

# On the Origin and Evolution of Corporate Culture

## PRELIMINARY AND INCOMPLETE

Eric Van den Steen\*

April 7, 2003

### Abstract

Where does corporate culture come from? In trying to answer that question, this paper also suggests a new perspective on the relationship between corporate culture and performance.

The paper does essentially two things. First, it develops a simple model of corporate culture, based on the notion of culture as shared beliefs or assumptions, as suggested by the management literature. Culture, in this model, evolves from the common experiences of a firm's members. To partially validate the model, I show that it captures the following stylized facts. The culture of a firm is heavily influenced by the initial beliefs of the founder or early leader, and can persist even long after that founder or early leader is gone. External succession of a CEO is more likely to lead to a change in corporate culture than internal succession. Otherwise identical firms may develop very different cultures. Suboptimal cultures may persist, even if the members of the organization know that their culture is almost surely suboptimal.

Second, the paper uses this model to study the key link between culture and performance. It shows that, *even without any effect of culture on performance*, there will be a positive correlation between the two. This result questions some received wisdom about corporate culture.

## 1 Introduction

This paper starts from a simple stylized fact: organizations tend to be more homogenous in terms of beliefs, values, and behavioral patterns than society at large. Moreover, these shared beliefs and norms are fairly constant over time, giving the organization a kind of ‘personality.’ In the management literature, such shared beliefs and values are often referred to as ‘corporate culture’ (Burns Stalker 1961, Schwartz and Davis 1981, Peters and Waterman 1982, Donaldson and Lorsch 1983, Schein 1985, Kotter and Heskett 1992).<sup>1</sup>

---

\*MIT-Sloan School of Management (evds@mit.edu). This paper benefitted greatly from two extensive discussions with Bob Gibbons, the discussion by Ben Hermalin, and very useful remarks from George Baker, Oliver Hart, Ed Lazear, Jesper Sørensen, Birger Wernerfelt, the participants in the NBER organizational economics conference and the MIT organizational economics lunch, and the seminar participants at USC, and from the guidance of Bengt Holmström and John Roberts. REMARK: this paper does not yet reflect the suggestions by the USC seminar participants, Kevin Murphy in particular.

<sup>1</sup>Section 2 discusses more extensively the concept ‘corporate culture.’ While there are some important alternatives, this definition of culture as shared beliefs (or mental models) and values that drive behavioral norms is the most prevalent in the management literature (e.g. Burns Stalker 1961, Schwartz and Davis 1981, Peters and Waterman 1982, Donaldson and Lorsch 1983, Schein 1985, Kotter and Heskett 1992).

The phenomenon has important economic implications. Homogeneity in values and beliefs simplifies delegation (Aghion and Tirole 1997, Dessein 2001), communication (Crawford and Sobel 1982), and coordination (Crémer 1993), but may reduce the incentives to collect information (Van den Steen 2002b). The persistence of such beliefs might help explain the persistent differences in performance of firms in the same industry (Mueller 1990, McGahan 1999). Indeed, the business press often explains business successes and failures by referring to the reigning corporate culture.<sup>2</sup>

Homogeneity in beliefs and values is also important for agency theory in general, because the latter tries to deal precisely with heterogeneity in objectives. In particular, the above observations suggest that stable organizations might have an innate ability to minimize agency problems by generating internal homogeneity. In other words, organizations might prevent agency problems rather than solving them. Might this be one of the advantages of a firm over a market?

To understand these issues, their implications, and the potential to manage them, we need to understand the origin of this homogeneity. Two explanations come up quite naturally.

1. Over time, the members of an organization develop similar mental models and beliefs since they share the same experiences. This leads them to hold similar values and to act in similar ways.
2. Organizations select and retain people who have beliefs and values that are similar to their own. Such sorting also leads to similar behavioral patterns.

This paper focuses on the first explanation, while the sorting story is studied in Van den Steen (2002a).

**The model** The paper studies a model in which a number of firms are each faced with a multi-armed bandit problem.<sup>3</sup> For reasons I discuss in more detail in Section 2.4, I focus on the case that the actions and outcomes of a firm are observed by the members of that firm only. Members of the same firm thus have identical experiences, which differ from these of members of other firms. As a consequence, the beliefs of a firm's employees will converge more than in society at large. Having similar beliefs, the employees of a firm will tend to act in similar ways. The firm thus develops a culture in the sense of shared beliefs that drive similar behavior.

After establishing these basic facts, I try to partially validate the model by showing that it passes the 'minimal test' of exhibiting certain stylized facts regarding culture. I show, for example, how the beliefs of the founder have an important influence on a firm's culture, that a culture can persist even in the view of complete turnover, that structurally similar firms may develop very different cultures with different levels of performance, that outsider successors to a CEO are more likely to cause change than insiders, that suboptimal cultures may persist, even if the members of the organization realize that their culture is almost surely suboptimal. I also discuss how turnover and multiple decision-makers would play out in the model.

---

<sup>2</sup>Some famous examples include the aggressive culture of Microsoft, the sales culture of Oracle, the rebel culture of Apple, the cowboy culture of Enron. Incompatible cultures are considered an obstacle to mergers and acquisitions, while culture is also cited as a major hurdle for corporate change.

<sup>3</sup>In a bandit problem, a firm has to choose repeatedly a course of action among  $N$  different possibilities. The expected returns of the actions are unknown, but firms learn them over time by observing realized returns. When choosing an action, a firm trades off immediate returns (from choosing the action which seems to have the highest expected payoff) against future benefits from learning (by choosing other actions).

After this partial validation, I draw one important new insight from the model. I show in particular that there will be a correlation between cultural strength and performance, *even when the homogeneity of beliefs has no direct impact on performance*. This casts some doubt on the culture-based literature that cites such correlations as proofs that culture is valuable and then goes on to explain why it is valuable. While my analysis does not imply that homogeneity is not valuable, it suggests that care must be taken when interpreting such correlations.

Before proceeding to the literature review, it is useful to make two remarks regarding the information and beliefs in this model. The first is that the model assumes, as mentioned earlier and discussed in more detail in Section 2.4, that the information about actions and outcomes is difficult to communicate. People only ‘get it’ through personal experience, personal observation of others, or through a long process of learning from stories and symbols. The reason for focusing on this type of information is the fact that information that is easy to communicate will diffuse quickly to society at large and cannot therefore be the basis of a distinctive culture. Hence, I exclude such factual information from the model. The second remark is that agents in this model are commonly known to have differing priors. The reason for the assumption is that differences in beliefs are at the heart of this theory and assuming common knowledge of differing priors is the most transparent way to model this.<sup>4</sup> Appendix A discusses this in more detail and considers the implications and complications of imposing a common prior.

**The Literature.** The papers that are most closely related are Crémer (1993) and Lazear (1995). Crémer (1993) defines culture, following Schein (1985), as a stock of shared knowledge and argues that it improves the efficiency of information processing. Most of the paper is an informal discussion on communication, but it also presents an interesting team-theoretic model that shows how shared information may improve coordination. While it starts essentially from the same definitions as the current paper, there are three key differences. First, Crémer seems to focus on factual information that is relatively easy to share (at some fixed cost). Second, the model does not really deal with the genesis and evolution of culture, but focuses on its effects. Some of Crémer’s comments suggest that the content of culture is a simple managerial choice. Third, except for the initial cost of sharing the information, more culture is always better in his model. As a consequence, the model has no natural explanation why culture might be dysfunctional or why firms in the same industry with similar history might end up having very different cultures.

Lazear (1995) also defines culture as ‘shared beliefs, values and technology’ and considers a ‘genetic’ evolutionary model of corporate culture, built on the assumption that culture is contagious but that management can influence the reproduction rates of the different ‘genes.’ This work is complementary to the current paper in that it provides an alternative perspective how culture might evolve. It also contains a broad overview of the non-economic literature on the topic. The strength and weakness of the paper lies in its level of abstraction. Without much micro-foundations, the paper assumes that cultural traits are like genes, that they are contagious, and that management can affect the transmission of culture. The most important difference, however, is that the paper, like most, implicitly assumes that homogeneity is valuable. As a consequence the model has no natural explanation for dysfunctional cultures or for the fact that structurally similar firms might end up with very different cultures.

An important alternative model of culture is that presented by Kreps (1990) and further clarified

---

<sup>4</sup>The assumption is formally not required for the theory if we are careful as to ‘who knows what when.’ But the model becomes awkward because I have to make people disagree without openly doing so. See Appendix A for a further discussion.

and interpreted by Hermalin (2001). As Hermalin points out, there are in fact two notions of corporate culture in Kreps' paper. The first notion, and the way the paper is usually interpreted, is that of culture as a mechanism to coordinate in the presence of multiple equilibria. Culture is thus a pure convention, analogous to driving on the left or on the right. Note that such culture is very fragile. While this idea of equilibrium selection is the most common interpretation of the paper, it is actually not how Kreps himself defined culture in that paper. Kreps' own definition, and the second notion of culture according to Hermalin, is that of culture as a reputation for dealing in a specific way with unforeseen contingencies. According to Kreps, such reputation is useful to protect employees against abuses of authority. Both Kreps and Hermalin refer to 'trust' in a repeated game as the way to think about this, but then with the future stage-games being unforeseen. As Hermalin points out, however, this notion is based on a fair amount of 'hand-waving,' since the idea of unforeseen contingencies and especially the idea of having a reputation for how you will deal with things that are by definition not yet known, defies economic modelling, at least for the time being. The most visible difference with the current paper is its definition of corporate culture. While I defer a complete discussion of alternative definitions to Section 2.2, it is useful to make two remarks at this point. First, a complex construct as culture can have more than one valid definition, depending on the context. So there is not necessarily a need to argue which definition is 'the right one.' Moreover, arguments on definitions are always to some extent semantic arguments. Second, the definitions might actually be closer than it seems at first, as discussed later. Probably a more important difference is that the Kreps-Hermalin model is essentially agnostic with regard to the origin of culture, which is exactly the focus of this paper. Overall, that paper presents an interesting perspective on culture that is in many respects complementary to the current work.

Apart from his important discussion of Kreps' model, Hermalin (2001) also summarizes and reinterprets other existing research, and adds to it by linking the topic with insights in other fields of economics, such as IO. In one model, for example, he assumes that culture is a fixed cost investment that lowers the variable cost and then derives industry-level implications. Carrillo and Gromb (1999) model corporate culture as production technologies for which employees can make specific investments and show that a risk that the firm will change technology leads to the coexistence of a strong-culture and a weak-culture equilibrium. Rob and Zemsky (2002) present a theory in which firms differ in the stationary state levels of cooperation among their employees, which they link to the notion of corporate culture. It is also interesting to note that Greif (1994) considers 'cultural beliefs' an important part of a national culture. This paper is also related to the work of Prescott and Visscher (1980). In particular, the current model could be extended such that culture is in part organization capital in the sense as they define it. As mentioned earlier, Van den Steen (2002a) considers the sorting story as the origin of culture. Two other related papers along these lines are Van den Steen (2002d) and Van den Steen (2002b). The first discusses how a manager's strong beliefs, or vision, may lead to sorting in the labor market and thus lead to homogenous beliefs within the firm. The second discusses some further costs and benefits of homogeneity of beliefs.

Weber and Camerer (2001) present experimental results that bear on the phenomenon culture. They let pairs of people ('firms') develop, through trial and error, a homemade language for solving problems, which they interpret as the firm's culture. They then merge groups and show that their performance declines after the merger. A key difference with the current theory is that culture in their sense has an obvious performance advantage.

The contribution of this paper, then, is to present a model of corporate culture as shared beliefs in formal economic terms, to partially validate it by comparing comparative statics to stylized facts,

and to derive new insights from it regarding the relationship between culture and firm performance.

An important implicit insight of the paper is that firms and organizations have a very natural tendency to generate homogeneity, even absent management intervention. Culture lives a bit a life of its own. This has important economic implications since homogeneity in beliefs and values seems to be a key factor in agency problems.

The next section discusses the phenomenon corporate culture in more depth. It gives some typical examples of culture, reviews key definitions of culture, and proposes the measure that will be used throughout the paper. Section 3 describes the formal model that is the basis of the analysis. Section 4 considers the genesis and evolution of culture. It uses multi-armed bandit theory to discuss the choice of action. It shows that firms will eventually settle on an action, that otherwise identical firms might end up settling on different actions, and that the beliefs of employees are more homogenous within than across firms. Section 5 considers the influence of the manager and of managerial succession, and shows how cultures may persist in the face of complete turnover. Section 6 studies the correlation between culture and performance. Sections 7 through 9 consider some other aspects of culture, potential issues with the model, and testing strategies. Section 10 concludes. The appendices contain a discussion of the use of differing priors and some of the proofs.

## 2 Corporate Culture

When coming into contact with an organization, people are often struck by the fact that members of the organization seem to act and think similarly, but differently from members of similar other organizations. It is as if each organization has its own ‘personality.’ Moreover, this ‘personality’ often remains remarkably constant over time. Even when many of the original members are gone, the new generation thinks and acts in very much the same way as their predecessors. It is essentially this character of an organization, which some have more than others, that has been called its ‘culture.’ Given the rather vague phenomenon, it is not surprising that there are many divergent definitions in the literature. Moreover, as the term became more popular, it also began to live a life of its own. Lazear (1995) provides a survey of alternative definitions. Rather than trying to replicate such survey here, the purpose of this section is to present the view in the management literature on which this paper is based, and compare it to some key alternatives.

### 2.1 Examples of corporate culture

To fix ideas it is useful to start out with simply presenting some examples, drawn from personal experience and case descriptions, of corporate culture.

The first example is a comparison between the Brussels offices of Arthur D. Little and McKinsey in the mid to late nineties. These local offices served similar clients, were started at about the same time and were similar both in size and in personnel composition.

Arthur D. Little’s consultants proudly stated that their firm was an organized chaos or chaotic organization and that it had as many strategies as there are consultants. While formal training existed, every team really went its own way. Data analysis was not so important, but listening to the client was key. Conclusions were often backed up by quotes from clients, or by stories. It was important to have an open mind and not to come too quickly to conclusions. There were very few formatting standards for presentations. Performance evaluations were done every few months via informal 5-minute chats. Every consultant was responsible for his or her own staffing via a market-based system. People took lunch while working in their office. Team lunches were exceptional.

Arthur D. Little called itself ‘the Company.’

McKinsey’s well-developed consulting methodology, on the contrary, guided each study pretty closely. A new study started by collecting all ‘knowledge’ about similar studies that had been conducted in the past in other offices. From the start of the study, consultants were supposed to think in terms of final client recommendations. Any conclusion was backed up by data. There were strict formats for the presentations, decided upon by a global committee of senior directors. There was a clear one-firm policy: the process, rules and systems should be similar all around the globe. Consultants got evaluations every 6 or 12 *weeks*, using extensive and formal evaluation forms. Staffing was centralized and future assignments were chosen to improve on weaknesses. Consultants spent nearly all their time at the client site. Lunch, and often even dinner, were taken as much as possible with the team. McKinsey called itself ‘the Firm.’

These were two firms that were essentially in the same business but worked in very different ways. There were no obvious structural limitations or legacy systems that prevented one to switch to the other’s model. Both were aware of the differences. In fact, these differences in behavior reflect differences in opinions and beliefs among the most senior people about the relative importance of ‘the one best way of doing things,’ individual creativity, teamwork, the most effective process of doing consulting, etc. A telling fact is that new consultants at McKinsey receive a copy of the book *Perspective on McKinsey*, by Marvin Bower, McKinsey’s factual founder. The book is accompanied by a memo from Bower, urging ‘not to give or loan copies to people outside the Firm.’ The book essentially gives Bower’s perspective on the ‘lessons that I believe might be learned from our successes, mistakes, and failures, [...]’ Some McKinsey people refer to it as ‘the bible.’ It seems a conscious effort to influence the beliefs of new employees. The fact that it is explicitly for internal use only, is a clear statement that this is not just posturing for the outside world, but valuable information from which new employees can and maybe should learn.

Cultures also comes in less functional forms. Some companies, such as Enron, encourage their people to be aggressive and push limits, even if it gets them close to legal limits. Other companies, including some government administrations, have implicit cultural beliefs that initiative creates personal risks without any upside. Some firms have a strong ‘nine to five’ culture while in others people always stay late, even if they don’t have anything to do. Cultural differences can also relate to the importance of consensus, the treatment of new employees, the level of confrontation, the level of cooperation, the competitiveness, the implicit importance and status of engineers versus marketers, the existence of reserved parking spots for top management, open or closed door policies, etc.

Note that cultures can also develop along other dimensions than firms. We can talk, for example, about a sales culture versus a production culture, or about the culture of academic economists as opposed to that of academic sociologists or engineers. Each of these groups have a set of common experiences they go through.

## 2.2 Definition of corporate culture

Since culture is a complex social phenomenon, it has multiple dimensions and therefore multiple potential definitions, that all have their value in the right context. While I will discuss some alternative definitions below, I will focus here on the most prevalent definition in the management literature, which is also the basis of this paper. In particular, most sources in the management literature define corporate culture as shared values, beliefs and assumptions that drive behavioral

norms and ‘the way we do things around here.’<sup>5</sup> It should be mentioned up front that most of the management literature considers values and beliefs to be almost interchangeable.<sup>6</sup>

The idea of culture as shared beliefs or values goes back at least to Burns and Stalker (1961) who, in their seminal discussion of ‘organic’ versus ‘mechanistic’ organizations, define culture as ‘a dependable constant system of shared beliefs.’ Other early contributions were the work of Baker (1980) and Schwartz and Davis (1981) who defined culture respectively as ‘some interrelated set of beliefs, shared by most of their members’ and ‘a pattern of beliefs and expectations that is shared by the organization’s members.’ A key impetus in popularizing the notion of culture was the bestseller *In Search of Excellence* by Peters and Waterman (1982), who defined culture as ‘shared values’ but stress that this includes ‘basic beliefs.’ Donaldson and Lorsch (1983), which is often considered a seminal work on corporate culture, do not mention the word culture, but talk instead about managerial beliefs. Most of these authors suggest that a culture can have subcultures.

Probably the most cited perspective on corporate culture is that of Schein (1985). He defines culture as having three levels. The most visible, but most superficial, level is that of culture as a pattern of behavior. It is ‘the way things are done around here,’ the norms, the stories, the symbols. These behavioral patterns reflect a second, deeper, level of culture, which are the firm’s shared values. Shared values are, on their turn, driven by the third and most fundamental level of culture: shared assumptions.<sup>7</sup> Kotter and Heskett (1992) base their definition on Schein (1985), but eliminate the distinction between beliefs and values.

All these definitions essentially consider the norms, stories and symbols as expressions of these shared beliefs or assumptions. A typical behavioral norm would be, for example, whether office doors are open or closed. What seems to be a pure behavioral norm may in fact reflect beliefs about the importance of accessibility, interaction, and group work versus confidentiality, privacy, and getting your work done. This does not mean that the behavioral implications are unimportant. On the contrary, all these authors study culture because of its behavioral implications. But to develop a systematic theory of that behavior, these authors suggest that we must look deeper, to values and beliefs.

**Important alternative definitions** There are a number of important alternative definitions of corporate culture, that differ from the above definition to various degrees. I will discuss the most important:

- Joanne Martin’s 3 perspectives on corporate culture (Martin 1992)
- Culture as pure conventions
- Culture as norms

Martin (1992) identifies 3 perspectives that according to her have dominated research on organizational culture:

---

<sup>5</sup>Although a complete review is outside the scope of this paper, it seems that at least 80% of the management and organizational literature subscribes to the view of culture as shared beliefs and values.

<sup>6</sup>The issue is familiar to economists: you need some fairly strong assumptions to separate beliefs and utilities into expected utility. On a more informal basis, it is easy to see how valuing the environment is closely related to beliefs about where our planet is headed.

<sup>7</sup>Schein’s formal definition is as follows: ‘A pattern of shared basic assumptions that the group learned as it solved problems of external adaptation and internal integration, that has worked well enough to be considered valid and, therefore, to be taught to new members as the correct way to perceive, think, and feel in relation to those problems.’

1. According to the Integration perspective all members of the firm share an organization-wide consensus (in the sense of shared beliefs).
2. The Differentiation perspective explicitly recognizes and focuses on the fact that there are subcultures.
3. The Fragmentation perspective has a ‘focus on ambiguity as the essence of organizational culture.’

The definition in this paper essentially combines the first two perspectives. Some beliefs are shared by the whole organization while others may be shared by specific subgroups only. Elements of both are in fact found in nearly any discussion of corporate culture. The Fragmentation perspective, on the contrary, is more complex. In fact, Martin recognizes that ‘the Fragmentation perspective is difficult to discuss with clarity.’ Martin considers it to be founded in the postmodernist tradition. Some would call it an anti-theory, a theory about the non-existence of theory. Martin explicitly confirms that the Integration perspective, of which this paper is essentially part, ‘has become the dominant view of organizational researchers and practitioners.’

Another important view is that of culture as simple conventions. This is also the first notion of culture in the Kreps-Hermalin model. In the presence of multiple equilibria, culture defines which equilibrium will be followed. When we can drive on the left or on the right, culture defines the expectations and thus the norm. Such culture is very fragile, however. In principle, such conventions can be changed in a split second by mutual agreement. If the United Kingdom legislates tomorrow that cars should drive on the right, does that change the country’s culture by law? Note, more importantly, that this notion is closely related to the notion of culture as shared beliefs. In particular, a necessary condition for having a convention is that there are shared beliefs or expectations about behavior. Second, since payoffs will play an important role in the equilibrium selection process, shared beliefs may play an important facilitating role. Take for example a coordination game discussed by Hermalin (2001) that has two possible equilibria: junior people defer to senior people versus senior people give way to junior people. Hermalin interprets culture as the norm that defines which equilibrium will be followed. Such norm does not develop in a vacuum, however. One would expect the norm to be built on more general beliefs and values in the organization, for example, about the importance of youthful energy and initiative versus the more mature wisdom and balancedness.<sup>8</sup> Note also that culture as language or meaning is closely related to the idea of conventions.

Another closely related idea is that of ‘culture as norms.’ This differs from ‘culture as convention’ in that norms also have some moral element or an idea of punishment. The second notion of the Kreps-Hermalin model can be interpreted along these lines. In particular, while the paper interprets the equilibrium of the ‘trust game’ as *the firm* having a reputation for fairness or reciprocity, an equally valid interpretation of the model is that *the employees* imposed such a norm on the firm and threaten to punish the firm if it does not live up to that norm (i.e., ‘honor our trust, or else’). This notion is also related to that of culture as shared beliefs since norms are usually founded in shared beliefs on what is right and what is wrong.

---

<sup>8</sup>Note that even what seems to be the purest of conventions, which side of the road to drive on, has structural origins. For an extensive study of this ‘convention,’ see Kincaid (1986).



## 2.3 A Measure for Culture

Given the definition of corporate culture as shared beliefs, there are essentially two interpretations for the notion of the ‘strength’ of a culture:

1. Internal homogeneity: the degree to which people within the firm share the same beliefs.
2. External heterogeneity: the degree to which these common beliefs are different from the beliefs of the population at large.

Consistent with most of the literature, I will reserve the term ‘strength’ for the first concept (i.e., for the degree of homogeneity of beliefs within the firm). To denote the second concept, the degree of difference with the population at large, I will use the term ‘distinctiveness.’

This paper uses as the measure of cultural strength the likelihood that two randomly selected employees of the firm agree on the optimal action. The exact expression will be introduced once the model is further detailed. As I will show later, a nice aspect of this measure is that it can be interpreted as the limit probability that two randomly selected employees will ‘do the same thing.’ It is thus very closely related to ‘the way we do things around here.’

There are obviously alternative ways to implement the idea of ‘shared beliefs,’ each with its particular advantages and disadvantages. The ‘right’ measure of culture is probably best based on a theory as to why culture matters. The implicit theory behind the measure that I use here is, loosely speaking, one of coordination between randomly selected members of the firm. The value of this theory relative to others is an empirical matter. Moreover, as culture has multiple facets, there are probably also multiple useful and valid ways of defining it.

## 2.4 The Difficulty of Communicating Information

As mentioned in the introduction, I assume that a firm’s experience is observed only by the firm’s current members and cannot be communicated. It should be uncontroversial that some information is indeed difficult to communicate. Try the following ‘thought experiment.’ Observe a person on the street during one second. Now try to describe that person to someone else in such detail that the other would be able to recognize the person with the same ease as you yourself would. This would take a tremendous amount of time, if it is at all possible.

This idea that some information is difficult to communicate is well captured by Confucius’ ‘Tell me and I will forget, Show me and I will remember, Involve me and I will understand!’ New members of firms (or new members of a profession) are often told to follow a specific methodology. As they discover the merits of the method and start using it independently, they become in fact socialized.

Note that this is essentially a top-down way of communicating. It would be very difficult for an employee to communicate his experience to his boss by making his boss do something in a specific way. This may be one reason why it is so difficult to transfer culture from one organization to another and why it is often said that cultural change has to come from the top.

## 2.5 The role of learning in the formation of culture

Culture in this paper gets developed over time through shared experiences. The management literature mentions this idea explicitly. Take for example Schein’s (1985) definition of corporate culture as ‘a pattern of shared basic assumptions that the group learned as it solved problems of external adaptation and internal integration, that has worked well enough to be considered valid

and, therefore, to be taught to new members as the correct way to perceive, think, and feel in relation to those problems.’ Schein deems the process of learning so important that he includes it in his definition of culture. Donaldson and Lorsch (1983) claim that ‘although the founders’ personal beliefs lie at the heart of the belief system, corporate history also plays an important part in shaping current beliefs. As the founders and their successors manage by their principles, their experiences lead them to modify the system through the process of incremental change.’ While I will discuss below the idea that the founders’ beliefs matter, this quote indicates again the role that the firm’s experience plays in the formation of shared beliefs. Schwartz and Davis (1981) state that ‘culture reflects what has worked in the past.’ In their analysis of how cultural change can occur, Kotter and Heskett (1992) observe that ‘the importance of results cannot be overstated. These new cultures grew in a cycle that was driven by successful results.’ All these remarks suggest that culture is developed to a large extent through joint learning from the company’s experiences.

### 3 The Model

Consider a set of  $F$  firms with  $J$  employees and one manager each. A typical agent is denoted  $(f, i)$ . Agent  $(f, m)$  is the manager of firm  $f$ , while agents  $(f, 1)$  through  $(f, J)$  are regular employees of firm  $f$ .

The firms are involved in an infinitely repeated game. The stage game (in some period  $t$ ) consists of two steps. First, the managers of all firms each simultaneously choose an action  $a_f^t$ , out of a set  $A$ , that their firm will undertake in that period. The action set  $A$  is common to all firms, fixed over the game, and has  $N$  elements  $a_n$ . In the second step, the firm gets a payoff which depends on the action as follows.

Action  $a_1$  is equivalent to doing nothing: it gives a constant payoff of zero in any period to any firm that chooses  $a_1$  in that period. For all other actions, the period- $t$  payoff to firm  $f$  that chose action  $a$  in period  $t$  is  $R_{f,a}^t = \rho_a + \epsilon_f^t$  where  $\rho_a$  is the (unknown) average payoff of action  $a$  and  $\epsilon_f^t$  is a random variation. It is commonly known among all agents that the variations  $\epsilon_f^t$  are independently and identically distributed according a standard normal distribution; i.e.,  $\epsilon_f^t \stackrel{\text{i.i.d.}}{\sim} N(0, 1)$ . The average payoffs  $\rho_a$  are unknown to the agents, but each agent has his own subjective prior beliefs about the  $\rho_a$ ’s. The prior beliefs of agents are commonly known to differ. Agent  $i$ ’s prior belief about  $\rho_a$  is that  $\rho_a$  is normally distributed with mean  $r_{a,i}^0$  and variance one; i.e.,  $N(r_{a,i}^0, 1)$ . The true average payoffs  $\rho_a$  and the means of the priors  $r_{a,i}^0$  for all agents and all actions are given at the beginning of the game. For the purpose of analysis, I will assume that, empirically, the  $\rho_a$  and  $r_{a,i}^0$  are independently and identically distributed according to standard normal distributions:  $\rho_a \sim N(0, 1)$  and  $r_{a,i}^0 \sim N(0, 1)$ . Remember that these are not priors, but empirical distributions that just happen to reflect the distributions in the population.

A firm’s actions and payoffs are observed by the firm’s members and by them only. The objective of the manager of firm  $f$  is to maximize her firm’s discounted payoff. The discount factor equals  $\delta$  for all firms.

It is commonly known that agents have no private information beyond what they observe about the actions of their firm. An agent will not be able to infer anything from the fact that other agents in his firm hold beliefs that are different from his own. Assume moreover, at this stage of the analysis, that agents cannot observe each other’s posterior beliefs or communicate any of their experience to others. Section 6.2 discusses an important modification of this assumption.

As the game has been formulated, the employees are purely observers: they do not undertake any actions. One could assume, for example, that their presence is needed to implement the action but that all their actions are contractible. Section 4.3 shows that this model can be usefully interpreted as the limit case of a situation in which all agents undertake actions and discusses how it can be further generalized to other multi-decision-maker situations.

The fact that employees are just observers has an important implication: similarity in the employees' beliefs is performance-neutral. Culture has no effect on performance in this model.

**The measure for culture** Let  $\tilde{a}_{(f,i)}^t$  denote the action that has the highest payoff according to agent  $(f, i)$  in period  $t$ .

**Definition 1** *The cultural strength of firm  $f$  at the beginning of period  $t$ ,  $S_f^t$ , is*

$$S_f^t = \frac{\sum_{i=1}^J \sum_{j=1}^J I_{\tilde{a}_i^t = \tilde{a}_j^t}}{J^2}$$

where  $I$  is the indicator function.

Implicit in this definition is the idea that coordination between random employees' actions is a key effect of culture. Further research on the role of homogeneity should bring more clarity whether this is indeed the best way to measure culture.

## 4 Basic Analysis

### 4.1 Choice of Actions

A first step in the analysis is to figure out what actions each manager will take. Because the different firms' actions and payoffs are independent and the firms don't observe each others' actions or payoffs, it is sufficient to study one firm in isolation. To that purpose, fix some firm  $f$ . To simplify notation, I will drop all references to firm  $f$ . Whenever there is no risk of confusion, I will also drop the reference to the agent.

The manager's problem is to choose in each period  $t$  an action  $a^t \in A$  to maximize the firm's discounted profits. Taking an action has two effects: a direct contribution to profits and information about the payoff of that particular action. It is thus not always optimal to simply choose the action with the highest expected payoff: the longer term gain from learning about some other action might outweigh the temporarily lower performance. This type of problem is known as the multi-armed bandit problem, which has been studied extensively. In particular, the model corresponds to a multi-armed bandit with  $N$  independent arms. In this case, the well-known result of Gittins and Jones (Gittins and Jones 1974, Gittins 1989, Berry and Fristedt 1985) applies.

To translate the result of Gittins and Jones and some of its key implications to this specific context, note that at the beginning of each period  $t$ , the manager has beliefs about each action's payoff  $\rho_a$  that are a combination of her prior beliefs and the information she has acquired by undertaking action  $a$  in some of the preceding periods. In particular, consider a situation in which the manager's prior was  $N(r_a^0, 1)$ , she has chosen action  $a$  in  $K$  periods, and the payoffs in these periods were  $\{x_k\}_{k=1}^K$ . In that case, the manager's posterior belief (at the beginning of period  $t$ ) is normally distributed with mean  $r_a^t = \frac{r_a^0 + \sum_{k=1}^K x_k}{K+1}$  and variance  $\sigma_{a,t}^2 = \frac{1}{1+K}$ .

Let  $\epsilon_n(s)$  denote the variation  $\epsilon^t$  when the firm samples action  $n$  for the  $s$ 'th time. Let  $(\Omega, S, P)$  denote the probability space generated by these variations  $\epsilon_n(s)$ , the prior beliefs of all agents, and the average performances of all actions. Let  $\omega$  denote a typical element of  $\Omega$ . Note that, given a rule for choosing actions (as I will derive immediately in the proposition) that is based on the mean and variance of these returns, this probability space is equivalent to the one generated by the  $\epsilon_f^t$ . It results in a simpler analysis, however.

Let  $(\mathbf{r}^t, \sigma_{\mathbf{t}}^2) = (r_{a_1}^t, \dots, r_{a_N}^t, \sigma_{t,a_1}^2, \dots, \sigma_{t,a_N}^2)$  denote the manager's current beliefs about the  $N$  actions (including the constant one) at the beginning of period  $t$ .

**Proposition 1** *1. There exists a function  $g$  from  $\mathbb{R} \times [0, \infty)$  to  $\mathbb{R}$ , such that in each period  $t$  the optimal action  $a^t$  satisfies*

$$a^t = \operatorname{argmax}_{a \in A} g(r_a^t, \sigma_{a,t}^2)$$

- 2. There exists almost surely a  $T^*(\omega)$  such that for all  $s, t \geq T^*(\omega)$ ,  $a^s(\omega) = a^t(\omega) = a^*(\omega)$ ; i.e., after a certain time the firm plays the same action forever after.*
- 3. For any action  $a_k$  with  $\rho_k > 0$ ,  $P[a_k = a^*] > 0$ . For any action  $a_k$  with  $\rho_k < 0$ ,  $P[a_k = a^*] = 0$ .*
- 4. With positive probability,  $a^*(\omega) \neq \operatorname{argmax}_{a \in A} \rho_a$ ; i.e.,  $a^*(\omega)$  does not have the highest expected payoff in the set  $A$ . In the limit as  $N \rightarrow \infty$ , any agent knows that there exists almost surely an action with a higher payoff than  $a^*(\omega)$ .*

**Proof :** The first part of the proposition follows directly from the theorem of Gittins and Jones, by the following facts:

- The problem under consideration is a multi-armed bandit with independent arms and geometric discounting. The theorem of Gittins and Jones implies that in this case, the desirability of an arm can be completely characterized by a number (the 'Gittins index') that depends only on the beliefs about that arm (and on the discount factor) and not on the beliefs about any of the other arms.
- The mean and variance are sufficient statistics for the posterior beliefs of the manager about any arm at any point in time.

The function  $g$  in the theorem projects the (sufficient statistics of) the manager's beliefs about a particular action into its corresponding Gittins index. The optimal choice among the actions is then the action with the highest index.

The second part of the proposition adapts the proof of Rothschild (1974) to the current context. Since the proof is long and does not provide much insight in the analysis of this paper, it is relegated to the appendix. The rest of the proposition builds on these results and is therefore also treated in appendix. ■

The first part of the proposition says that we can characterize the attractiveness of an action by a number that depends only on the action's characteristics and not in any way on the other actions. Stated in a different way, there exists a complete (non-strict) ordering of all possible alternative actions. This simplifies enormously the problem of choosing the optimal action. The calculation of these 'Gittins indices,' however, remains a difficult task. For the model under consideration no analytically closed expressions are known, but it is possible to derive the indices numerically (Gittins 1989).

The second part of the proposition says that a firm will eventually settle on an action and play that action forever. The intuition behind this argument is very simple. First of all, with only a finite number of actions and an infinite number of periods, at least one actions must be used infinitely often. Second, assume (by contradiction) that there is more than one action that is used infinitely often. After sufficient time, the expected payoffs of both actions are know with arbitrary precision. But then there is no gain any more from experimenting, and one can simply choose the higher-payoff action forever. I will call this action the firm’s ‘eventual action’  $a^*$ . This will thus also be the firm’s eventual culture.

The third part of the proposition shows that any action with positive returns may end up being the eventual action. It will lead (in the following part of the proposition) to the important observation that firms in identical circumstances will likely end up having different ‘eventual cultures.’ Moreover, it shows that these cultures may have very different performance levels.

The last part of the proposition says that the eventual action is not necessarily the one with the highest payoff. Even worse, the very last part of the proposition says that, with many alternative actions, the firm’s agents will be almost sure that there are better ways of doing things, but stick to their current course of action anyways. This might explain the persistence of performance differences among firms that can observe each other’s performance.

## 4.2 Emergence of culture

Members of the same firm have identical experiences, which differ from these of members of other firms. As a consequence, employees of the same firm will tend to agree more on the optimal action than randomly selected members of society.

To see this more formally, consider two firms  $f$  and  $g$ . Pick two employees,  $i$  and  $j$ , of firm  $f$  and one employee,  $k$ , of firm  $g$ . Let again  $\tilde{a}_i^t$  denote the action that has the highest payoff according to agent  $i$  in period  $t$ .

**Proposition 2** • For any period  $t$ ,

$$P[\tilde{a}_i^t = \tilde{a}_j^t] \geq P[\tilde{a}_i^t = \tilde{a}_k^t]$$

*i.e., employees of the same firm are more likely to agree on the optimal action than employees of different firms.*

- As  $N \rightarrow \infty$ ,  $P[a_f^* = a_g^*] \rightarrow 0$ ; *i.e., the probability that one firm’s eventual culture is different from that of any other firm converges to one.*

**Proof :** Consider the first point of the proposition. It is sufficient to show that for any particular action, the probability that 2 employees of the same firm consider that action the best is higher than for 2 employees of different firms. To that purpose, pick any action, say  $\hat{a}$ . Let employees 1 and 2 be employees of the same firm, while 3 is an employee of a different firm. In this whole analysis, I will drop the reference to the time period ( $t$ ), in order to simplify notation.

Let  $X_{\hat{a},1} = \{\omega : r_{\hat{a},1} \geq r_{a,1}, \forall a \in A\}$  denote the event that employee 1 considers  $\hat{a}$  to be the best action. It is sufficient to show that  $P[X_{\hat{a},1} \cap X_{\hat{a},2}] \geq P[X_{\hat{a},1} \cap X_{\hat{a},3}]$  or  $P[X_{\hat{a},1} | X_{\hat{a},2}]P[X_{\hat{a},2}] \geq P[X_{\hat{a},1} | X_{\hat{a},3}]P[X_{\hat{a},3}]$ . Since ex-ante all employees are symmetric, it must be that  $P[X_{\hat{a},2}] = P[X_{\hat{a},3}]$ , so that it is sufficient to show that  $P[X_{\hat{a},1} | X_{\hat{a},2}] \geq P[X_{\hat{a},1} | X_{\hat{a},3}]$ .

Conditional on the  $\rho$ ’s, there is no information in  $X_{\hat{a},3}$  that is relevant to  $X_{\hat{a},1}$ . It is thus sufficient to show that, conditional on the  $\rho$ ’s,  $P[X_{\hat{a},1} | X_{\hat{a},2}] \geq P[X_{\hat{a},1}]$  or  $P[X_{\hat{a},1} \& X_{\hat{a},2}] \geq P[X_{\hat{a},1}]P[X_{\hat{a},2}]$ . Note that both employees are now from the same firm. So I will drop the index referring to the employee whenever a variable is common to all employees in the same firm.

The inequality  $r_{\hat{a},1} \geq r_{\bar{a},1}$  can be written (with  $K_{\hat{a}}$  denoting the number of times  $\hat{a}$  has been selected by the management of the firm)

$$\frac{K_{\hat{a}}\rho_{\hat{a}} + \sum_j^{K_{\hat{a}}} \epsilon_{\hat{a}}^j + \epsilon_{\hat{a},1}^0}{K_{\hat{a}} + 1} \geq \frac{K_{\bar{a}}\rho_{\bar{a}} + \sum_j^{K_{\bar{a}}} \epsilon_{\bar{a}}^j + \epsilon_{\bar{a},1}^0}{K_{\bar{a}} + 1}$$

or

$$\frac{K_{\hat{a}}\rho_{\hat{a}} + \sum_j^{K_{\hat{a}}} \epsilon_{\hat{a}}^j}{K_{\hat{a}} + 1} - \frac{K_{\bar{a}}\rho_{\bar{a}} + \sum_j^{K_{\bar{a}}} \epsilon_{\bar{a}}^j}{K_{\bar{a}} + 1} \geq \frac{\epsilon_{\hat{a},1}^0}{K_{\hat{a}} + 1} - \frac{\epsilon_{\bar{a},1}^0}{K_{\bar{a}} + 1}$$

Conditional on  $K_{a_n}$  and  $\rho_{a_n}$ , the RHS is the sum of two normally distributed random variables, which I will denote  $n_{\bar{a},1}$  and  $n_{\hat{a},1}$  respectively. The LHS is also a random variable, which I will denote  $x_{\bar{a}}$ , since it is not dependent on the identity of the employee. So the inequality then becomes  $x_{\bar{a}} \geq n_{\bar{a},1} + n_{\hat{a},1}$ . The event  $X_{\hat{a},1}$  corresponds to  $N$  such equations (if I include the trivial one). Analogously,  $X_{\hat{a},2}$  corresponds to  $N$  analogous equations  $x_{\bar{a}} \geq n_{\bar{a},2} + n_{\hat{a},2}$  where  $x_{\bar{a}}$  is identically the same random variable as in  $X_{\hat{a},1}$ , while  $n_{\bar{a},2}$  is independently and identically distributed as  $n_{\bar{a},1}$  and idem for  $n_{\hat{a},1}$  and  $n_{\hat{a},2}$ .

Let  $f$  denote the joint density of the LHS's. Let  $F_{\hat{a}}$  denote the distribution of  $n_{\hat{a},1}$ , and thus also of  $n_{\hat{a},2}$ . We then get

$$\begin{aligned} & P[X_{\hat{a},1} \& X_{\hat{a},2}] \\ &= \int \cdots \int P[x_{a_1} \geq n_{a_1,1} + n_{\hat{a},1}, \dots, x_{a_N} \geq n_{a_N,1} + n_{\hat{a},1}, x_{a_1} \geq n_{a_1,2} + n_{\hat{a},2}, \dots, x_{a_N} \geq n_{a_N,2} + n_{\hat{a},2}] \\ & \quad f(x_{a_1}, \dots, x_{a_N}) da_1 \cdots da_N \\ &= \int \cdots \int \left[ \int F_{a_1}[x_{a_1} - n_{\hat{a},1}] \cdots F_{a_N}[x_{a_N} - n_{\hat{a},1}] dn_{\hat{a},1} \right]^2 f(x_{a_1}, \dots, x_{a_N}) da_1 \cdots da_N \\ &\geq \left[ \int \cdots \int \left[ \int F_{a_1}[x_{a_1} - n_{\hat{a},1}] \cdots F_{a_N}[x_{a_N} - n_{\hat{a},1}] dn_{\hat{a},1} \right] f(x_{a_1}, \dots, x_{a_N}) da_1 \cdots da_N \right]^2 \\ &= P[X_{\hat{a},1}]^2 \\ &= P[X_{\hat{a},1}]P[X_{\hat{a},2}] \end{aligned}$$

which thus proves the first part of the proposition.

*For the second and last point,* note that for a set of actions  $A$  with a given set of corresponding average performances  $\rho_a$ , the event that firm 2 ends up having action  $\bar{a}$  as its eventual course of action is equally likely but independent from the event that firm 1 ends up having action  $\bar{a}$  as its eventual course of action. From this and the earlier conclusion that the probability of any particular action being a firm's eventual action converges to zero as  $N \rightarrow \infty$ , it follows that the probability that both firms have the same action as their eventual course of action converges to zero as  $N \rightarrow \infty$ . ■

The first part of the proposition says essentially that a culture will develop in each firm: the beliefs within a firm become more homogenous than throughout society at large. If we defined a cultural strength for society, the proposition would say that a firm's culture is stronger than that of society.

The last part of the proposition explains the important stylized fact that structurally similar firms may end up with very different cultures (and thus very different performance). This is especially likely when there are many different ways to get things done.

So this part of the analysis establishes the first essential point of the paper: joint experience will lead to shared beliefs. Culture emerges, in the sense of culture as shared beliefs. I will try to

validate the model as much as possible by considering how well it fits some of the stylized facts. In fact, the last proposition did already some of that by showing that similar firms may develop different cultures, and that firms may stick to a culture that they know is suboptimal. Section 5 will do more of that by showing how culture is influenced by the manager and how culture may persist over time. But first, I want to discuss what happens when there are multiple decision makers and what the effect of turnover would be.

### 4.3 Multiple decision makers

An apparent limitation of the model is the fact that only the manager takes actions. The other employees are passive observers. As mentioned earlier, however, the model can be interpreted as the limit case of a model with all employees undertaking actions.

Consider in particular the following extension of the original model. Let all employees undertake actions in all periods. Only the outcomes of the manager's actions are observed. The firm's payoff is  $\beta$  times the payoff from the manager's action plus  $(1 - \beta)$  times the average payoff from the employees' actions. The manager and the employees all care about the firm's total profits.<sup>9</sup>

Consider now the limit as  $\beta \rightarrow 1$ . In this case, the manager will not change her actions from before. Employees will in each period simply choose the action that they think has the highest payoff. The following proposition says that employees' actions will be more similar within than across firms.

**Proposition 3** *Two employees of the same firm are more likely to undertake the same action than two randomly selected employees.*

**Proof :** Since employees can't learn from their actions, they simply choose the action they believe has the highest payoff. The proposition then follows immediately from proposition 2. ■

The key results of the paper seem to hold in fact in a lot more generality, including cases where all agents' actions are observed. Such models, however, are essentially multi-person multi-armed bandits, which bring us far outside the scope of this paper.

Note that this multiple decision maker perspective provides a nice alternative interpretation for this paper's measure of culture: in the limit, it is the probability that two randomly selected employees 'do the same thing.' This is thus very close to 'the way we do things around here.'

### 4.4 Turnover and culture

A key concern in this context should be the effect of turnover on culture. At first sight, it seems that turnover would weaken the culture, since new employees do not possess the same information as the old ones. The proof of proposition 2 would imply that with less shared information, the culture should be weaker.

The situation is much more complex, however. For one thing, new employees do not really start from scratch: they will learn from the firm's earlier experience in at least 3 ways.

First, they will deduce information from the firm's actions. If the firm undertakes  $\hat{a}$ , that means that it probably had good experiences with  $\hat{a}$  in the past. The exact inference, however, depends on all the information the new employee has, including, for example, how long the firm has been in operation. Three things are important to note. First, this inference process will usually create

---

<sup>9</sup>The results would be the same if I assumed that the employees cared only about the profit from their own activities.

homogeneity because it aligns the new employees' beliefs with those of the existing employees. Second, the employee will revise this deduced information with each subsequent observation. If, for example, the firm switches next period to a different action than that has different implications for the prior experience than if the firm sticks to the same action. Third, as already suggested by the point above, the inference process is tremendously complex and analytically prohibitive.

A second way how a new employee may learn from the past experience of the firm is by (imperfectly) observing the beliefs of his colleagues. In fact, the inferences from the firm's actions are a special case of this.

Finally, and most directly, the employee can learn from earlier experience through stories from his colleagues about the past successes and failures of the firm. In fact, the sociological literature often interprets the symbols and stories that are part of the culture as essentially this: communicating to new members the fundamental beliefs and values of the firm through rich stories and anecdotes. The new employee has to be aware, however, of the biases in such stories: his colleagues are likely to tell the stories to try to influence him, typically towards their own beliefs.

Not only can a new employee learn from his older colleagues, but there is likely also sorting going on in the hiring process. In particular, firms are likely to hire employees with beliefs that are similar to their own (Schein 1985, Van den Steen 2002a). This would further increase homogeneity in the face of turnover.

Finally, the way how turnover affects the firm's optimal level of investment in socialization (assuming culture is useful) is ambiguous: with new incoming members the effect of socialization is likely to be higher, but the payoff from such investments are reduced when employees stay shorter.

All in all, the balance is complex and the effect of turnover is a complete research topic on itself. I do believe that overall turnover is likely to weaken the culture. But it should not be overlooked that some firms with extremely high turnover, such as consulting firms, manage to build extremely strong cultures.

## 5 Managers and Culture

### 5.1 The Manager's Direct Influence on Culture

The management literature suggests that an organization's culture is influenced by the beliefs of founders and early managers (Donaldson and Lorsch 1983, Schein 1985, Kotter and Heskett 1992, Baron et al. 1999). This section considers this effect in the context of the current model.<sup>10</sup> The first proposition shows that the manager's prior has indeed an important influence on the firm's eventual behavior. Let  $a^*$  denote a firm's eventual action or culture, where I drop the dependence on the state  $\omega$ .

**Proposition 4** *For any action  $\hat{a}$ , the probability that  $a^* = \hat{a}$  increases in  $r_{\hat{a}}^0$ , the mean of the manager's prior about the payoff of  $\hat{a}$ .*

**Proof :** This follows from lemma 6. ■

---

<sup>10</sup>An alternative perspective of a manager's influence on culture is presented in Van den Steen (2002d). In that model, a manager's strong beliefs or vision cause sorting in the labor market, which leads to a homogeneity in the firm's beliefs.



## 5.2 Managerial Succession

The last subsection makes clear that the manager has an important influence on the firm’s eventual culture. This raises the issue what happens when the firm changes managers. Accounts in the business press and management literature suggest that change is more likely when the new CEO is an outsider.

To study this issue more formally, consider the following extension of the original setup. At the end of each period, there is a small probability  $p$  that the manager will be replaced by a new one. I will compare the case that the new manager is an insider (i.e., one of the firm’s current employees) with the case that that new manager is an outsider (i.e., some randomly drawn individual from another firm). Assume that insiders know the whole firm’s history in terms of actions and outcomes, while outsiders have no information about the firm’s prior actions but do have the experience as an employee of another randomly drawn firm. Each manager maximizes the discounted payoff under his or her management. Let  $a_1^*$  and  $a_2^*$  denote respectively the last action under the first manager and the eventual action under the second manager.

**Proposition 5** *As  $N \rightarrow \infty$  and  $p \rightarrow 0$ , the probability that  $a_1^* = a_2^*$  converges to zero for an outsider-successor, but to some value weakly larger than  $1/2$  for an insider-successor: An eventual change in culture is more likely under an outsider successor than under an insider successor.*

**Proof :** Note first that, for any  $\epsilon > 0$  and as  $p \rightarrow 0$ , the time it takes to succession is almost surely larger than  $T^*$  plus any time needed to get the belief regarding  $\rho_{a^*}$  within  $\epsilon$  of the true value.

Consider now the case of the insider successor, denoted  $I$ . Consider first the modified problem that  $I$ ’s choice set is restricted to  $A \setminus a_1^*$ . Let  $\hat{a}$  denote  $I$ ’s eventual action in this modified problem. In the unmodified problem  $I$ ’s eventual action  $a_2^*$  must then be either  $a_1^*$  or  $\hat{a}$ . Note that  $I$ ’s eventual action can be  $\hat{a}$  only if  $\rho_{a_1^*} \leq \rho_{\hat{a}}$ . Furthermore, by lemma 7,  $P[\rho_{a_1^*} \geq \rho_{\hat{a}}] \geq 0.5$ . So it follows that the probability that  $a_2^* = a_1^*$  must be weakly larger than  $1/2$  in the limit.

Consider next the outsider-successor. Since there is no information transfer, it is as if we consider a similar succession in a completely different firm. An immediate extension of proposition 1 implies then that as  $N \rightarrow \infty$ , the probability of the eventual actions being identical goes to zero. ■

Simulations suggest that this effect can be really large. Table 1 gives the results from 5 simulations of 200 firms each, in which one manager runs the firm for 150 periods and then gets succeeded by either one of her employees or an outsider who runs the firm for another 150 periods. The numbers in the table represent the percentage of cases in which the new manager eventually chose the *same* action as the old manager. Clearly, outsiders are *much* more likely to change the culture than insiders.

The effect is essentially one of forced learning: during her tenure, the original manager implicitly chooses what her successor learns. When the successor takes over, he might try out a few changes, but if he doesn’t quickly find an action that performs really well, he will fall back on the proven strategy of his predecessor. Combined with the earlier result, this implies that a manager’s beliefs may determine a firm’s culture even after the manager is gone.

## 5.3 Turnover and Culture (2)

One of the striking things about corporate culture is indeed its persistence over time. In particular, even after all the original members are long gone, the firm may still have the same shared beliefs.

Simulation	Insider	Outsider	Avg. % Performance difference
1	74	15	21
2	69	13	20
3	72	18	17
4	71	18	18
5	70	14	22

Table 1: Percentage of cases in which the action at the end of 150 periods under the successor-CEO is identical to the action at the end of 150 periods under the original CEO, and the average percentage performance difference. The data represent 5 simulations of 200 firms each. The number of actions was 100.

The above analysis on managerial succession suggests one reason why that may happen. The simulation results in table 2 put some more flesh on this assertion. They show the outcomes from the following simulation, run five times. A manager runs the firm for 150 periods. At that point one of his employees succeeds him, while all other employees get replaced by new ones who did not observe the firm’s earlier history and who do not make inferences about it.<sup>11</sup> After another 150 periods, the second manager is succeeded by one of her employees. The results in the table indicate for 5 simulations the percentage of cases (over a sample of 100 firms) that the third and last manager settles on the same action as the very first one.

For the purpose of the argument it is important to note that at the end of the simulations, none of the original employees is still with the firm. Neither the third manager nor any of his employees ever worked with the first manager. Nevertheless, their behavior is still very much determined by that original manager’s beliefs.

To control for the fact that this might simply be convergence on an optimal action, I ran the simulation simultaneously for two firms and report in the second column the percentage of cases that the two firms end up settling on the same action. If the results were due to a convergence on the optimal action, then the numbers should be similar. It is clear from the comparison that there is persistence of culture that is not due to a convergence to the optimal action. Moreover, the average performance difference between the firms, reported in the third column, is substantial.

Simulation	Culture Persistence	Persis- tence	Identical to other firm	Avg. % Performance Difference
1		59	10	34
2		50	14	18
3		52	12	15
4		67	16	15
5		55	15	13

Table 2: Percentage of cases in which the eventual action under the final manager is identical to the eventual action under the original manager, and percentage that the eventual action is identical to that of some other firm. The data represent 5 simulations of 100 firms each. The number of actions was 300.

---

<sup>11</sup>I make this assumption for analytical reasons. Taking into account such inferences would actually strengthen the result.

## 6 Culture and Performance

The interest in the phenomenon of corporate culture has been driven to a large extent by its potential impact on corporate performance. In particular, works such as these of Deal and Kennedy (1982), Peters and Waterman (1982), or Collins and Porras (1994), have popularized the notion that culture is a key driver of strong performance. It is unclear, however, whether this popular notion is really correct. The case studies in these management books are far from random since they explicitly selected ‘excellent performers.’ This leads to selection bias.<sup>12</sup> More systematic evidence is rare. In fact, the three more systematic studies that I encountered (Kotter and Heskett 1992, Burt, Gabbay, and Holt 1994, Sørensen 2002), all use the same data set. Since there are some important issues with this data set (and with some of the analyses, see Lazear (1995)), this evidence is far from conclusive. Nevertheless, the data combined with the case studies make a suggestive case for a correlation between culture and performance.

Probably the most interesting result of this paper is the fact that such correlation is in fact to be expected, *even when culture has no positive effect on performance*. The fact is that firms with good performance will develop a strong culture. There are at least two mechanisms, which I discuss immediately, through which this may happen.

1. The extreme-value effect
2. The communication or socialization effect

The intuition for the extreme-value effect is as follows. Consider a firm which has stumbled upon a high-performing action. All its employees will observe how well the action performs. Many will become convinced that this is indeed the optimal action. This results in a high level of agreement and thus a strong culture. In a firm that hit only upon mediocre actions, on the contrary, disagreement as to what is the best action will persist and thus culture will be weak.

The communication or socialization effect is even more transparent but requires a slight extension of the setup. Up to this point, I have assumed that all employees of the firm perfectly observe all realizations. This will generally not be the case. In particular, the flow of information may depend on the firm’s investment in communication and socialization systems. Firms that have discovered high-performing actions will invest more in such systems, because the upside from communication is larger. Analogously, employees will be more eager to learn about how things work when they see that the firm’s methods are very effective. As a consequence, there will be more homogeneity in beliefs about the optimal action in firms that perform better.

Both these effects seem so straightforward that one may wonder about the gains from a formal analysis. The opposite turns out to be true. While the analysis quickly becomes very involved, it shows clearly that the situation is more complex than it seems.

The main issue is the following. What matters for the level of agreement is the difference in perceived performance between the highest and any subsequent actions. The simple intuition would be that this difference increases as the value of the higher one increases. This is not necessarily the case. Consider, for example, two random variables that are independently and identically

---

<sup>12</sup>While there is a clear selection bias in Deal and Kennedy (1982) and Peters and Waterman (1982), the problem is better hidden in the case of Collins and Porras (1994). These authors claim that they have eliminated the bias since they included randomly selected firms with average performance that otherwise matched some key characteristics of the original firms. It is easy to see, however, that including such firms does not eliminate the issue that the results may be driven by an increase in the variance rather than in the mean. This kind of selection is especially important since strong cultures may increase the variance of the firm’s results.

distributed according to the density in figure 1. In this case, the expected difference between the higher and the lower does not always increase in the higher value. Conditional on the first-order statistic being 10, the expected value of the second-order statistic is 5, giving a difference 5. When the first-order is 11, however, the expected value of the second-order increases to 10, for a difference of only 1. This contradicts the ‘simple’ intuition that the difference between best and second-best increases as the value of the best increases.




Figure 1: Density for which the expected difference between first and second order statistic does not always increase when the value of the first order statistic increases.

This may seem like an extreme case. Note, however, that qualitatively the same kind of effects play with the normal distribution when the highest value is close to (or smaller than) the mean. So we should not be surprised to see this effect play a real role. Indeed, the simulations below suggest that this might be happening for smaller values of  $\rho$ .

There is a second, mainly analytical, issue. While the random variables are ex-ante independent, they become ex-post dependent because the action choices depends on the earlier realizations. The ex-post distribution of beliefs about, for example,  $a_2$  will depend on the realization of the other actions. If  $\rho_{a_3}$  and  $r_{a_3}^0$  are extremely high, then it is unlikely that  $a_2$  will be ever selected. The density of  $a_2$  will be approximately the prior. If, on the contrary, all actions other than  $a_2$  have extremely low values for the priors, then  $a_2$  may be selected forever (or until it happens to generate a very negative result). In fact, the ex-post joint density is analytically completely intractable.

## 6.1 The extreme-value effect

Nevertheless, the extreme-value effect does hold for ‘large’ values. To avoid the issue of dependency among the random variables, the following proposition proves the effect for the situation immediately after period 1. To simplify the analysis, I will assume, in this and what follows, that  $a_1$  is an action like all others; i.e., it does not have a constant payoff any more. Let  $R$  denote the performance in the first period, and remember that  $S_2$  denotes the expected strength of the firm’s culture at the start of period 2.

**Proposition 6** *There exist  $\hat{R}$  such that  $S_2$  increases in the first period’s performance  $R$  for  $R \geq \hat{R}$ .*

**Proof :** To simplify notation, I drop the superscript for the period. All beliefs are prior beliefs. Let wlog. the action selected in the first period be  $a_N$ . Randomly select an employee  $i$  and denote his prior beliefs by  $(a_n)_{n=1}^N$ . Note that conditional on  $a_N$  being selected the prior beliefs of employees are still independent normal variables since it was the manager and not the employees who selected the action. Let

$P$  denote the probability that this particular employee considers  $a_N$  the best action at the start of period 2.

$$P = P \left[ r_{a_n} \leq \frac{r_{a_N} + R}{2} \forall n \right] = \int P \left[ r_{a_n} \leq \frac{x + R}{2} \forall n \right] f(x) dx = \int F^{N-1} \left( \frac{x + R}{2} \right) f(x) dx$$

which increases in  $R$  and converges to 1 as  $R \rightarrow \infty$ .

Note now that, by symmetry for all actions but  $N$ , the probability that the employee considers (at the start of period 2) any action other than  $a_N$  the best action, is simply  $\frac{1-P}{N-1}$ . Since all employees are identical and their prior beliefs are independent, the probability that two randomly selected employees agree on the optimal action is  $P^2 + (N-1) \left( \frac{1-P}{N-1} \right)^2$  which has derivative for  $P$

$$2P - \frac{1}{N-1} 2(1-P) = \frac{2}{N-1} ((N-1)P + P - 1) = \frac{2(NP - 1)}{N-1}$$

which is positive for  $P \geq \frac{1}{N}$ .

This implies that there exists a  $\hat{R}$  such that for  $R \geq \hat{R}$ , the probability that two agents agree on the optimal action increases in  $R$ . This concludes the proof. ■

The extreme effect is actually pretty strong. Figure 2 shows the results of a typical simulation. In




Figure 2: The extreme value effect. Performance in function of cultural strength (in probability) for a simulation of 100 firms with 100 employees each, that operated for 150 periods.

this case, I let 100 firms operate for 150 periods and then plot their cultural strength (considering 100 employees) in function of their performance, both in the last period. Above a performance of 2, there is clearly a very strong correlation with the strength of culture. The overall correlation between cultural strength and performance in this particular simulation was 0.59, and is typically between 0.5 and 0.65. These are very high numbers, as can be expected from the graphical result. Note also from the figure that up to a performance of 2, the correlation is actually slightly negative. This is true for nearly all such simulations. It suggests that the counterexample of figure 1 is not a pathological case, but a real phenomenon.

## 6.2 The communication or socialization effect

To study the communication effect, I need to extend the model to allow the spread of information to depend on the firm's investments. To that purpose, consider the following modifications.

- Employees originally do not observe the manager's actions and the outcomes of these actions.
- At some point in time, the manager can invest in a socialization and communication system, at a cost  $(1 - \beta)c$ . The details will be discussed below.
- As in Section 4.3, each employee takes an action  $a \in A$  in each period. The outcomes of the actions are not observed. The profit of the firm is  $\beta$  times the profit from the manager plus  $(1 - \beta)$  times the average profit from the employees. Both the manager and the employees try to maximize firm profit.<sup>13</sup>

Consider the case  $\beta \rightarrow 1$ . The manager still takes actions as before, while the agents just choose in each period the action they consider best.

There is a complicating factor in this analysis: how to separate analytically the communication effect from the extreme-value effect? From the analysis in Section 4 it followed that more investment in a communication system will strengthen the culture. Hence, the best solution seems to focus directly on the link between performance and the level of investment in communication. This is what I will do here.

To avoid again the problem of ex-post dependency, consider the following simple case. Assume that after the first period, the manager can communicate the nature of the first action and its outcome at a cost  $(1 - \beta)c$ . No further information can or will be communicated. The manager knows the distribution of beliefs among employees but not the precise beliefs of the employee to whom he considers communicating.

**Proposition 7** *There exists a  $\hat{R}$  such that the probability of investment increases in the first realization  $R$  for  $R \geq \hat{R}$ .*

**Proof :** For notational convenience, I will drop the superscript references to the period (all values will be prior beliefs) and denote the manager as agent  $m$  and the employee as agent 1. Assume wlog. that the manager chooses action  $a_1$  in the first period and that the result is  $R$ .

Let  $V^0$  denote the expected per-period payoff from the agent's action (from the manager's perspective) when the manager does not communicate any information. Since all actions are equally likely (from the manager's perspective) without communication, that value is

$$V^0 = \frac{r_{a_1,m} + R}{2N} + \sum_{i=2}^N \frac{r_{a_i,m}}{N}$$

Let  $P(R)$  be the probability (from the manager's perspective) that the employee will undertake  $a_1$  when he hears that the first return is  $R$ . Note that  $P(R) = P\left[\frac{R+r_{a_1,1}}{2} \geq r_{a_i,1} \forall i\right] = G^{N-1}(R)$  where  $G$  is the distribution of  $2r_{a_i,2} - r_{a_1,2}$ . The function  $P(R)$  is thus strictly increasing and converges to 1 as  $R \rightarrow \infty$ . Let  $\hat{R}$  be such that  $P(\hat{R}) \geq \frac{1}{N}$ .

---

<sup>13</sup>Again, the results would be unchanged if I assumed that employees only cared about the payoffs from their own actions.

Let  $V^1$  denote the expected per-period payoff from the agent's action (from the manager's perspective) when the manager does communicate the information. We then have that

$$V^1 = \frac{r_{a_1,m} + R}{2}P(R) + \sum_{i=2}^N \frac{r_{a_i,m}}{N-1}(1 - P(R))$$

so that

$$V^1 - V^0 = \left( \frac{r_{a_1,m} + R}{2} - \sum_{i=2}^N \frac{r_{a_i,m}}{N-1} \right) \left( P(R) - \frac{1}{N} \right)$$

So the probability of communication is

$$P \left[ \left( \frac{r_{a_1,m} + R}{2} - \sum_{i=2}^N \frac{r_{a_i,m}}{N-1} \right) \left( P(R) - \frac{1}{N} \right) \geq c \right]$$

where all but  $R$ ,  $N$ , and  $c$  are random variables. When  $R > \hat{R}$  so that  $P(R) - \frac{1}{N} > 0$ , this increases in  $R$ . This proves the proposition.  $\blacksquare$

A simulation can again give an impression of the impact of this phenomenon. In the simulation below, I let the manager decide at the end of period 150 whether he installs a communication system or not. The communication system will communicate all future actions and realizations. In making the decision, the manager considers his beliefs, knows the distribution of prior beliefs of his employees (but not their specific beliefs), and assumes that employees do not make any inferences about earlier realizations (for example, because they are not aware that the firm operated earlier). The analysis determines the expected gain from communicating future information, assuming that by period 150 all firms have settled on their eventual course of action. Figure 3 shows how the expected gain from communication is related to the performance. There is clearly a strong positive correlation, which for this particular simulation is 0.82 and typically exceeds 0.8. Note also that for relatively low-performing firms, communicating information *reduces* expected performance from the manager's perspective. The reason is that, when the manager's actions have a mediocre performance, employees will believe they know better and undertake other actions, which the manager considers inferior. In this case, there does not seem to be a negative correlation for low levels of performance, although the correlation is clearly lower than for high levels of performance.

## 7 Other aspects of culture

There are many important aspects of culture that I have not dealt with in this paper, but for which the model could have interesting implications.

**Evolution of culture** Some culture researchers have suggested that older organizations may have stronger cultures. It sounds intuitive that this should indeed be the case in the current model. The fact that the ex-post distributions are unknown, complex, and highly interdependent makes an analytical approach to this question very difficult. Simulations clearly suggest, however, that culture gets indeed stronger over time in this model.




Figure 3: The communication effect. Performance in function of the gain that the manager expects from communication. The simulation consisted of 200 firms with 100 employees each, that operated for 150 periods.

**Culture and Change** The management literature has suggested that a strong culture may hinder change. While there might be direct causal links in more complex models (that have, for example, coordination problems), the current paper suggests that there might also be non-causal effects. In as far as the strong culture is due to extremely high outcomes in the past, the employees might have very strong beliefs that the current course is the right one and be slow to adapt to changed circumstances.

**Culture and variability** Sørensen (2002) argues that a strong culture should reduce the variability of performance and shows that this holds indeed in the data set of Kotter and Heskett (1992). The current model suggests at least one causal link for such relationship. If employees are randomly selected to undertake actions, more homogenous beliefs about what to do will definitely cause more systematic behavior. It would be interesting to see whether there might also be non-causal effects.

## 8 Discussion of the model

There are clearly some aspects of the model that raise questions. The purpose of this section is to try to deal with some of the most prominent ones.

A key assumption of the model is that the outcomes are observed only by the members of the firm. The idea is that information that can be easily communicated to other firms will get communicated. It can thus not be the basis for lasting differences such as culture. Making this extreme assumption makes things more transparent. The qualitative results would probably not change if I were to allow communication at a cost or noisy observations.

The assumption of differing priors, which is discussed in more detail in appendix, is not necessary for the conclusions, as long as employees cannot (perfectly) observe each other's beliefs and get some private signals. The complexity of the analysis would increase enormously, however.



The measure of culture is the probability that two randomly selected employees agree on the optimal action but does not consider the beliefs of the manager. While this definition seems close to the management literature and is necessary for comparative statics on managerial beliefs, there are also analytical reasons to exclude the manager. Since the firm's course of action is the one that the manager most prefers, his beliefs about that course of action are distributed differently than those of the other employees, and depend on the number of available actions. Including the manager adds a whole other set of considerations and makes it difficult to obtain analytical results. It does not seem that any of the results would be qualitatively different if I included the manager.

The model assumes that employees don't move between firms. Note first that allowing employees to move between firms raises the issue on what inferences they can draw about their new firm's past experiences by, for example, observing the firm's actions. Since I assumed that direct communication of information is impossible, however, there seem to be no real effects from an employee moving to another firm, unless that employee becomes the manager of the other firm. This is because regular employees are passive observers in this model. This happens to fit well with the idea that any real change in culture must come from the top (Donaldson and Lorsch 1983, Schein 1985, Kotter and Heskett 1992).

Note that the model leaves out many aspects that would actually make it easier to obtain the conclusions of the paper. In particular, the following elements would typically strengthen the conclusions:

- Sorting effects.
- Investments (by the firms and by the employees) that are complementary to the action choice.
- Any kind of interaction among employees that would create value from homogeneity. A need for coordination would be an example.

The reason to leave them out of the analysis is both to simplify the model and to show that such effects are not necessary to obtain the results.

## 9 Testing strategies

The most direct test of the theory would try to measure corporate culture and then test the different aspects of the theory: how culture is influenced by the manager's beliefs, whether strong performance leads to a strong culture, whether culture indeed persists over different generations of employees, etc. The difficulty is of course to find a way to measure culture (and a manager's beliefs). Directly questioning employees and managers about the strength of culture has very limited validity, but may be a useful first step. Much better would be to question employees and managers on their opinions about optimal actions or to observe their actions, assuming these actions are really their own choice and not dictated by, for example, authority or incentives.

There are also some indirect empirical tests, although the potential for confounding effects is larger. Some of these indirect implications are:

- Firms change more under outsider-successors to the manager than under insider-successors.
- People who worked a substantial time for the same organization tend to act in similar ways. This could be tested in the style of Bertrand and Schoar (2001).

- Firms that exchange a lot of managers and employees, or firms that are located close together, would tend to have more similar cultures and act in more similar ways.

In some of these cases, it may be difficult to separate learning effects from sorting effects, but that is not such a problem if we consider it a joint test of the learning and sorting theories of culture (Van den Steen 2002a).

A final, and potentially very promising, approach are experimental analyses in the style of Weber and Camerer (2001). A possible experiment would be to have groups of people repeatedly play a partially random game and see whether groups develop internally homogenous beliefs about the optimal actions and whether such beliefs can persist over generations of players.

## 10 Conclusion

This paper shows how corporate culture, defined as shared beliefs, can evolve from the common experience of a firm's members. It partially validated the model by showing that it is consistent with some key stylized facts. The culture of the firm is influenced by the original manager's beliefs and can persist even long after that manager is gone. Outside successors are more likely to change the culture than insiders. Firms in similar circumstances can develop different beliefs.

A key insight of the paper is that a correlation between culture and performance is to be expected even if culture has no performance benefits. From a practical point of view, the paper thus raises the question whether the perceived benefits of a strong culture might actually be due to an inverse causality.

Implicitly the paper also suggests that firms have a natural tendency to develop a culture in the sense of shared beliefs. Management can try to influence this process, but only to some extent.

I believe that the results in this paper suggest that taking the sociological model of culture as shared beliefs seriously might be a promising approach for economics to develop a systematic theory of corporate culture that contributes to our understanding of the behavior and performance of organizations. It might also provide a useful perspective on agency problems and incentive systems.

## A A Note on ‘Differing Beliefs’ in Economic Modeling

The model in this paper differs in one respect from most economic models: the agents knowingly entertain differing beliefs (without having private information). The reason for this assumption is pragmatic: differences in beliefs are at the heart of the issues studied here, and assuming common knowledge of differing beliefs is the most transparent and parsimonious way to study this question.<sup>14</sup> Differing beliefs do not contradict the economic paradigm: while rational agents should use Bayes’ rule to update their prior with new information, nothing is said about those priors themselves, which are primitives of the model. In particular, absent any relevant information agents have no rational basis to agree on a prior.<sup>15</sup> Harsanyi (1967), for example, observed that ‘by the very nature of subjective probabilities, even if two individuals have exactly the same information and are at exactly the same high level of intelligence, they may very well assign different subjective probabilities to the very same events’. The best argument for the traditional use of common priors is Aumann’s (1987) argument that they allow us to ‘zero in on purely informational issues’. Conversely, differing priors allow us to zero in on the implications of open disagreement and differing beliefs. Note, finally, that the existence of the winner’s curse, any disagreement on the correct action when utilities are aligned, and nearly all evidence on bounded rationality and biases in decision making imply empirically the existence of differing priors.

Van den Steen (2002c) considers this issue in more detail.<sup>16</sup> Among other things, it argues against the idea that differing priors might allow us to explain anything, it discounts the theoretical possibility that agents will make very large or infinite bets, and shows that the epistemic foundations for Nash equilibria in the sense of Aumann and Brandenburger (1995) extend to this context with differing priors on the payoff-space.

**The model with common priors** Nearly identical results obtain under a common prior. In particular, assume that:

- Agents have a common prior.
- All agents get a private signal about each action.
- All employees of a firm observe the manager’s signal.
- Agents cannot in any way observe each other’s beliefs.

Note that this model, makes the agents behave as if they have differing priors. While it sticks pro forma to the tradition of a common prior, it does so artificially.

With multiple decision makers (or if we do not allow the non-CEO agents to observe the CEO’s signal), this model increases enormously in complexity. Agents now have to back out signals from actions. While this becomes quickly completely intransparent, it seems that all results should still hold. In this case, keeping with the tradition comes at a real cost either in terms of an important increase in analytical complexity or in the inability to analytically obtain some results that do, however, hold.

The theoretical support for differing priors mentioned at the start of this section combined with these very practical reasons have led me to use a model with differing priors.

---

<sup>14</sup>While formally most of the analysis can be done under the common prior assumption, such analysis would miss the essential point: that, *holding information constant*, the strength of beliefs is an important influence; that it can be optimal to have a CEO who has stronger beliefs than the board *even if he does not have more information*.

<sup>15</sup>A traditional argument against differing priors is that rational people would argue until they reach agreement. There are 3 reasons why this argument does not hold. First of all, any fact that is presented in the discussion brings with it potentially differing priors about the relevance of the fact or the precision of the data. This leads immediately to an infinite regress of arguments. Second, what matters for economic analysis is not so much whether agents can theoretically come to agreement but whether they do at the time they have to take an action. Finally, as argued in the main paper, much information is difficult, if not impossible, to communicate. Note also that under the common prior assumption, agents with aligned utility functions should never disagree on the optimal actions to undertake. On a sinking ship or in the heat of a battle (where survival is supposedly the common objective), there should always be perfect agreement on the optimal course of action. This seems to contradict our intuition.

<sup>16</sup>See also Morris (1995), the discussion between Gul (1998) and Aumann (1998), or Yildiz (2000).

**Utility versus beliefs** An alternative way of modelling would be to incorporate beliefs in the utility function. Rotemberg and Saloner (2000) and Hart and Holmström (2002) can be interpreted in this way.

While this approach may seem more general, there are a number of issues. First of all, there are real differences between utility and beliefs. Utilities can be influenced by incentive contracts and beliefs by new data, but not the other way around. One implication is that beliefs provide a stronger commitment: financial incentives based on firm performance can undo biased preferences, but will not change beliefs about what the firm has to do to maximize that performance. Second, the benefits of stock-based compensation also differ. The use of stock options as a cheap way to compensate ‘believers’ does not work when there are only utility differences. A third implication relates to the dynamics of the situation. As more information becomes available, beliefs will change but utilities won’t. The two models thus have different implications for change and resistance to change. Overall, modelling beliefs via utilities risks missing out on important phenomena.

Apart from these analytical differences, there are obvious differences in interpretation. Since the phenomena in this paper are more naturally described in terms of beliefs, it makes sense to model beliefs explicitly.

## B Proofs

### B.1 Basic analysis

Continuation of the proof of proposition 1.

**Proof:** Prior to starting the proof, we need to define the probability spaces as they will be used in this and other proofs. As mentioned in the main text, let  $\epsilon_n(s)$  denote the variation  $\epsilon^t$  when the firm samples action  $n$  for the  $s$ 'th time. Note that we could express the manager's prior also as a variation  $\epsilon_n(0) = r_n^o - \rho_{a_n}$ , and then consider an 'extended' sequence of variations that includes the manager's prior.

Let  $(\Omega_{n,\epsilon}, S_{n,\epsilon}, P_{n,\epsilon})$  denote the probability space generated by the variations  $\epsilon_n(s)$  for action  $a_n$  only. Let  $(\Omega_\epsilon, S_\epsilon, P_\epsilon)$  denote the product space of the  $(\Omega_{n,\epsilon}, S_{n,\epsilon}, P_{n,\epsilon})$ . Let  $(\Omega, S, P)$  denote the probability space generated by the variations  $\epsilon_n(s)$ , the prior beliefs  $\mathbf{r}^0$  of all agents (in firm  $f$ ), and the average performances of all actions  $\rho$ . Let  $\omega_{n,\epsilon}$ ,  $\omega_\epsilon$ , and  $\omega$  denote typical elements of these probability spaces. Let, finally,  $(\Omega_{n,\bar{\epsilon}}, S_{n,\bar{\epsilon}}, P_{n,\bar{\epsilon}})$  denote the probability space generated by the variations  $\epsilon_n(s)$  and the manager's prior (i.e. the 'extended' sequence of variations mentioned above), for action  $a_n$  only, and let  $(\Omega_{\bar{\epsilon}}, S_{\bar{\epsilon}}, P_{\bar{\epsilon}})$  denote the product space of these  $(\Omega_{n,\bar{\epsilon}}, S_{n,\bar{\epsilon}}, P_{n,\bar{\epsilon}})$ .

I now start with *the second part of the proposition*. Let  $(\mathbf{r}, \sigma^2) = (r_1, \dots, r_n, \sigma_1^2, \dots, \sigma_n^2)$  denote the manager's current beliefs about the  $N$  actions (including the constant one). Theorem 2.2 of Gittins (1989) implies that an optimal policy exists, that the corresponding value function satisfies the functional equation of dynamic programming, that a policy is optimal if and only if it maximizes the right hand side of that functional equation of dynamic programming, and that the value iteration of the functional equation starting from a bounded approximation of the value function converges to the true value function.

Let  $u_n((\mathbf{r}, \sigma^2), x)$  be the updated values of the manager's beliefs given a prior (for the period considered)  $(\mathbf{r}, \sigma^2)$ , action  $n$  being chosen, and the return of that action being  $x$ .

The functional equation of dynamic programming can be written

$$V = \max_{n \in N} V_n$$

and

$$V_n = r_n + \delta \int V(u_n((\mathbf{r}, \sigma^2), v)) N_{r_n, \sigma_n^2}(v) dv$$

with  $N_{\mu, \sigma^2}$  denoting the normal density with mean  $\mu$  and variance  $\sigma^2$ .

For the value iteration, define

$$\begin{aligned} V^0(\mathbf{r}, \sigma^2) &= 0 \\ V^t(\mathbf{r}, \sigma^2) &= \max_{n \in N} V_n^t \\ V_n^t(\mathbf{r}, \sigma^2) &= r_n + \delta \int V^{t-1}(u_n((\mathbf{r}, \sigma^2), v)) N_{r_n, \sigma_n^2}(v) dv \end{aligned}$$

Theorem 2.2 in Gittins (1989) thus implies that  $\lim_{t \rightarrow \infty} V^t(\mathbf{r}, \sigma^2) = V(\mathbf{r}, \sigma^2)$ .

**Lemma 1** *The functions  $V^t$  and  $V_n^t$  are continuous on  $\mathbb{R} \times [0, \infty)$ .*

**Proof:** This follows by induction. The proposition holds for  $t = 1$  since a maximum of continuous functions is continuous. Assume now that  $V^{t-1}$  and  $V_n^{t-1}$  are continuous. Then  $V_n^t$  is continuous since  $u_n$  and  $V_n^{t-1}$  are.  $V^t$  is again continuous since it is a maximum of continuous functions. ■

**Lemma 2**  *$V^t \geq V^{t-1}$  and  $V_n^t \geq V_n^{t-1}$*

**Proof :** Let  $\hat{V}^t$  equal  $V^t$  if we restrict the choice set in the  $t$ 'th period to the constant (return zero) action. Clearly  $V^{t-1} = \hat{V}^t \leq V^t$ . Analogously for  $V_n^t$  and  $V_n^{t-1}$ . ■

**Lemma 3** *The functions  $V(\mathbf{r}, \sigma^2)$  and  $V_n(\mathbf{r}, \sigma^2)$  are continuous on  $\mathbb{R} \times [0, \infty)$ .*

**Proof :** Let  $\hat{r}(\mathbf{r}, \sigma^2)$  be the expected value of the maximum return given the distributions of returns characterized by  $(\mathbf{r}, \sigma^2)$ , i.e.  $\hat{r}(\mathbf{r}, \sigma^2) = E[\max_{n \in N} r_n \mid (\mathbf{r}, \sigma^2)]$ . Note that for some fixed  $\epsilon > 0$ ,  $\max(\hat{r}(\mathbf{r}, \sigma^2), \epsilon) > 0$ . Moreover, since  $\hat{r}(\mathbf{r}, \sigma^2)$  is a continuous function on  $\mathbb{R} \times [0, \infty)$ , so is  $\max(\hat{r}(\mathbf{r}, \sigma^2), \epsilon)$ . Let  $\bar{V}^t(\mathbf{r}, \sigma^2)$  be the expected discounted payoff from the first  $t$  periods of the original infinitely repeated game.

In that case

$$\begin{aligned} V(\mathbf{r}, \sigma^2) &\leq \bar{V}^t(\mathbf{r}, \sigma^2) + \delta^{t+1} \sum_{\tau=0}^{\infty} \delta^\tau \hat{r}(\mathbf{r}, \sigma^2) \\ &\leq V^t(\mathbf{r}, \sigma^2) + \delta^{t+1} \frac{1}{1-\delta} \hat{r}(\mathbf{r}, \sigma^2) \\ &\leq V^t(\mathbf{r}, \sigma^2) + \delta^{t+1} \frac{1}{1-\delta} \max(\hat{r}(\mathbf{r}, \sigma^2), \epsilon) \end{aligned}$$

or

$$\frac{V(\mathbf{r}, \sigma^2)}{\max(\hat{r}(\mathbf{r}, \sigma^2), \epsilon)} \leq \frac{V^t(\mathbf{r}, \sigma^2)}{\max(\hat{r}(\mathbf{r}, \sigma^2), \epsilon)} + \delta^{t+1} \frac{1}{1-\delta}$$

or

$$\left| \frac{V(\mathbf{r}, \sigma^2)}{\max(\hat{r}(\mathbf{r}, \sigma^2), \epsilon)} - \frac{V^t(\mathbf{r}, \sigma^2)}{\max(\hat{r}(\mathbf{r}, \sigma^2), \epsilon)} \right| \leq \delta^{t+1} \frac{1}{1-\delta}$$

Since this inequality is independent from  $(\mathbf{r}, \sigma^2)$ ,  $\frac{V^t(\mathbf{r}, \sigma^2)}{\max(\hat{r}(\mathbf{r}, \sigma^2), \epsilon)}$  converges uniformly to  $\frac{V(\mathbf{r}, \sigma^2)}{\max(\hat{r}(\mathbf{r}, \sigma^2), \epsilon)}$ . But then, together with the fact that  $\frac{V^t}{\max(\hat{r}(\mathbf{r}, \sigma^2), \epsilon)}$  is continuous, this implies that  $\frac{V(\mathbf{r}, \sigma^2)}{\max(\hat{r}(\mathbf{r}, \sigma^2), \epsilon)}$ , and thus  $V(\mathbf{r}, \sigma^2)$ , is continuous. This finally implies that  $V_n$  is continuous too. ■

Let now

$$A_{n,l} = \{(\mathbf{r}, \sigma^2) \mid V_n(\mathbf{r}, \sigma^2) > V_l(\mathbf{r}, \sigma^2)\}$$

The set  $A_{n,l}$  is the set of values for which it is strictly better to play action  $n$  than action  $l$ .

**Lemma 4** *The set  $A_{n,l}$  is open.*

**Proof :** The proof is identical to lemma 4 in Rothschild (1974). ■

Let

$$A_n = \{(\mathbf{r}, \sigma^2) \mid \forall l \neq n V_n(\mathbf{r}, \sigma^2) > V_l(\mathbf{r}, \sigma^2)\}$$

so that  $A_n = \bigcap_{l \neq n} A_{n,l}$  and thus  $A_n$  is also open.

I now prove *the second part of the proposition*. In particular, I will show that for any two arbitrary actions, almost surely at most one action will be played infinitely often under any optimal strategy. This implies the second part of the proposition.<sup>17</sup>

<sup>17</sup>Since it implies that one and only one of the actions will be played infinitely often. Note that at least one action will be played infinitely often since there are only a finite number of actions but an infinite number of periods.

Note that, given the assumptions on the distribution of the  $\rho_a$  and the outcomes, it is almost surely true that  $\rho_n \neq \rho_l$  when  $n \neq l$ . Let again  $A_{n,l} = \{(\mathbf{r}, \sigma^2) \mid V_n(\mathbf{r}, \sigma^2) > V_l(\mathbf{r}, \sigma^2)\}$ . Let  $\hat{\rho}_n$  and  $\hat{\rho}_l$  be the specific values (of  $\rho$ ) for actions  $a_n$  and  $a_l$  (i.e. they were drawn from a normal  $N(0, 1)$  distribution). Suppose (without loss of generality) that  $\hat{\rho}_n > \hat{\rho}_l$ . Then  $(\mathbf{r}, \sigma^2)$  with  $\hat{\rho}_n$  and  $\hat{\rho}_l$  as given and  $\sigma_n^2 = \sigma_l^2 = 0$  is clearly in  $A_{n,l}$ . Since  $A_{n,l}$  is open, there exist  $\epsilon > 0$  and  $\delta > 0$  such that  $(\mathbf{r}, \sigma^2)$  with  $r_n \in (\hat{\rho}_n - \epsilon, \hat{\rho}_n + \epsilon)$ ,  $r_l \in (\hat{\rho}_l - \epsilon, \hat{\rho}_l + \epsilon)$ ,  $\sigma_n^2 < \delta$  and  $\sigma_l^2 < \delta$ , is also in  $A_{n,l}$ .

Let now

$$Q_n = \{\omega_{n,\bar{\epsilon}} \in \Omega_{n,\bar{\epsilon}} \mid \lim_{s \rightarrow \infty} r_n(\omega_{n,\bar{\epsilon}}(s)) \in (\hat{\rho}_n - \epsilon, \hat{\rho}_n + \epsilon)\}$$

By the law of large numbers,  $P_n(Q_n \mid \rho_n = \hat{\rho}_n) = 1$ .

Now I claim that for any  $\hat{\omega}_{\bar{\epsilon}}$  with  $\hat{\omega}_{n,\bar{\epsilon}} \in Q_n$  and  $\hat{\omega}_{l,\bar{\epsilon}} \in Q_l$ , it cannot be true that under the optimal policy both action  $n$  and action  $l$  get played infinitely often. If they get both played infinitely often, then, for any  $\epsilon > 0$  and  $\delta > 0$ , from some period on, it will be the case that  $r_n \in (\hat{\rho}_n - \epsilon, \hat{\rho}_n + \epsilon)$ ,  $r_l \in (\hat{\rho}_l - \epsilon, \hat{\rho}_l + \epsilon)$ ,  $\sigma_n^2 < \delta$  and  $\sigma_l^2 < \delta$ . With the  $\epsilon$  and  $\delta$  from above, it follows that from that period on  $(\mathbf{r}, \sigma^2) \in A_{n,l}$ . But this implies that  $l$  will not get played any more after period  $t$ , which contradicts the assumption that both get played infinitely often. Since  $P(\hat{\Omega}) = 1$  for  $\hat{\Omega} = \{\omega_{\bar{\epsilon}} \in \Omega_{\bar{\epsilon}} \mid \omega_{n,\bar{\epsilon}} \in Q_n \text{ and } \omega_{l,\bar{\epsilon}} \in Q_l\}$ , this implies the proposition.

For the *third part of the proposition*, let  $\nu(x, m)$  denote the index number of an action that has been tried  $m-1$  times and which (after these  $m-1$  tries) has an expected payoff  $x$ . By Gittins (1989),  $\nu(x, m) = x + \nu(0, m)$ . Let now  $\underline{\nu} = -\nu(0, 1)$  and let  $\rho_k > 0$  for some  $1 < k \leq N$ . Note that with strictly positive probability  $r_n^o < \underline{\nu}$  for all actions with  $n \neq 1$ ,  $n \neq k$  and  $r_k^o \geq \rho_k$ . Since the starting index numbers of all the actions with  $n \neq 1$ ,  $n \neq k$  are now smaller than zero, they will never be selected. For  $k$  to be the eventual choice of action it suffices that with positive probability,  $r_{a_k}^t = \frac{r_k^o + \sum_{s=1}^t x_s}{t+1} > 0$  for all  $t$ . But this holds since this is a random walk with positive starting value and positive drift. For the opposite conclusion, let  $\rho_k < 0$  and assume that the firm eventually takes  $a_k$  as its course of action. By the law of large numbers, the mean of the posterior belief of the manager almost surely converges to  $\rho_k$ . Moreover, the variance of the posterior belief converges to zero. It follows that in the limit, the index number of  $a_k$  converges almost surely to  $\rho_k < 0$ . But once the index number is strictly smaller than zero, the manager prefers to follow the constant zero-return action. This contradicts the assumption that  $a_k$  was the firm's eventual course of action.

For the *fourth and last part of the proof*, note the well-known fact that a bandit problem with a constant action is a stopping problem: once the constant action is selected once, it will be selected forever. It is thus sufficient to show that with positive probability the constant action will be played at least once although one or more other actions have a higher reference return. For simplicity I present the argument when there is only one non-constant action. The index number of the constant action is zero (since a constant action with zero payoff is the calibration point for the Gittins index). Using the results in Gittins (1989), the index number of the other action is

$$\nu(0, n+1) + \bar{x}$$

where  $n$  is the number of times the action has been selected and  $\bar{x}$  is the average observation (including the prior). Since the returns upon selection of an action are normally distributed, the return upon selection of the non-constant action will with positive probability be so negative that its index number turns negative, no matter what the average return of the action is. It follows that the manager will select the constant action the next period (and forever after), even when its constant return is lower than the reference return of the variable action.

For the second part of part 4, it suffices to show that the probability of any particular action (including the action with the highest payoff) being the eventual action goes to zero as  $N \rightarrow \infty$ . To show this, assume (by contradiction) that the probability of some  $a_n$  with average payoff  $\rho_n$  becoming culture is always larger or equal than  $\epsilon > 0$ . Then, by lemma 6 below, that must be the case for all actions with  $\rho \geq \rho_n$ . As  $N \rightarrow \infty$ , the number of such actions is almost surely larger than  $1/\epsilon$ . This implies that the probability of any such action becoming the culture is larger than one, which is clearly a contradiction. ■

**Lemma 5** For a given  $\hat{\omega}$ , prior beliefs  $\mathbf{r}^0$ , and returns  $\rho$  on all but one action  $n$ , the number of times the action  $n$  is chosen prior to some time  $\hat{t}$  increases in the average return  $\rho_n$  of the action, and vice versa.

For a given  $\hat{\omega}$ , returns  $\rho$ , and prior beliefs  $\mathbf{r}^0$  on all but one action  $n$ , the number of times the action  $n$  is chosen prior to some time  $\hat{t}$  increases in the manager's prior  $r_n^0$  regarding the action, and vice versa.

**Proof :** Note that the  $\epsilon_n(s)$  (that correspond to this  $\hat{\omega}$ ) can be combined with the priors to specify the outcomes directly in terms of the resulting index numbers. The manager's choice of action is simply that action which has the highest current index number. An increase in  $\rho_n$  (without changing the prior  $r_n^0$ ) raises all index numbers for action  $n$  and more so as they are further away in time.

Consider two values for  $\rho_n$ :  $\check{\rho}_n < \tilde{\rho}_n$ .

We first show that this implies that prior to time  $\hat{t}$  the number of times action  $n$  is chosen increases. The key is to show (by induction) that a certain index number for  $n$ , say the  $s$ 'th, cannot be selected at a later time when  $\rho_n = \tilde{\rho}_n$  than when it was  $\check{\rho}_n$ . This is trivial for the first index number since the index number is unchanged from before. Assume now that the  $(s-1)$ 'th number was selected at the same or an earlier time than under  $\check{\rho}_n$ . We have to prove that the  $s$ 'th will not be selected any later than before.

Suppose, by contradiction, it were selected at a later time, then at the time that it was selected under  $\check{\rho}_n$ , all index numbers are identical, except for that of action  $n$  which is higher. Under  $\check{\rho}_n$  the index of  $n$  is lower than it is now, and nevertheless  $n$  was chosen last time and thus must have been the highest index number and not now. This is a contradiction.

An analogous argument shows that the number of times the action is selected decreases as  $\rho_n$  decreases.

This also implies that when the action has been selected more often prior to some time  $\hat{t}$  and all is identical except  $\rho_n$ , then it must be that  $\rho_n$  is higher.

The argument for the second part of the proposition is completely analogous. ■

**Lemma 6** The probability that an action  $\hat{a}$  becomes a firm's eventual action  $a^*$  increases in  $\rho_{\hat{a}}$  and in the manager's prior  $r_{\hat{a}}^0$ .

**Proof :** Let  $I^*$  be an indicator function that  $\hat{a}$  is the eventual action. Lemma 5 implies that, for a given  $\hat{\omega}$ , prior beliefs, and returns on all but  $\hat{a}$ ,  $I^*$  increases in  $\rho_{\hat{a}}$ . The result then follows by taking expectations. The argument for the prior is identical. ■

## B.2 Managers and culture

**Lemma 7** For any actions  $\tilde{a} \neq \hat{a}$ ,  $E[\rho_{\tilde{a}} - \rho_{\hat{a}} \mid \tilde{a} = a^*] \geq 0$  and  $P[\rho_{\tilde{a}} \geq \rho_{\hat{a}} \mid \tilde{a} = a^*] \geq 0.5$ .

**Proof :** By lemma 8 and the properties of FOSD, it suffices to show that the distribution of  $\rho_{\tilde{a}}$  FOSD that of  $\rho_{\hat{a}}$ .

Note now that, with  $g$  denoting the prior distribution of  $\rho_a$

$$f(\rho_{\tilde{a}} \mid \tilde{a} = a^*) = \frac{P[\tilde{a} = a^* \mid \rho_{\tilde{a}} = v]g(v)}{\int P[\tilde{a} = a^* \mid \rho_{\tilde{a}} = u]g(u) du}$$

while

$$f(\rho_{\hat{a}} \mid \hat{a} = a^*) = \frac{P[\hat{a} = a^* \mid \rho_{\hat{a}} = v]g(v)}{\int P[\hat{a} = a^* \mid \rho_{\hat{a}} = u]g(u) du}$$

Let

$$\frac{P[\tilde{a} = a^* \mid \rho_{\tilde{a}} = v]}{\int P[\tilde{a} = a^* \mid \rho_{\tilde{a}} = u]g(u) du} = h_1(v)$$



and

$$\frac{P[\tilde{a} = a^* \mid \rho_{\tilde{a}} = v]}{\int P[\tilde{a} = a^* \mid \rho_a = u]g(u) du} = h_2(v)$$

By lemma 6,  $h_1$  increases in  $v$  (from 0 to some number). Analogously,  $h_2$  decreases in  $v$  (from some number to 0). Let  $\tilde{v}$  be defined by  $h_1(\tilde{v}) = h_2(\tilde{v})$ . Such  $\tilde{v}$  is unique and  $h_1 \leq h_2$  iff  $v \leq \tilde{v}$ .

So you need to show now that  $h_1(v)g(v)$  FOSD  $h_2(v)g(v)$  or

$$\int_{-\infty}^x [h_1(v) - h_2(v)]g(v) dv \leq 0$$

Clearly, the above holds for  $x < \tilde{v}$ . Moreover, for  $x \geq \tilde{v}$ ,

$$\begin{aligned} & \int_{-\infty}^x [h_1(v) - h_2(v)]g(v) dv \\ &= - \int_x^{\infty} [h_1(v) - h_2(v)]g(v) dv \\ &\leq 0 \end{aligned}$$

This proves the result. ■

**Lemma 8** *If  $f(x)$  first order stochastically dominates  $g(y)$  (so that  $F(x) \leq G(x)$ ), then  $P[x \geq y] \geq 0.5$ .*

**Proof :** Note that

$$\begin{aligned} P[x \geq y] &= \int_{-\infty}^{\infty} \int_{-\infty}^x g(y) dy f(x) dx \\ &= \int_{-\infty}^{\infty} G(x) f(x) dx \\ &\geq \int_{-\infty}^{\infty} F(x) f(x) dx \\ &= \int_{-\infty}^{\infty} F(x) dF(x) \\ &= \int_0^1 u du = 0.5 \end{aligned}$$
■

## References

- AGHION, P., AND J. TIROLE (1997): "Formal and real authority in organizations," *Journal of Political Economy*, 105(1), 1– 29.
- AUMANN, R. J. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1– 18.
- AUMANN, R. J. (1998): "Common Priors : A Reply to Gul," *Econometrica*, 66(4), 929– 938.
- AUMANN, R. J., AND A. BRANDENBURGER (1995): "Epistemic Conditions for Nash Equilibrium," *Econometrica*, 63, 1161– 1180.
- BAKER, E. L. (1980): "Managing Organizational Culture," *Management Review*, 69(7), 8– 13.
- BARON, J. N., M. D. BURTON, AND M. T. HANNAN (1999): "Engineering Bureaucracy: The Genesis of Formal Policies, Positions, and Structures in High-Technology Firms," *Journal of Law, Economics, and Organization*, 15(1), 1– 41.
- BERRY, D. E., AND B. FRISTEDT (1985): *Bandit Problems: Sequential Allocation of Experiments*. Chapman and Hall, New York.
- BERTRAND, M., AND A. SCHOAR (2001): "Managing with Style: The Effect of Managers on Firm Policies," Working paper.
- BURNS, T., AND G. M. STALKER (1961): *The Management of Innovation*. London : Tavistock.
- BURT, R. S., S. M. GABBAY, AND G. HOLT (1994): "Contingent Organization as a Network Theory: The Culture Performance Contingency Function," *Acta Sociologica*, 37, 345– 370.
- CARRILLO, J. D., AND D. GROMB (1999): "On the Strength of Corporate Cultures," *European Economic Review*, 43, 1021– 1037.
- COLLINS, J. C., AND J. I. PORRAS (1994): *Built to last : Successful Habits of Visionary Companies*. Harper & Row, New York.
- CRAWFORD, V., AND J. SOBEL (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431– 1452.
- CRÉMER, J. (1993): "Corporate Culture and Shared Knowledge," *Industrial and Corporate Change*, 3(3), 351– 386.
- DEAL, T., AND A. A. KENNEDY (1982): *Corporate Cultures: The Rites and Rituals of Corporate Life*. Addison Wesley, Reading MA.
- DESSEIN, W. (2001): "Authority and Communication in Organizations," Forthcoming in the Review of Economics Studies.
- DONALDSON, G., AND J. W. LORSCH (1983): *Decision Making at the Top : The Shaping of Strategic Direction*. Basic Books, New York.
- GITTINS, J. C. (1989): *Multi-Armed Bandit Allocation Indices*. John Wiley & Sons, New York.
- GITTINS, J. C., AND D. M. JONES (1974): "A Dynamic Allocation Index for the Sequential Design of Experiments," in *Progress in Statistics*, ed. by J. Gani, et al., and et. al. North-Holland, Amsterdam.
- GREIF, A. (1994): "Cultural beliefs and the organization of society : A historical and theoretical reflection on collectivist and individualist societies," *Journal of Political Economy*, 102(5).
- GUL, F. (1998): "A Comment on Aumann's Bayesian View," *Econometrica*, 66(4), 923– 927.
- HARSANYI, J. C. (1967): "Games with Incomplete Information Played by 'Bayesian' Players, Parts I, II and III," *Management Science*, 14, 159– 182.
- HART, O., AND B. HOLMSTRÖM (2002): "Vision and Firm Scope," Working Paper.
- HERMALIN, B. E. (2001): "Economics and Corporate Culture," in *The International Handbook of Organizational Culture and Climate*, ed. by C. L. Cooper, S. Cartwright, and P. C. Earley. John Wiley & Sons, New York, Working paper, University of California-Berkeley.
- KINCAID, P. (1986): *The Rule of the Road: An International Guide to History and Practice*. Greenwood Press, Westport, CT.
- KOTTER, J. P., AND J. L. HESKETT (1992): *Corporate Culture and Performance*. Free Press, New York.
- KREPS, D. M. (1990): "Corporate Culture and Economic Theory," pp. 90– 143.

- LAZEAR, E. P. (1995): "Corporate Culture and Diffusion of Values," in *Trends in Business Organization : Do Participation and Cooperation Increase Competitiveness*, ed. by H. Siebert, pp. 89– 133. J.C.B. Mohr/Paul Siebeck, Tübingen.
- MARTIN, J. (1992): *Cultures in Organizations: Three Perspectives*. Oxford University Press, New York.
- MCGAHAN, A. M. (1999): "The Performance of US Corporations: 1981-1994," *The Journal of Industrial Economics*, 47(4), 373– 398.
- MORRIS, S. (1995): "The Common Prior Assumption in Economic Theory," *Economics and Philosophy*, 11, 227– 253.
- MUELLER, D. C. (1990): "The Persistence of Profits in the United States," in *The Dynamics of Company Profits: An International Comparison*, ed. by D. C. Mueller. Cambridge University Press, Cambridge.
- PETERS, T. J., AND R. H. WATERMAN (1982): *In Search of Excellence : Lessons from America's Best-run Companies*. Harper & Row, New York.
- ROB, R., AND P. B. ZEMSKY (2002): "Social Capital, Corporate Culture and Incentive Intensity," *Rand Journal of Economics*, 33(2), 243– 257.
- ROTEMBERG, J. J., AND G. SALONER (2000): "Visionaries, Managers, and Strategic Direction," *Rand Journal of Economics*, 31(4), 693– 716.
- ROTHSCHILD, M. (1974): "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 9, 185– 202.
- SCHEIN, E. H. (1985): *Organizational Culture and Leadership*. Jossey-Bass Publishers, San Francisco, CA.
- SCHWARTZ, H., AND S. DAVIS (1981): "Matching Corporate Culture and Business Strategy," *Organizational Dynamics*, pp. 30– 48.
- SØRENSEN, J. B. (2002): "The Strength of Corporate Culture and the Reliability of Firm Performance," *Administrative Science Quarterly*, Forthcoming.
- VAN DEN STEEN, E. J. (2002a): "Birds of a Feather ... : Formation and Effects of Organizational Culture," Working paper, MIT-Sloan.
- (2002b): "Costs and Benefits of Homogeneity," Working paper.
- (2002c): "Notes on Modelling with Differing or Heterogeneous Priors," Working Paper, MIT-Sloan.
- VAN DEN STEEN, E. J. (2002d): "Organizational Beliefs and Managerial Vision," Working paper, MIT-Sloan.
- WEBER, R. A., AND C. F. CAMERER (2001): "Cultural Conflict and Merger Failure: An Experimental Approach," USC Center for Law, Economics, and Organization Research Paper C01-10.
- YILDIZ, M. (2000): "Sequential Bargaining without a Recognition Device," working paper.