Decision Problems with Imperfect Recall," *Games and Economic Behavior* 20: 3-24.

- [36] Piketty, T. (1995), "Social Learning and Redistributive Politics," *Quarterly Journal of Economics*
- [37] Propper, Sutton, Whitnall, and Windmeijer (2008), "Incentives and Targets in Hospital Care," Bristol CMPO D.P. 08/20, <http://www.bristol.ac.uk/cmppo/publications/papers/2008/wp205.pdf>
- [38] Scharfstein, D. and J.Stein (1990), "Herd Behavior and Investment," *American Economic Review*
- [39] Schoar, A., (2007), "CEO Careers and Styles," http://www.mit.edu/~aschoar/CEOCareersStyle_v2.pdf
- [40] Spier, K. (1992), "Incomplete Contracts and Signaling," *RAND Journal of Economics* 23 (3): 432-443.
- [41] Squintani, F. (2003), "Moral hazard, renegotiation, and forgetfulness," *Games and Economic Behavior* 44: 98-113
- [42] Stiglitz, J.E. (1975), "Incentives, Risk , and Information: Notes Towards a Theory of Hierarchy," *Bell Journal of Economics*, 6(2): 552-579
- [43] Tirole, J. (2009), "Cognition and Incomplete Contracts," *American Economic Review* 99(1): 265–294.
- [44] Whyte, W.F., M. Dalton, D. Roy, L. Sayles, O. Collins, F. Miller, G. Strauss, F. Fuerstenberg, and A. Bavelas (1955), *Money and Motivation: An Analysis of Incentives in Industry*, Harper & Row.
- [45] Wilson, A. (2002) "Bounded Memory and Biases in Information Processing," mimeo