Organizational codes in the lab

Colin F. Camerer

Division of HSS 228-77
Caltech
Pasadena CA 91125
camerer@hss.caltech.edu

Roberto Weber

Social & Decision Sciences
Carnegie-Mellon University

25 November 2003. **This paper is both incomplete [esp Appendices and section 4] and too long**. It was prepared for advance circulation at the NBER Org Econ conference, December 12-13, 2003. **Please do not circulate** but comments are welcome. This research was supported by NSF grant SES-???. Comments were received from audiences at the Yale Law School LEO workshop (September 2003, the USC LEO workshop (June 2002), and the Stanford SITE conference (August 2003). The project benefitted from terrific research assistance from Ming Hsu, Galen Loram, Alice Oh, and Scott Rick, and "extreme" computer programming (as well as insight) from Charlie Hornberger and John Lin.

1. Introduction

Organizational culture is usually defined as "the way business is done around here", or a core set of values which are expressed by stories, rituals, persona, symbols, and language. This paper studies the linguistic dimension of culture— *codes*— experimentally. A code is a set of terms that a group creates endogeneously to concisely label complex situations. An advantage of studying codes, as a metaphor for culture, is that they are easy to measure and can be created rapidly, in the short time frame available in the lab, compared to corporate values.

In the experiments, individual members of a laboratory group (or "firm") view a set of identical pictures . Think of the pictures as scenarios representing possible workplace events. Figure 1 shows an example (an office scene).

Some members of the group ("managers") know which pictures are "targets" and have to describe them to others, rapidly and accurately, so that these uninformed group members ("employees"), who are being told which pictures to select, choose the right pictures. To do so, they must together develop a code that assigns distinctive words or phrases to pictures, so that every member of the group can quickly understand which picture is being referred to. For example, groups called the office scene in Figure 1 "Lady with typewriter" (at the foreground), "Cubeville", or "Uday Rao" (since the man standing on the right looked like their professor, Uday Rao).

The experimental codes are like the specialized language workers in an organization use to describe scenes that are not mutually viewable to each other, in order to facilitate action (like a police dispatcher using "Code Six"—officer down – and other terms to rapidly describe key features of a crime scene to a police officer who is speeding toward the scene).

Obviously the code development in the lab is a sped-up version of the way code naturally evolves (like a time-lapse photograph that compresses the process of a rose blooming into 3 minutes). During a boom in interest in organizational culture, Business Week (1986) reported the following story: " 'This corporate culture stuff is great' the chairman raved at dinner following the talk. Then, turning to his president, he demanded,

`I want a culture by Monday!'" The story is amusing because the impatient chairman doesn't seem to understand that cultures (usually) grow slowly.

However, our experiments *are* like creating a culture instantly, `by Monday'. The point of creating such instant codes is not to replicate the slow-moving process by which actual cultures are created in naturally-occurring environments on the same time-scale. Instead, the idea is more like using fruit flies to study genetics because their rapid breeding time enables the production of many generations quickly. For studying culture, it is useful to gather many observations rapidly (for statistical purposes) and to be able to impose exogeneous treatments and study their effects.

Figure 1: A picture seen in the experiment (coded as "Lady with typewriter", "Cubeville", "Uday Rao")

A. Culture and codes

It is important to understand how codes relate to culture. Language codes are a *facet* of culture. Schein (1983) defines culture as "the *basic assumptions* and *beliefs* that are shared by members of an organization."[1] Codes are one expression of those values. In universities, "dead wood" is a derisive term to describe tenured professors who are no longer productive (typically on any dimension). The term expresses the organizational value that faculty should repay their tenure "gift" with some kind of service or output; unproductive tenured faculty weigh the organizational "tree" down and take rather than give. At Disney's theme parks, all workers are called "cast members". This code emphasizes an organizational value— workers are in a service business where they provide live entertainment or entertainment support.

One way to articulate values is in the form of a giant employee manual. The manual could tell employees what to do in every situation. The problem is that in complex, fluid organizations there are a huge number of situations. Using a manual only works if employees can read, remember, and understand it. But the human brain evolved in an ancestral past in which people hunted and gathered in small bands. In the face of evolved brain structure, the ideal expression of corporate values is not likely to be a written rulebook. Instead, effective cultures take advantage of the way in which brains process and remember information (which is rarely in the form of memorizing manuals or rulebooks). Therefore, culture is usually expressed by *stories*, *personification* (exemplars of "heroes" and "villains" who reflect the values and betray them), *symbols, rituals*, and *code*.

Stories express culture with cognitive economy because human memory dually codes memories into "episodes" (time-stamped single events) and "semantic memory" (the projection of the set of recalled episodes into categorical facts about the world) (see Martin, 19??). *Personification* of culture is cognitively economical because various brain structures make it easy to store information about what to do in the form of traits of people. (Some examples of personification are given below when we discuss founder effects.) *Symbols* are icons that are stored in visual representations and linked to episodes and people. (For example, many Nike employees, known as "Ekins"—the word Nike

[1] To an economist, organizational "values" are agreed-upon tradeoffs among organizational goals, like profit and customer satisfaction. A contractual way to think about values is as a system of rules that tell employees what to do (i.e., what behavior is appropriate or `valued') in different situations.

spelled backwards-- have the famous Nike "swoosh" logo tattooed on their ankles.) *Rituals* create episodic memories and are often personalized. (Think of Silicon Valley "beer busts" in which the (paper-)billionaire CEO shows up every Friday afternoon in shirtsleeves to drink beer with loading-dock workers and secretaries. This ritual expresses, in a memorable way, an egalitarian culture in which every employee's work is valued.) *Code* is the specialized language used to describe the situations employees face.

Code is one facet or expression of culture. For instance, the fact that the Navajo language contains no words for superior, subordinate, boss, or hierarchy reflects a deeply held cultural value of equality (Ott, 1989). When Time Warner employees call AOL employees "Assholes Online" (Klein, 2003) their derisive codeword expresses a contempt for what they perceive as an aggressive, money-grubbing AOL corporate value they do not share. When casino employees call a heavy gambler a "whale", their code illustrates the casino's values— a customer is like a fish, and the goal is to catch big fish.

Besides being a facet of culture, codes share some basic properties with other expressions of culture. Code is socially shared, endogenous, and both sticky (because there is a shared incentive to use previous codes to understand each other) and fluid (because a social group is free to create new codes, and do). Language code also shares the latter properties with other kinds of culture—the bedrock "values"—which are harder to create instantly in the lab. By studying how codes are developed and change, and their efficiency properties, we hope to learn something about how other elements of culture might develop, change, and solve organizational coordination problems more cognitively efficiently than thick manuals allow.

Based on this link between a group or organization's culture and its codes, we explore whether differences in the kinds of culture or codes a group has produce systematic differences in the kinds of outcomes that group experiences. We use two new studies to explore such relationships.

The first study compares the effectiveness of code creation by a single "founder" compared to an egalitarian consensus that emerges when different subjects contribute to the development of the code. Having one person creating code may be more efficient—a founder advantage – because the code may be more coherent or thematic, and easier to remember. On the other hand, if different people each have ideas about what to call the

pictures, an egalitarian process that chooses the best code for each picture will be better than any one player's personal code.

In fact, having a single culture-making founder is efficient. However, we noticed that even though participants in the founder-dominant culture earn more money than their identical-twin counterparts in an egalitarian group, they are less satisfied in their jobs (according to self-reports). Therefore, in a follow-up experiment we measure how much participants in founder and egalitarian cultures contribute in a public good contribution game after creating code. The public good contributions are a measure of "camraderie", team spirit, or how much people help each other out. Even though the subjects in the "egalitarian" condition (who each get to create bits of code, rotating among themselves) earned less money in the code phase, they contribute more in the public good game. It appears tentatively that while egalitarianism in control of code-creation is bad for efficiency, it is good for creating camraderie in a later interaction that requires cooperativeness.

The second study explores whether having a history of flexibility or variance in code-creation is an advantage in facing a new situation. A typical intuition about organizations is that those with strong cultures—i.e., those with a consistent set of shared values—may be handicapped in adapting to new environments when new code is needed to describe new situations. We bring this challenge to life experimentally and find that strong cultures do worse in confronting new situations.

2. Previous thinking

a. Codes and economic theory

The idea that shared codes are an important part of economic organization is traceable at least to Arrow (1974), who anticipated limits on communication and mutual rationality as constraints on organizational scope (see also Marschak, 19?? on the economics of language). Arrow insightfully thought of codes as components of social capital that were productive, shared, and (because they are shared) may be difficult to replicate and transfer.

Later, researchers in business strategy interested in differences in firms, and why some firms succeed while others fail, took up the problem of how firms could create organizational assets which can create advantages that are sustainable in the face of competition (e.g, Barney, 19??). One feature of assets which create rents is "social complexity". Loosely speaking, a socially complex asset cannot be taken away from a firm by a private worker and replanted elsewhere cheaply, because no single worker understands the asset well enough, or can overcome the coordination and bargaining problems of assembling a team to transport the shared asset.[2] Codes are a good example of socially complex assets because language is inherently social, and cannot be recreated without time-consuming shared experience.

Despite Arrow's clear explication, few theories have tried to carefully define codes and how they can create value. Cremer (1993) wrote on this topic:

> "If some type of information is to be transmitted repeatedly, then it is efficient to generate a specialized code to transmit ideas or fact that recur repeatedly. For instance, sailors find it convenient to speak of a 'sheet' rather than a 'rope that is used to flatten the sail.' Similarly, in organizations, buildings, jobs, and types of clients will be designated by their own code word. Very often this word will not be new, but rather an existing word used to refer to a very specific object or phenomenon."

The important paper by Cremer (1993) emphasized that codes are an asset which create returns—but codes are also an asset with special and interesting properties.

An earlier breakthrough paper is Kreps (1990) (see also Camerer and Vepsalainan, 1988). He followed Schelling (1960) in emphasizing the fact that people can often solve pure coordination problems, in which people want to match strategies, but don't care which match is picked. Kreps defined culture as focal principles that tell employees how to solve pure coordination ("matching") problems. His idea is illustrated by an advertisement for Southwest Airlines, shown in Figure 2. The ad shows a hypothetical quiz given to employees which asks what to do when a Customer has an hour layover in the airport and needs baby formula (based on an actual episode). Answer "C"—"Go find some formula and pick up a coloring book for her older child" is the right

[2] Silicon Valley is perhaps the exception that proves the rule. Many spinoff firms have resulted in which modest-sized teams spin off to another firmand prosper. These spinoffs have succeeded because the assets of the firm are largely human capital (in the famous phrase, the assets "walk out the door every night"). California state laws also limit no-compete clauses in employment termination, which enable employee teams that leave to start up effectively when they leave.

answer (an example of how stories encapsulate cultural values). Working at Southwest is essentially a series of matching games which are very much like answering the question posed in Figure 2. (The fact that the word "customer" is capitalized into "Customer" is also a subtly coded way to express Southwest's corporate values: viz., customers are important; they are "Customers", not merely "customers".)

The organizational research on culture is actually not that far removed from the economists' conception. Schein (1983, p 6) emphasizes that the culture's "assumptions and beliefs are *learned* responses to a group's problems of *survival* in its external environment and its problems of *internal integration*." [italics his] Translating into game theory code: Culture is an equilibrium behavior which adapted (*learned*) or evolved (*survived*) and achieves coordination (*internal integration*).

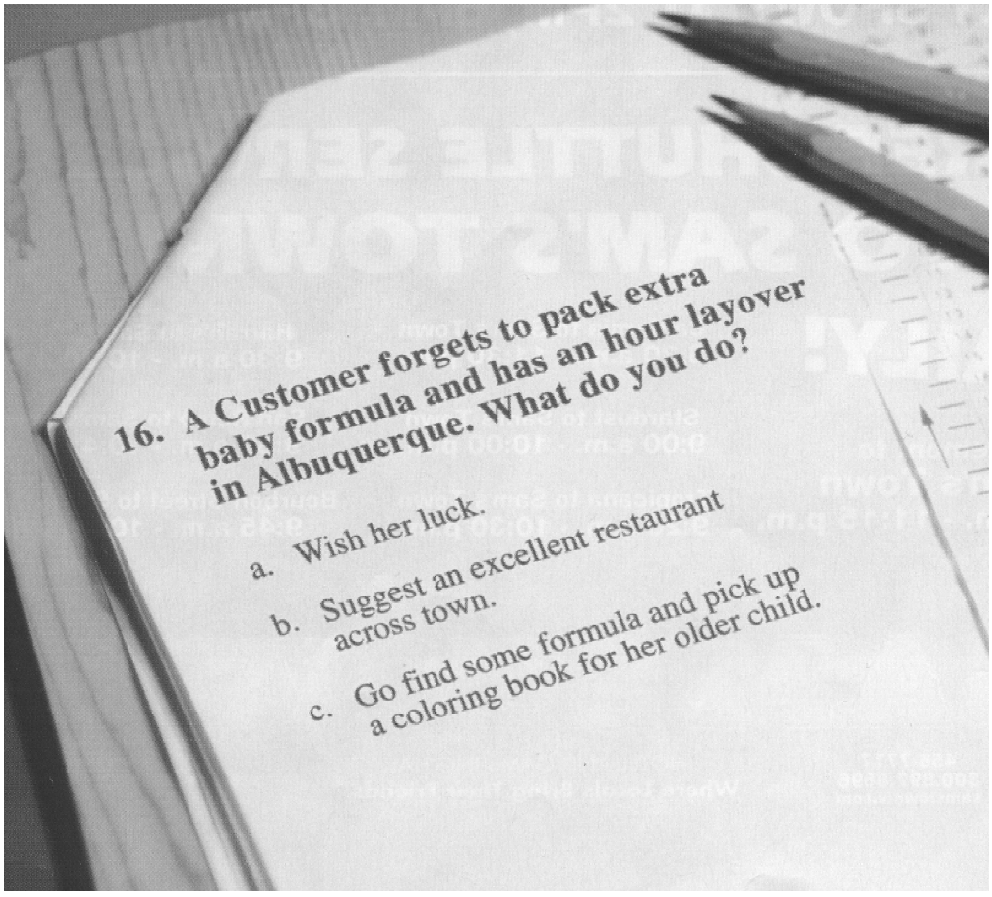

Figure 2: A Krepsian employee quiz at Southwest Airlines

Garicano, Prat and Cremer (2003) explore further properties of codes and how they are integrated. They define codes as mappings from a large set of perceived set of states into a narrower set of linguistic terms, presumably constrained by working memory and the need to communicate in a common (and hence, coarse) language.[3,4] They show that when the processing costs of mapping states onto word falls, firms should first integrate (using common codes) then create hierarchies in which "translators" bridge gaps between workers with different codes. Wernerfelt (in press) uses a similar concept and shows when common codes are more efficient.

Like the papers cited above, our experimental paradigm focusses on games in which organizational participants have a common goal—to communicate situations of organizational interest, through codes, rapidly and accurately. Since the participants have a common interest, they are "teams", in the jargon of Marschak and Radner (19??).

Of course, a huge amount of organizational theory concerns the conflict of interest between principals, and the agents who take actions that affect those principals. Team structures ignore these incentive conflicts. There is no doubt that incentive conflicts are important. But even when interests are aligned in teams, mundane and interesting problems of communication loom large in determining which organizations work well and which dysfunction. (An example is Heath, 20??, who studies "coordination neglect" in which teams underinvest in facilitators who do not directly contribute to production but purely coordinate production.[5]) Furthermore, starting with the base case of team incentives creates a foundation onto which incentive conflicts can

[3] In one respect Garicano et al's approach is exactly the opposite of ours. Our subjects are clearly able—rapidly—to assign a unique code to each of up to 24 pictures. Partly this is because they have to pick exactly the right picture so they need unique codes, but we were also struck by how adept college students (at 3 different schools) were at creating codes across a wide range of pictures. Presumably the coarse encoding (categorization) Garicano et al have in mind kicks in when there are thousands of pictures and some differences among them are economically irrelevant (e.g. whether a customer is paying with five one-dollar bills or one five-dollar bill). Then it is inefficient to have a unique code for each situation.

[4] Jackson and Fryer (2003) study the closely-related problem of optimal categorization (minimizing within-category difference and maximizing between-category difference), with applications to labor market discrimination. Their work could be seen as a basis for optimal code creation.

[5] A canonical example of a facilitator is a "coxswain" in rowing. The coxswain, usually a light-weight person, sits in the front of the boat and shouts out "stroke" at very regular timed intervals. In the language of game theory, the coxswain is a human correlating device who synchronizes the strokes of the rowers and perhaps motivates them. (Without a coxswain, rowers do not perfectly synchronize so some end up pulling their paddles out of the water while others are putting their oars in the water.) The fact that boats use coxswains— who add weight and don't row, slowing down the boat— is testimony to the advantage of having a person who purely coordinates.

be grafted to see how language-coordination (team) and conflict problems interact, in a later phase of research.

3. A picture-naming paradigm for studying organizational codes

In our experiments, code was created with an experimental paradigm that required subjects to develop names for picture. The two studies described in this paper used two kinds of pictures—scenes of people working in offices (e.g., Figure 1); and scenes taken from internet sites, of people in Asian cities.

In the experiments, groups of participants see a common set of pictures (see screenshots below). One person, designated the "manager", knows which subset of pictures are targets (which pictures are targets typically varies randomly in each experimental round). The manager's job is to describe the pictures to the other subjects—called "employees"—so that the employees can choose the right target pictures quickly.

The payoff function is typically a starting payment of (US) $.60 per round, minus a penalty for the time taken and a penalty for mistakes (i.e., for an employee choosing the wrong picture). For example, in study 2 below, employee participants lost $.005 for each second it took to correctly choose *all* the pictures (i.e., $.30/minute), and $.50 for each mistake. A manager's payoff in a round was the average of all the employees' payoffs. The payoff scheme provides a crisp way to measure efficiency of cultural codes. Good codes permit employees to choose the right pictures (minimizing mistakes) rapidly (minimizing the time penalty), earning the most money.

Before proceeding, it is useful to cement an example of what picture-naming corresponds to in naturally-occurring organizations. Imagine a cop sitting in a car listening to a dispatcher. The dispatcher is receiving information from a victim who describes a scene to her. Suppose the victim has been stabbed by her husband and is bleeding from a shoulder wound; the husband is fleeing on foot. The dispatcher and cop have some shared understanding of the wide range of scenes that could be described by victims (the husband could still be in the house; the husband might still brandish the knife as he marches down the street, shouting; the husband might have threatened to return; the "knife" might really be a fork; the victim might not know her attacker; and so forth).

It is as if the dispatcher and cop have in front of them an array of snapshots labelled by features the cop should know about in deciding how to respond—except that the snapshots are in their brains, rather than in front of them. The dispatcher's job is to quickly describe to the cop which picture applies, using a code that balances the need for relevant detail with the need for a very rapid response.

Now picture a casino manager who gets a phone call from an events coordinator who has met a high-roller at the airport. The high-roller might be in a bad mood, or might have told the coordinator he isn't going to gamble that night, or will stay a week, or has brought a "special friend" (or two). The event coordinator has a range of snapshots in mind and sees one of them in front of her. Her job is to convey to the casino manager which snapshot the casino manager should conjure up, to decide how to treat the high-roller.

Here's a third example: A waitress shouts, amid the din of a busy kitchen, a special request to the chef. The waitress has a `snapshot' in mind (a particular customer, his or her request, the likelihood of a big tip, whether the customers come regularly, and so forth) and tries to find the words to get the chef to `see' the same snapshot.

One more: A location scout is in Pittsburgh looking for suitable remote locations for a movie. The director is looking for a circa 1950's location with the right architecture, climate, and pool of down-and-out movie "extras" to shoot a scene in a dive bar. The location scout calls the fussy director on his cell phone and describes the scene in front of him.

The police dispatcher, event coordinator, waitress, and location scout all face the same problem. They and their counterparts share a common visual vocabulary of scenes and the person on the spot must convey which of the many possible scenes they are facing, to coordinate, rapidly, the right move to make. (Of course, the dispatcher and cop aren't literally looking at pictures, but they have episodic memories with crucial features of each picture stored in their brains, along with some information about how much nuance the other person knows and cares about.)

The picture-naming paradigm has the essential features of all four situations: Both players know the scene they are naming is one of a common set. The person who sees the scene (i.e., who is "on the scene") must convey its essence to another person, who must

choose how to respond rapidly and accurately. Codes develop. Good codes are based on shared understanding, and produce efficiency through speed and accuracy. Appendix B gives some lively examples of code and how they can prove useful in organizations to create fast, accurate coordinated activity and for other purposes (like marking insiders and outsiders, and relieving tension or disgust).

It is important to note that the experiments described in this paper are more exploratory, or pre-theoretical, than many other types of experiments which carefully test theory. Little formal theory guides the design (though theory could *follow* the data-gathering, rather than vice versa.) Instead, the idea is take a naturally-occurring phenomenon we have some loose understanding of, and see if we can bring it to life in the lab. The laboratory paradigm both gives a clear shape to the phenomenon, and provides a platform that future experiments can build upon. The complaint that "culture is more complicated than your experiment" is a legitimate critique if-- and only if-- it can be cast in the form of an alternative experimental design which includes the interesting complicating features.

4. Early experimental studies of organizational culture

Very few studies in economics explicitly attempt to create organizational culture in the laboratory.[6] This is not too surprising, since this requires creating the kind of "socially complex asset" described in a previous section among a group of unfamiliar laboratory subjects. Therefore, to create something like culture in the laboratory requires having a group of subjects interacting repeatedly with each other in a way that allows them to develop such an asset – usually producing only one observation per group.

One very simple form of such a collective asset is the ability of a group (or it's failure) to coordinate effectively. The ability of a group's members to take actions independently in a manner that complements the actions of others is an important part of what social assets such as culture allow (see Kreps, 1990). Therefore, the same groups of subjects repeatedly playing a coordination game represents one simple way in which such an asset can be created.

[6] There is an older tradition in organizational research (Cyert and March, chapter 4, 1963), Leavitt, Weick...more here

Knez and Camerer (1994, 2000) studied organizational issues in "weak-link games" in which a player's payoff is an increasing function of the minimum effort level chosen by the group, and decreases in how much extra effort the player put in (above the minimum). The effort levels that emerge can be thought of as work norms. They found that when two small groups merged into a large group, productivity fell. However, they also found that publicly announcing a small bonus for group efficiency improved productivity dramatically.

Weber (2003) also used weak-link games to explore how the ability of a group to solve such coordination problems is affected by the way in which a group grows, using this as a metaphor for how growing firms encounter coordination and cultural problems related to size. Simple laboratory "firms" have a difficult time coordinating action when they have many members. However, starting with a small group and growing slowly allows a large group to emerge which is efficient. It is also crucial that new entrants be aware of the group's history – reflecting the importance of training and acculturation programs regularly used by real-world firms.

In another study that creates simple laboratory "firms," Weber and Camerer (2003) use the same code paradigm as here to show what happens when small laboratory firms are allowed to simultaneously develop efficient, but idiosyncratic, codes and are then merged. The results replicate many aspects of real-world merger failure arising from cultural conflict: decreased performance due to differences in codes, overly-optimistic forecasts of post-merger performance (meaning that subjects fail to understand how difficult it will be to combine two cultures), and attribution of blame to members from the other pre-merger firm.

5. Study 1: Centralization

Cultures vary by how much power is either centralized and hierarchical, or decentralized and egalitarian. These differences influence the extent to which discussion is allowed, input from subordinates is welcome, or decisions are made by subordinates.

In an extremely centralized organization, orders flow from the top down. Little discussion, input, or dissent, is allowed from subordinates (e.g., the U.S. Army). People

in such firms think "the only way to manage a growing business is to supervise every detail on a daily basis" (Schein, 1983, p. 14). Lawrence and Lorsch (1967) discuss an organization where "The traditional method of resolving conflict in this company is by taking it upstairs" (p. 120).

At the other extreme is a decentralized, egalitarian, or participatory culture (for the purposes of our experiments, we use these terms synonymously). In these organizations, power and authority are shared (relatively) equally and formal leaders authority to subordinates and seek their input. This culture is expressed by the belief that "ideas can come from anywhere in this organization, so we must maintain a climate of total openness" (Schein, 1983, p. 14). Lawrence and Lorsch (1967) describe an organization in which "Top management has told [lower-level managers], 'We want you to decide what is best for your business, and we want you to run it. We don't want to tell you to run it.'" (p. 144). In a striking example of how code expresses culture, Ott (1989, p. 20) notes that "the Navajo language contains no words meaning superior, subordinate, boss, or hierarchy."[7]

The advantages of centralization are related to the idea that a charismatic founder creates a culture which can give a firm a long-lasting advantage. For example, Schein (1983, p. 221) notes that "cultures do not start from scratch". He writes

> Founders usually have a major impact on how the group defines and solves its external adaptation and internal integration problems. Because they had the original idea, they will typically have their own notion, based on their own cultural history and personality, of how to get the idea fulfilled. (p 210).

Ott (1989) discusses an accounting firm where the founder strongly believed that customers knew little about accounting matters and should be ignored. The founder's belief created a clear cultural "rule" that clients should not be allowed to participate in the accountant's work. Southwest chairman Herb Kelleher said "personality is culture" and talks freely about how he encouraged Southwest's values which expressed his own (cite) As noted earlier, having a firm's culture which is an expression of its founder's (or founders') own values— "personification"— can be cognitively efficient (taking advantage of the well-developed human capacity to store and organize knowledge about

[7] TBA: Field studies of centralization and profitability: What do we know?

other people in the form of episodic memories and extracted traits). If Disney workers know a lot about the founder's values, it may be easy for workers to answer the question "What does the company want me to do in this situation?" with the heuristic "What would Walt [Disney] do?" Similarly, there is little doubt what the Playboy magazine empire stands for, because it is personified by the behavior of founder Hugh Hefner himself— the playboy and corporate icon.

The personification heuristic can even survive the founder's death, making founders potentially immortal corporate assets. For example, at the University of Chicago in 1993 or so one of us (CC) participated in a tenure case at which George Stigler's opinion of the candidate's work was mentioned. Stigler wasn't been able to attend the meeting...because he had been dead since 1991. It sounds preposterous to let a dead man influence a tenure case. But Stigler's writings and philosophy embodied a lot of Chicago values—he was verbal, caustic, clear-thinking, not overly mathematical, knew a lot of facts, and believed in markets. Asking "What would George think?" was to ask "What do we think?"

A. Experimental design

The experimental design has two conditions. In the "fixed-manager" condition, one subject is always the manager, who sees the target pictures and creates code to describe them. Three other subjects are always employees. In the "rotating-manager" condition, the role of manager rotates in a cycle across the four subjects (i.e. one subject is the manager each period and the other three are employees).

These treatments are a simple way of creating the organizational distinctions discussed above. The fixed condition has specialization of roles, centralization and hierarchy (because one manager does almost all the talking), and measures founder effects because the single manager is like a founder.[8] The rotating condition has no role-specialization, is decentralized (everyone gets to manage), and is egalitarian.

The first phase of the experiment lasts 20 periods. The performance measures are speed of naming pictures and mistakes, which are combined in the overall payoff

[8] An interesting further treatment would be to run longer experiments in which the founder leaves and others take over (mimicking the persistent afterlife of Disney and Stigler), to see if the advantage of having a single founder persists or not.

measure. Good code makes money. If there are positive founder/centralization effects, then subjects will earn more in the fixed condition. On the other hand, if rotating among different subjects provides an opportunity for the best pieces of code to be assembled, then subjects will earn more in the rotating condition.

The second phase of the experiment asks whether centralized or decentralized cultures are better at dealing with change. Many studies in business strategy suggest that the dynamic capabilities of the firm—their ability to adapt to change—are influenced by the extent to which information, perspectives and ideas are shared. Most studies show that more sharing leads to more adaptability (Iansiti and Clark, 1994; Henderson, 1994; Cohen and Levinthal, 1990; Lenox, 2002; cf. Lawrence and Lorsch, 1967).

To study this experimentally, in the second phase of 15 periods, the groups have either environmental change (new pictures) or growth (one new employee).

B. Methods

Subjects were Carnegie Mellon undergraduate and graduate students who had not previously participated in any other experiment involving this culture paradigm.[9]

Groups of either four (New Picture sessions) or five (New Employee sessions) subjects arrived and were given numbers randomly. In New Employee sessions, the subject number 5 was asked to leave and return in 30 minutes, in exchange for a payment of $5 for punctuality.

The remaining participants received instructions on performing the culture task (available from the authors). They all received the same eight numbered pictures on four sheets of paper (arranged in different positions, so the manager could not say "pick the upper left one"). The pictures showed office environments (Figure 1; and see Weber and Camerer, 2003).

[9]Subjects were recruited from an e-mail list of interested students. The e-mail announcement requested that only participants fluent in spoken English sign up While subjects were waiting for the experiment to start, the experimenter also announced this request and asked if everyone waiting was a fluent English speaker. In spite of this, there were two sessions (one Fixed, one Rotating) in which one subject showed somewhat limited command of spoken English. While it would have been nice to avoid these instances, a more intrusive measure of English language skills was rejected because of the complicated and arbitrary nature of such tests and because it might affect (especially rejected) subjects' attitudes towards the experimenter and behavior in future experiments. The results do not change if we omit these two sessions.

In each period, the manager knew which *four* (4) of the 8 pictures were target pictures in that round. The four subjects were seated in tables facing in opposite directions. Each employee had a stopwatch. When a period started, the manager began describing the four pictures to the employees and all employees started their stopwatches (the experimenter also kept time). Once a particular employee wrote down four picture numbers, he or she stopped the watch and recorded the time. Once all the employees were done, the experimenter read the correct numbers aloud and subjects recorded whether or not the number they had placed in a particular position was correct, keeping track of the number of mistakes.[10] The payoffs to each employee for completing the task in each period were equal to \$0.60 minus \$0.01 for each second that it took to complete the task and a \$0.10 penalty per mistake. The manager received the average earnings for the three employees.

The first 20 periods of this task took about 30 minutes. In the fixed sessions, participant 1 was always the manager. In the rotating sessions, participant 1 was manager in period 1, participant 2 was manager in period 2, and so forth.

After the 20th period, subjects got a brief new instruction (see Appendix). In some sessions ("New Picture"), the employees were told that they would perform another 15 periods of the task in the same manner as before, but that they would receive a new set of pictures, also of office environments. In the remaining sessions ("New Employee"), participant 5 was added to the group for the remaining 15 periods. Participant 5 had received a copy of the instructions upon returning to the experiment and had the opportunity to ask questions about the task. In both kinds of sessions, the role of manager either remained fixed or continued to rotate as it had in the previous 20 periods. In New Employee sessions, participant 5 was always in the role of employee.

At the end of both periods 20 and 35, every participant wrote down "the word or phrase that best describes the term your group has used to refer to the picture." We conducted a total of 34 sessions using 160 subjects (see Table 1).[11]

[10] The experimenter stood between the subjects at all times, from a position where he could easily see if a subject was writing when he or she was not supposed to. This situation did not arise.

[11] [[add in dates and subject pools and any idiosync variations across sessions]]

	New Employee	New Picture
Founding manager	12 (60)	5 (20)
Rotating manager	12 (60)	5 (20)

Table 1. Study 1 number of sessions (subjects) by condition (fixed vs. rotating)

C. Results

We first explore differences between fixed and rotating groups over the first 20 periods (pooling across New Employee and New Picture groups[12]). The clear finding is that the fixed groups earn more.

Figure 1 shows average completion times in the fixed and rotating conditions. Completion times start at around one minute and end at around 15 seconds, following a typical graceful learning curve (well-fit by a power law). In the first period there is no difference in completion times between the fixed and rotating groups, but after that a

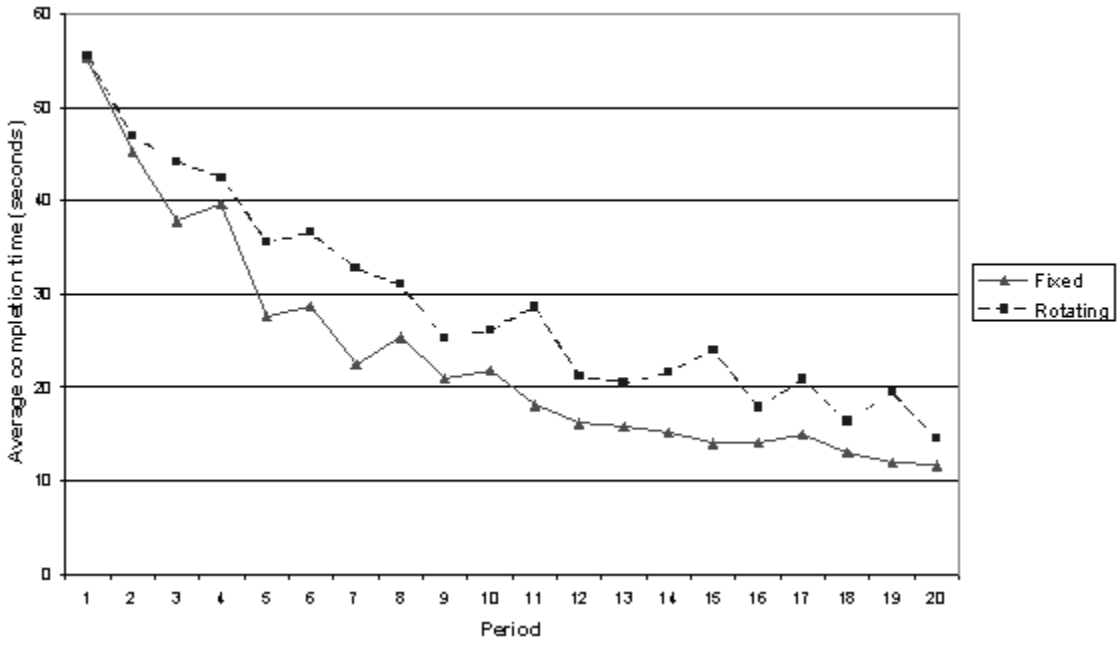

Figure 3 : Average completion time, periods 1-20, fixed vs. rotating

[12] There are no significant differences between the New Employee and New Picture groups within each fixed or rotating condition. The lack of a difference is not surprising since the instructions and procedures over the first 20 rounds were almost identical for New Employee and New Picture sessions. In fact, the only difference was that subject in the New Employee condition knew that another subject had drawn the participant number 5 and had left the room. While they might have anticipated that the participant would be returning later, the fact that we find no difference indicates that this consideration was not an important influence on behavior.

clear advantage of the fixed groups (thick line) emerges (which is significant in all four 5-period blocks at p<.05 by t-test). The average number of mistakes in the two conditions starts at .50 per period and drops to .10 by period 20, and there is no perceptible difference between the two conditions (see periods 1-20 in Figure 7 below).

Figure 4 shows that average earnings[13] are persistently higher in the fixed condition. (This is no surprise because the completion times are lower in the fixed condition and there is no difference in the numbers of mistakes.) The differences are significant in all four 5-period block at p<.04 (two-tailed t-test).

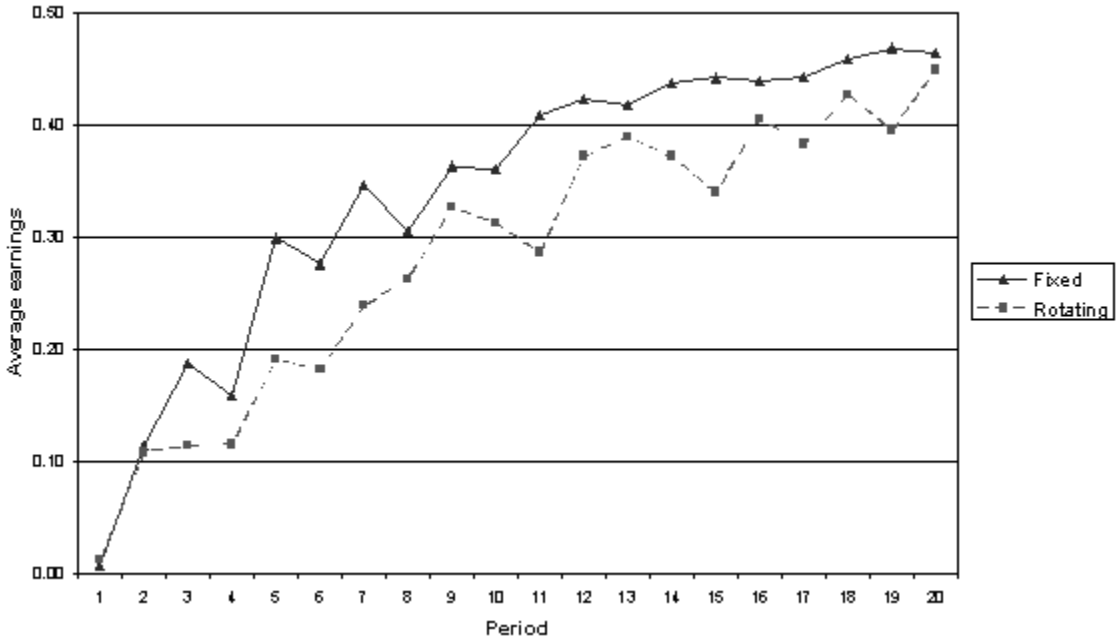

Figure 4: Average earnings, periods 1-20, fixed vs rotating groups

D. Introducing new pictures

The fact that fixed groups are faster is also manifested by the length of phrases written down by subjects at the end of the 20th period which are consensus descriptions of the code that emerges describing the pictures. The fixed-group code length averages 20.6

[13] Figure 4 plots the average earnings across the three employees in each group. Note that this is also how the manager is paid.

characters per picture, while among the rotating groups the average is 23.8 (p=.07 using a one-tailed test, with the session as a unit of observation).

In ten of the sessions (5 fixed, 5 rotating), we introduced a new set of office pictures after the 20th period. Figure 5 shows the average earnings per period for the two different kinds of group in these sessions. Rotating groups continued to perform worse when the new pictures are introduced.

As Figure 5 shows, rotating groups continued to perform worse when new pictures were introduced. The fixed groups have higher average earnings in every period after the introduction of new pictures. In fact, the gap in earnings widens. (The fixed groups also have faster completion times in the early New Picture periods 21-25, $p<.05$ by a one-tailed t-test in each period). It appears that the fixed groups do not only learn to create code quickly in the first phase of 20 periods, but this code-making capability enables them to describe new pictures more quickly.

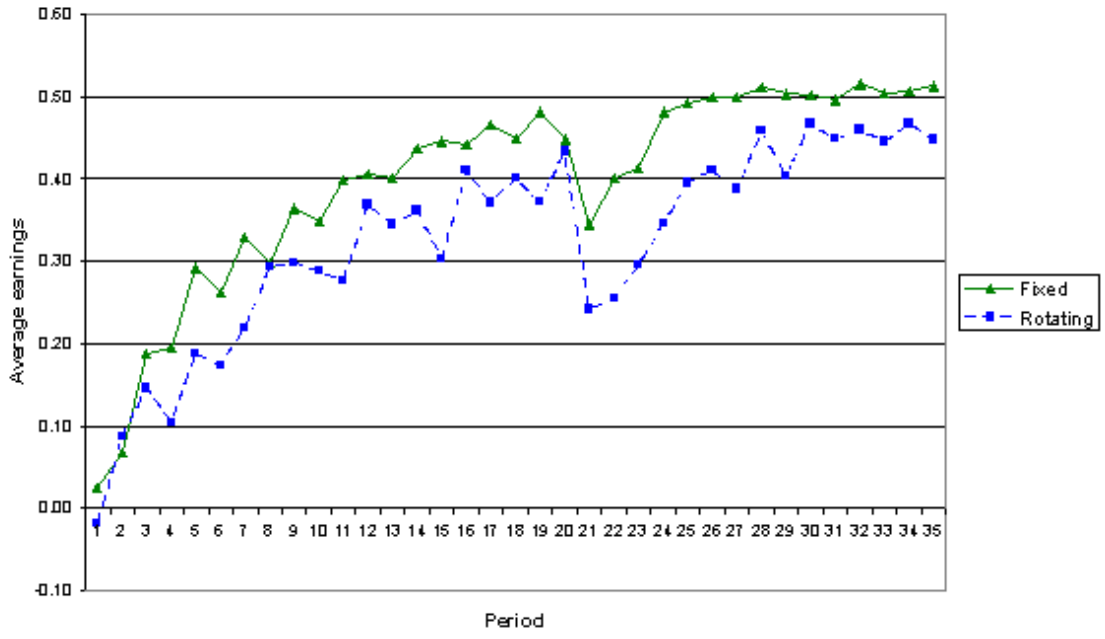

Figure 5: Average earnings, New Picture conditions (new pictures introduced in round 21)

E. Introducing new employees

In 24 sessions (12 fixed, 12 rotating), a new employee was added to the group after the 20th period. This subject remained in the role of employee (i.e., he or she never was a manager).

In seeing what happens, it is useful to distinguish between he incumbent employees who were previously in the firm in periods 1-20 ("old" employees) and the entrant employees who started in period 21 ("new" employees). Figure 6 shows the average earnings of both kinds of employees. The positive earnings advantage of fixed groups continues for old employees after period 21 (though it is only marginally significant in the three new 5-period blocks at $p<.065$, $p<.06$, $p<.03$ one-tailed).

However, being in a fixed group is a slight earnings *handicap* for new employees who start work in period 21 (although only weakly significant in the period 21-25 block at $p<.09$, one-tailed). The earning handicap is not due to slower completion times (in fact, the new-fixed employees are slightly faster than the new-rotating ones). It is entirely due to a large difference in the number of mistakes, shown in Figure 7. Even though these

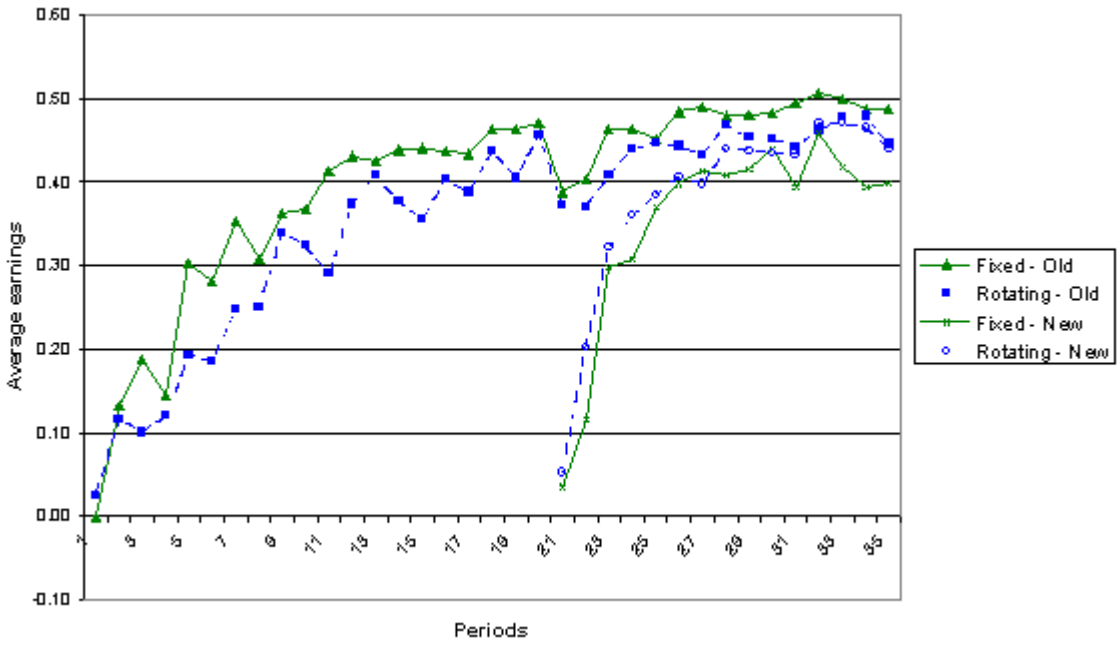

Figure 6: Average earnings in New Employee conditions by type of employee

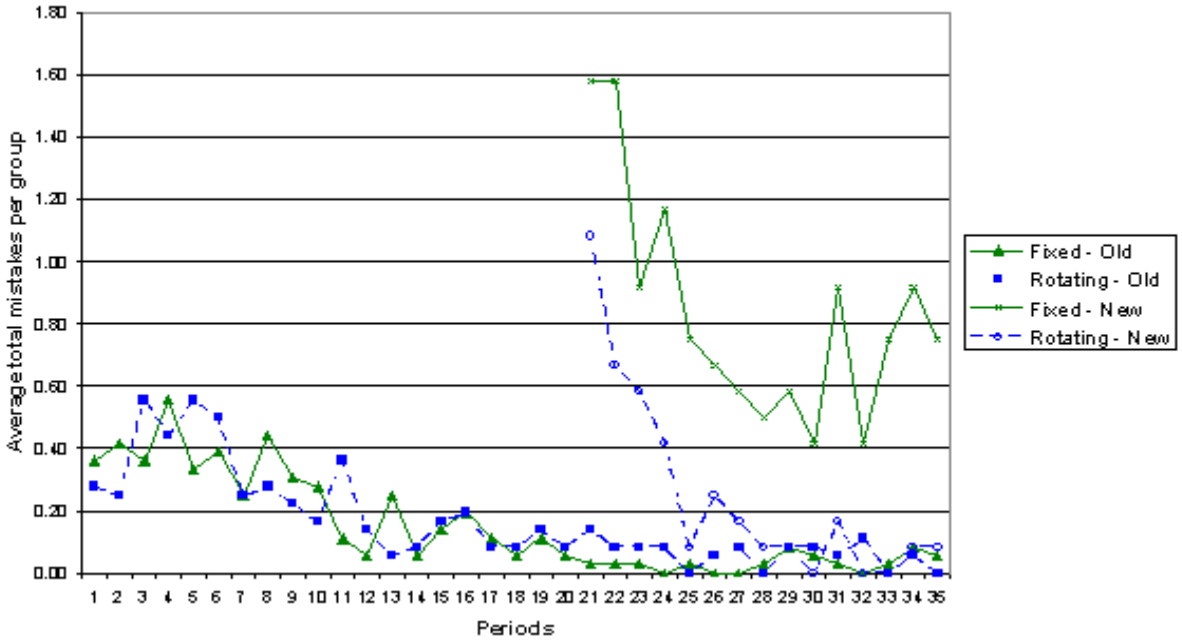

Figure 7: Average number of mistakes among new and old employees

mistakes are rare and variable, the difference between new and old employees in the fixed groups is modestly significant in all three 5-period blocks after period 20 ($p < 0.05$, $p < 0.07$, $p < 0.03$, one-tailed t-test).

The difference in mistakes is partly driven by the fact that in three of twelve fixed groups, the manager essentially gave up on incorporating the new employee (leaving him or her to guess at some of the pictures and make mistakes), which never happened in the rotating groups. In these three sessions, the new employees averaged 1.8, 2.2, and 3.4 mistakes per period over the last five periods.

Thus, introducing new pictures or a new employee does not reverse the superior performance of old employees in the fixed groups. However, the culture of rotating groups appears to be better suited for incorporating new employees. It appears that in the fixed groups, the single manager (who has been creating code for 20 periods) is just not good at incorporating a new employee and often gives up, which creates a lot of mistakes. An image springs to mind of a gruff, cantankerous old-time CEO who has always (in periods 1-20) been surrounded by lifelong employees who knows what he means. The old-timer doesn't "suffer fools gladly", and treats the new employee as a fool.

F. Camraderie and public goods contribution

In running the fixed and rotating sessions, subjects appeared to have more fun in the rotating sessions—despite the fact that the fixed-condition subjects earned more money. One possibility is that subjects like to talk, rather than listen (at least some of the time), and being manager means you get to talk more. This preference may be related to the value of control and concepts of employee empowerment and—literally—"voice" (cites). Another possibility is that rotating the manager role creates a stronger sense of group identity.

One way to measure apparent "team spirit" behaviorally is to play a game after the picture-naming phase which gets closer to concepts of organizational values. We use a linear public-good contribution game (e.g. Ledyard, 1995). Subjects are endowed with 50 tokens. Each token they keep earns them $.05. Each token they contribute earns the *group* $.10, which is shared equally. Since the $.10 public good is divided four ways, the private return from investing a token is $.025. Since this is less than the $.05 from keeping it, selfish players will free-ride and not contribute anything. But if everyone contributes everything, the group payoff is maximized.

In an organizational context, contribution in these games is a behavioral measure of "team spirit" (see Sugden, 20??), camraderie, or how much people will "pitch in to help others out". It provides a simple way to begin to study the link between what happens during the process of code-development, and its effect on group identity or expression of social preference.

We conducted 12 more sessions (six fixed, six rotating) with the same protocol above. After 20 periods they played the contribution game. The completion times and earnings in the first 20 periods strongly replicated the results above (fixed and rotating participants earned averages of $.39 and $.31, respectively, and took 19 seconds and 27 seconds.)

Figure 8 shows a histogram of contributions (n=24 in each of the two conditions). The mean contributions are 26.8 (rotating) and 20.2 tokens (fixed) and the medians are 29.5 and 16. The difference is large in magnitude but not significant at conventional

levels (p=.11 by one-tailed t-test). (We are running more groups to test whether the effect is genuine.). Note that the rotating groups earned, together, $1.36 per group[14] more than the fixed groups. In the picture-naming phase, the fixed groups earned an average of $1.67 more. The fact that these two differences are close suggests an interesting tradeoff between the advantage of having a centralized founder in charge of making code (which earns the fixed group more) and the subsequent public-good advantage of the camraderie that seems to be created when code is created in an egalitarian way (which earns the rotating group more).

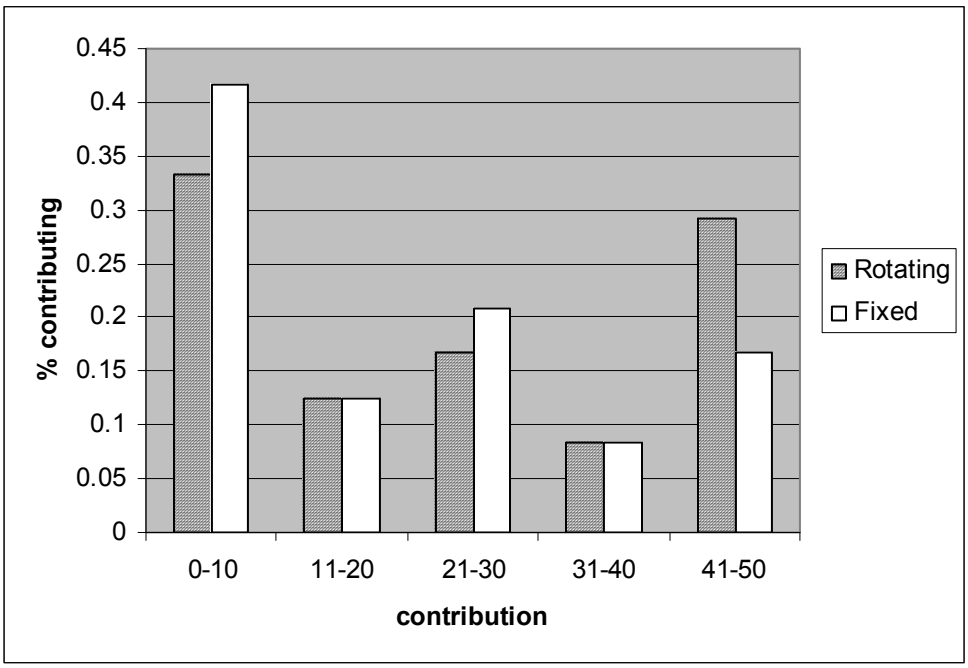

Figure 8: Public good contributions in rotating and fixed conditions

6. Study 2: Is rigid culture a handicap?

A classic issue in thinking about organizations is whether it is good to have a "rigid" or "strong" culture— a set of values, perhaps embodied by deeply "inside" code, which is strongly shared (i.e., low variance across a respondents) and preferrably transparent to outsiders. Kotter (197?) measured strength of culture with self-reported survey responses across companies, and found that strength was positively correlated

[14] That is, the marginal public over private return is $.05/token, the Rotators contributed 6.8 extra tokens per person, and each group had four people. The product (6.8)x(4)x($.05)=$1.36.

with accounting profitability (though see also Sorensen, 20??). Within the Krepsian framework, a strong visible culture is good because it coordinates activity and tells outsiders what the firms' values are (economizing on the cost of attracting, and hence retaining, the right employees).

Rigid culture is often associated with a business that has operated in a stable environment so that the culture has had a long time to take root and does not change much. (In our terms, this is like a business that has faced only a modest of set of scenarios and has not had to create many new cultural concepts to communicate about the scenarios or decide how to value behaivor in new situations.)

Examples like the American auto industry in the 1970's suggest that a rigid culture may be difficult to adapt to changing circumstances. One handicap is "encoding bias" (the tendency for "top-down" expectations about perceptions to filter out unexpected visual cues and bias encoding toward what is expected). A close kin of "encoding bias" is the Whorfian hypothesis—that lacking a language in which to name

Figure 9: An "Asian city" scene used in study 2

objects limits the cognitive ability to perceive differences in those objects (loosely put, language constrains perception). To the extent that encoding bias or Whorfian constraint shapes[15] what is perceived (or socially, can be talked about through mutually-understood perception), a rigid culture may ignore regime shifts in its strategic environment. Another handicap is that in a rigid culture, workers' abilities to create new cultural adaptations may atrophy. This idea suggests that what is important is the ability to nimbly create culture; a rigid culture has not had to develop such an ability.

Our second study explores whether a rigid culture is handicapped relative to a flexible one. To investigate this we created a design with two conditions. In the "flexible" condition, participants see random subsetsets of 8 pictures out of a total set of 24 pictures of Asian cities (Figure 9 is an example). In the first few periods, this group is often seeing brand new pictures and must work hard to invent 24 pieces of code.

In the contrasting "rigid" condition, sees the *same* 8 pictures every time (a single subset of the same 24 the flexible group sees).

Both groups go through a 30-period training phase. Not surprisingly, the rigid groups uses fewer characters, have a quicker completion time, and earn more money in the training phase, because they see the same 8 pictures every time and slide down the learning curve faster than the flexible-condition workers who often see new pictures that take longer to encode [add numbers here].

In the second test phase, both groups see 8-picture subsets of l6 pictures of office scenes (see Figure 1 above). Keep in mind that neither group has seen these office pictures before. The flexible group is used to seeing new pictures, because they saw different 8-picture subsets in their training phase. The rigid group never saw *any* new pictures because they had the same 8-picture subset in the training phrase.

The crucial test is whether the flexible or rigid groups do better on the new set of pictures, which require them to both create brand-new code for office scenes, *and* to adapt to a changing world in which they do not see the same 8-picture set each time. Since both groups are seeing the phase II office pictures for the first time, the natural null hypothesis is that they will develop equally fast and accurate codes. The natural

[15] In psycholinguistics this general proposition is known as the Whorf hypothesis: Add Ocasio here.

alternative hypothesis is that the flexible group has developed more codemaking capability and adapts better to a new (and variable) set of pictures.

The experiment used 63 pairs of subjects at UCLA and Caltech, from June to October 2003).. The experiments were run using CultureX software.[16] Subjects earned $.60 per period, minus a time penalty of $.005/sec (i.e., $.30/minute, an implicit wage of $18/hour) and $.50 for each mistake.

Figures 10-12 show average string lengths (i.e., the number of characters used to describe or talk about all four pictures), completion times, and dollar payoffs in the two conditions.[17] The Figures also show power learning curves which were fit to all three time series, with two dummy variables capturing the estimated difference in initial lengths, times, and payoffs and the steady-state differences (these are reported in Table 2).[18]

The regression specification used is $P_{it}= \alpha+\alpha_R*D_i(R)+(\beta+ \beta_R*D_i(R))f(t)+\varepsilon_{it}$ where P_{it} is the performance measure of the i-th group in period t, $D_i(R)$ is a dummy variable that equals one if group i is in the rigid condition (and zero otherwise), and $f(t)=e^{-\gamma(t-1)}$ is a power learning-curve function with $f(1)=1$ and $\lim_{t\to\infty} f(t)=0$. The coefficients α_R and β_R represent the differences between steady-state performance level and the intercept (as a difference from the steady-state performance level), respectively, in rigid groups relative to flexible groups. (We have not yet corrected for the likely autocorrelation and heteroskedasticity in ε_{it} in the results reported below.) Since most groups converge to common completion times and payoffs, we expect α_R to be small and insignificant. The key question is whether rigid groups are slower and earn less in the early periods (captured by positive β_R in completion times and string length, and negative β_R in payoffs).

[16] Subjects were recruited using web-based software developed by Walter Yuan at the Caltech SSEL Lab. Documentation, screenshots, instructions, and short .avi "movies" of two reenacted periods are available at http://ruffian.hss.caltech.edu/cx-docs/ .

[17] Due to early software complications, in many early sessions the software slowed down completion times due to congestion and slow backup in file creation. We took a conservative route by excluding all sessions in which completion times in the last five periods (26-30)were slower, on average, than in the preceding five periods (21-25). Including these sessions does weaken the significance of the result but not the sign of the effects. MORE HERE

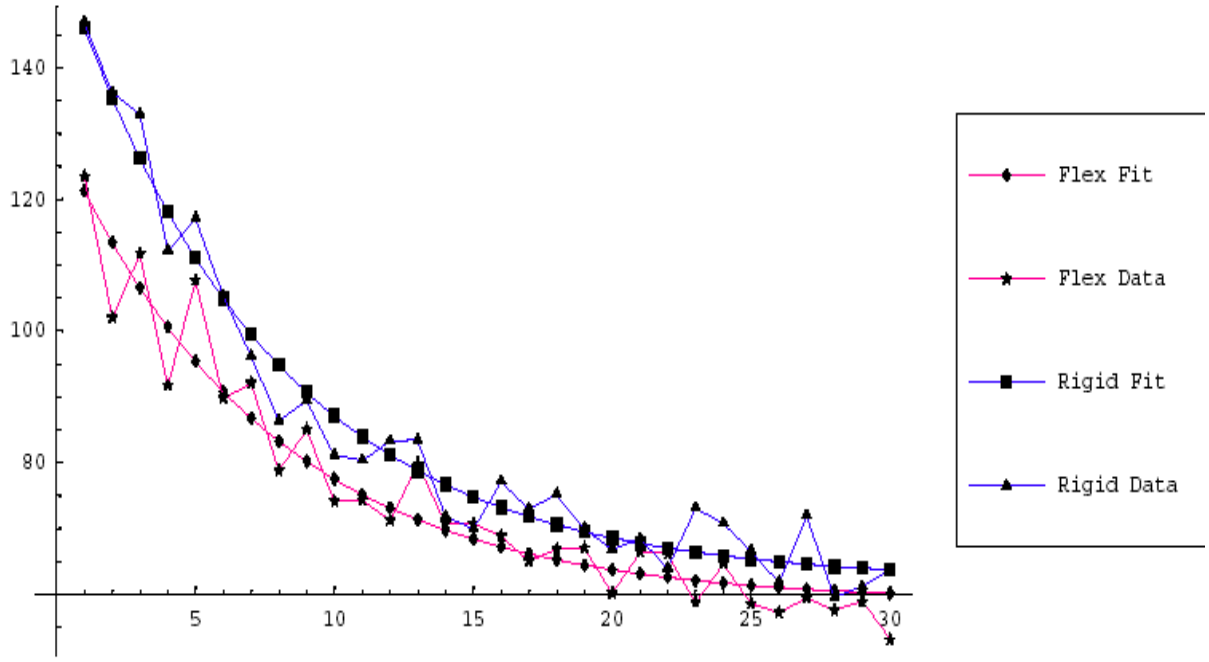

Figure 10: String lengths in characters per period (y-axis) across 30 periods (x-axis) for rigid groups (blue; top line) and flexible groups (pink; bottom line).

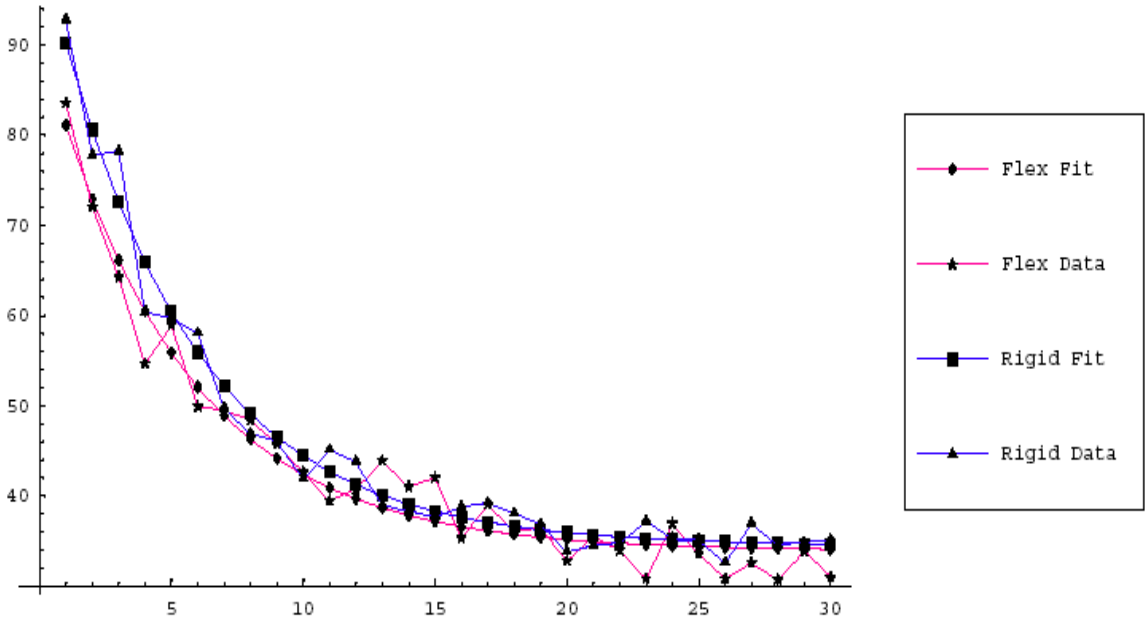

Figure 11: Average completion times (y-axis) across 30 periods (x-axis) for rigid (blue, top line) and flexible (pink, bottom line) groups.

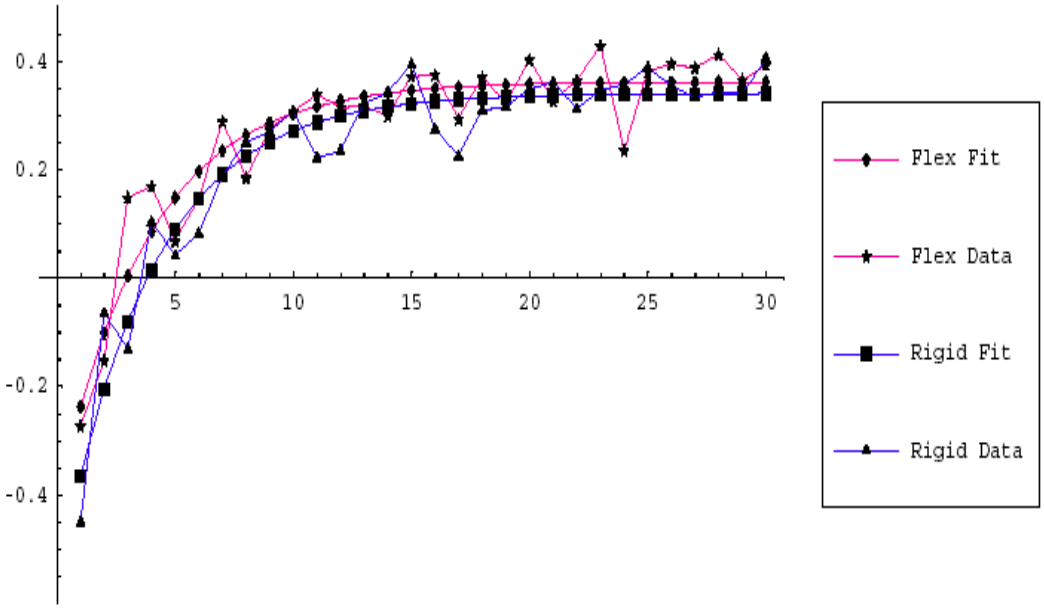

Figure 12: Payoffs in dollars per period (y-axis) across 30 periods (x-axis) for rigid (blue) and flexible (pink groups).

The three Figures, and summary statistics in Table 2, show there is a rigidity handicap at first, which fades over time. (All α_R coefficients are insignificant and all β_R are significant at $p<.05$). In the second phase, with new office pictures, the rigid groups use about 20 more characters at first, take about eight extra seconds to choose all four target pictures, and earn about $.10 less on average, per period. These differences are significant but largely disappear (and become insignificant) after 30 periods.

The results show a modest flexibility advantage (or a strong-culture handicap): Groups which got acclimated to seeing different pictures, and developed general

Dependent variable:	String length	Completion time	Payoffs
Intercept rigid bias β_R	21.54 (7.26)	8.65 (3.71)	-.$106 (.066)
Steady-state rigid bias α_R	3.15 (2.79)	.53 (1.20)	-$.022 (.019)

Table 2: Coefficients (standard errors) measuring rigid group bias

codemaking capability, were pithier, faster, and more profitable when they saw an entirely new set of pictures. The "strong" culture groups were handicapped when forced to develop new codes.

7. Conclusion

This paper uses a picture-naming paradigm to study organizational code. Codes are important because they are a facet of culture, and because they share some of the properties of other cultural values and expressions—they are socially shared, created endogeneously, persistent but changing, and can enhance or inhibit efficient economic action.

Using this paradigm, we investigated two issues:

First, is it better to have a single manager creating code (as in a hierarchical, centralized firm). or to have code-creation rotating among different people? The experiments with 34 groups (and another 12 in a replication) show that a fixed-manager regime is better for creating short and efficient code. However, there is tentative evidence that participants in the rotating-manager experiments had more fun—either because they like to talk and have more control when the code-making role rotates (or they don't like always being told what to do), or feel more like a group. This sense of team spirit or camraderie is expressed by modestly higher contributions in a subsequent public good game.

Second, is developing a code which is rigid—tailored to describe only a small set of scenarios—a handicap when a laboratory `firm' enters a new environment where the scenarios often change? The answer seems to be Yes. Rigid groups which have never been exposed to changing pictures as they develop code use more characters to describe new pictures, take longer to choose them, and earn less money at first. In the steady-state, however, the rigidity handicap disappears.

These studies are really just illustrations of how the picture-naming paradigm might be used to bring organizational phenomena to life. The broader point is that an experimental paradigm commits researchers to a definition of constructs which are otherwise often fuzzy (spawning extensive arguments over semantics), and to predictions about might happen. In the study of organizations, there is a surprising, long-standing gap

between organizational researchers (often schooled in sociology and social psychology) and economists. Concepts like culture, empowerment, team spirit, and charisma are ubiquitous in writings about business organizations, but conspicuous in their absence (with some important exceptions) from organizational economics. Experimental and theoretical paradigms provide one way to potentially bridge the gap.

Organization researchers are likely to say that we have not truly created an organization, or captured its essence. Fair enough. However, our view is that rules of effective scientific dialogue require such a critic to make some effort to say how the experiment should be improved to capture the essence that is missing. If that essence cannot be articulated in terms that are implementable in a laboratory, then it is not well-understood.

There are many directions for future research, which are relegated to Appendix A (along with some casual remarks about interesting things we observed doing these experiments). An obvious question is how the process of culture creation and persistence scales up. While our experiments last only a couple of hours, with mini-"organizations" consisting of only 2-6 people, in principle nothing prevents experimenters from scaling time and complexity. The software used in the latest studies is web-based and can potentially accomodate up to 50 or more subjects communicating with each other. Participants get tired after three hours or so, but an "overlapping generations" design could be used to extend the effective life of a culture to days or weeks.[19] So while it is premature to make any direct claims that the findings below will scale up to much larger groups and time-scales, it is scientifically possible to find out if they do.

[19] Cf. Schotter and Sopher's (in press) experiments on "advice" which use hundreds of "generations", with an outgoing generation giving advice to their "children".

Appendix A: Remarks & new directions

This section reports a few miscellaneous observations and ideas for new experiments which might inspire further work.

A. Observations

Bad code: Caltech incident...

Subitizing and ethnicity: Infants and nonhuman animals exhibit a remarkable ability to recognize the number of objects in a visual set when the number is roughly 4 or fewer. This ability is called "subitizing".[20] In the office pictures we used (e.g. Figure 1), and less so in the Asian cities (e.g., Figure 9) many participants instinctively described the pictures in terms of the number of people prominent in the picture.

In other cases, participants instantly picked out ethnicity or other physical cues in naming picture ("Asian girl", "the ugly blonde"). One conjecture is that when participants are under time pressure, rapid brain processes are used to grab the visual cues which first spring to mind (i.e., are processed most rapidly). These processes may reflect a preference for simple cues (e.g., subitizing of number) and failure of normal inhibitions. The rapid-naming paradigm could therefore be a useful paradigm for studying how people really describe the world around them, when time pressure overcomes their normal inhibitions (cf. Stroop tasks in psychology).

Myths and misnomers: Look again at Figure 9. On the computer screens of participants, the picture is about 2" wide and 1.5" tall[21] (All the groups coded this picture as "lecture", "classroom", "cinema", or "crowd inside". In fact, it is a scene of people sitting on traffic, on bicycles and motorbikes, in downtown Ho Chi Minh City (f/k/a Saigon). The code "crowd inside" is thoroughly incorrect as a literal description of the scene. It is a misnomer, or myth.

Historical examples of persistent misnomers are not hard to find. Spanish explorers called the large South American river they `discovered' the "Amazon" because they thought the muscular, long-haired natives they encountered there, at a distance, where the

[20] E.g. infants. GR book example.

[21] However, the software gives them the ability, with one click, to "zoom" the picture into a full-sized window roughly the size that it is shown in this paper.

legendary Amazon *women*. American "Indians" are so-called because misguided European explorers thought they'd landed in India.

The common mistake of the participants in calling a street scene "crowd inside" illustrates two points. First, once a code is adopted, it may blind observers to features of the scene which might otherwise lead them to self-correct their label (e.g., the lights—headlights of bikes and scooters—; the balloon in the upper right, and the two girls cozily sitting together on one motorbike in the lower left). Second, even if the code is technically wrong, since its purpose is to coordinate action, it persists (like the labels Amazon and American Indian). In the UCLA experiments, we asked some subjects leaving the experiment whether they realized the scene was a street scene rather than a "lecture" or "crowd inside" (showing them an enlarge version of the picture). One subject said she had zoomed on the picture and noticed it was a street scene, not a "lecture"; but she figured it was too late, and not worthwhile, to change its code name to something more accurate. The fact that language not only describes objects, but coordinates mutual understanding with some irreversibility, means many labels which are commonly used in language may be quite false, perpetuate the falsehood underlying them but persist anyway (e.g., "The King of Pop"?).

Gender: Our research assistant Galen Loram noticed that same-gender groups seemed to earn more in the flexible-rigid study 2. Statistical analysis showed a modest effect but the study was not designed optimally to find such it so further work is needed. If there is such an effect, it suggests a "Mars-Venus" dynamic in which same-gender groups have some background understanding of efficient rules of turn-taking, interruption, and subservience. The picture-naming paradigm could then be used to shed light on how strong these effects and their efficiency implications.

B. Other studies

Hypercode: The experimental paradigm permits us to study what happens when participants see the same pictures for up to 200 periods or so (or in an overlapping generations design, perhaps up to 1000 periods). Even in 30-period phases with modest picture sets, subjects sometimes develop slang which uses parts of phrases ("drag" for "dragon"; "stat" for "statue"; see the .avi `movie' at our site at

http://ruffian.hss.caltech.edu/cx-docs/ for example). This observation is of more interest as pure psycholinguistics, but it may have some use in thinking about organizational culture too. Also, it is possible that in very long experiments people create mutated new codes out of sheer boredom or to demonstrate skill. If so, the idea that a code sticks hard as an equilibrium, because of its coordinating properties, is wrong.

"Where's the culture?":

Secret code: An interesting psycho-linguistic problem arises when one group of participants conspires to communicate with each other without being understood by hostile eavesdroppers.(Think of parents spelling out words—"C-A-N-D-Y"—they don't want their young children to comprehend.) The conspirators need to create "secret code" which is impenetrable to outsiders. Our experimental paradigm can easily be stretched to fit this situation if another player listens to the manager and employees, and chooses pictures, but the manager and employee *lose* money if the eavesdropper picks the right picture. Clark and ? (20??) studied this situation with student subjects trying to name campus locations, so that eavesdropping students could not guess the location.. He found that students tended to use personalized indefinite references which were likely to be opaque to eavesdroppers (as in, "Remember where you first met Jason?" [insert real example here]).

In organizations, secret codes are needed when, for example, workers want to hide the content of their discussion from management. Other examples come from between-firm communications or illegal activity. For example, in the HealthSouth accounting fraud, the people privy to the fraud were called "team players" and they held "team meetings". The gap between reported and actual earnings was called "the hole" which was filled with "dirt"—fraudulently-reported revenues. In **Random Family**, a remarkable story of two generations of Brooklyn Heights extended families, young Latino drug dealers fear—correctly—that they are being wiretapped. One says on the phone "What if they elley-say aggies-bay in the otty-spay?"—pig-Latin for "What if they sell baggies [heroin] in the spot [our "turf"]?" His accomplice responds "You know what I'm a do, too. I'm a open up ackie-jays" [tr: "I-ll open up jackies [tr: a crack outlet]"]. A challenge with secret code is that simply using unusual terminology signals that secrecy

is being used, which is a big tipoff to an eavesdropper. (What kinds of accountants talk about "dirt"?).

Machine learning: Computer scientists have been studying code development in distributed learning systems. Steels (2003) and Christiansen and Kirby (2003) discuss how shared communication systems emerge among robots equipped with a cognitive apparatus, a sensori-motor system, and a body, can arrive at shared grounded communication systems. Their work could both draw on how human systems evolve, and in turn inform optimal code creation.

Appendix B: Interesting code

"Will you drink the Kool-Aid?" (Jim Jones→ Microsoft)

Army: SNAFU (situation all f****d up), BOGSAT (bunch of guys sitting around talk), collateral damage (civilian casualties)

Alaskan pilots: VFR (visual flying rules), IFR (I follow river)

Diners and restaurants. : "paint a bow-wow red" (hotdog with ketchup);

"whiskey down…and it's walking" (rye toast, to go). An order to microwave something is conveyed by saying "give this some radar love" (an allusion to the beloved hit by Golden Earring). Busy restaurants are hectic, stressful places—perhaps having a whimsical term relieves tension and improves productivity (i.e., code is therapeutic).

Hip-hop slang: "bling" (jewelry), "biters" (imitators), "411" (information), "Are u feeling me?" (understanding)

References (highly incomplete)

Clark, Herb. Audience design

Christiansen, Morten H. and Simon Kirby. 2003. **Language Evolution: Consensus and Controversies**, Trends in Cognitive Sciences 7(7):300-307

Cremer, Jacques; Luis Garicano; and Andrea Prat (2003). Codes in organizations, working paper URL HERE

Heath, Chip.

Sorensen, Jasper.

Steels, Luc. 2003. **Evolving Grounded Communication for Robots**, Trends in Cognitive Sciences 7(7):308-312

Wernerfelt, Birger.

Camerer and Knez, "Coordination, organizational boundaries, and fads in business practice," Industrial and Corporate Change, in press.

Camerer and Knez, "Creating 'expectational assets' in the laboratory: 'Weakest-link' coordination games," Strategic Management Journal, 15, 1994, 101-119.