CHAPTER SEVEN

Behavioral Economics of Organizations

By Colin F. Camerer and Ulrike Malmendier [1]

7.1 Introduction

This essay analyzes how behavioral economics can be *applied* to organizations, and how it can be *enriched* by thinking about the economic questions associated with economic organization.

Behavioral economics modifies the standard economic model to account for psychophysical properties of preference and judgment, which create limits on rational calculation, willpower, and greed (see, for example, Rabin 1998; Mullainathan and Thaler 2001; Camerer and Loewenstein 2004). The modified economics theory aims at providing parsimonious and psychologically sound explanations for empirical findings that the standard model has a tough time explaining. From a methodological perspective, behavioral economics is simply a humble approach to economics, which respects the comparative empirical advantages of neighboring social sciences and sees neighboring sciences as trading partners. The empirical regularity and constructs carefully explored by those neighboring fields are presumed to be an important input which should often trump the seduction of mathematically elegant economic theories that are empirically unmotivated.

Thinking about organizations naturally extends the definition of behavioral economics to include how socialization, networks, and identity shape individual behavior in organizations (see, for example, Akerlof and Kranton 2005; Gibbons 2004b). As such, the behavioral economics of organizations involve both a special challenge and opportunity. Consider how behavioral *finance* has developed as an academic discipline, and how behavioral organizational economics *could* develop. Until 1990 or so, finance was arguably the area of economics most hostile to the idea that psychological limits matter for the focus of the field's attention—namely, stock price movements. Since then, there has been a dramatic shift in the amount of careful

[1] Ideas from the NBER Organizational Economics conference in March, 2004, particularly Bob Gibbons's presentation, have been useful, as well as discussions with Dirk Jenter, Chip Heath, Sendhil Mullainathan, Andreas Roider, Bengt Holmström's discussion in Helsinki, and Peter Diamond's close reading.

attention paid to behavioral ideas. This is surprising because it has been argued that large stock markets are the ultimate domain in which highly rational traders should limit the influence of those who make mistakes. So why did academic asset pricers start to "misbehave" so fast? One reason was good data, which made it easy to test new behavioral theories against the rational incumbents. Another reason was the availability of a clear benchmark model (market efficiency) to argue with. Good data and a sharp benchmark enabled researchers to create a set of clear anomalies in asset pricing. The situation was reversed in corporate finance. Rational arbitrageurs cannot easily limit suboptimal corporate-finance decisions, since it is hard to short-sell a CEO or CFO. Quite the opposite occurs: top executives may be entrenched and hard to get rid of. Biases and mistakes in decision-making are thus much more likely to have a persistent and large effect in corporate finance than on asset pricing. Nevertheless, behavioral research in corporate finance took off considerably later than that in asset pricing, possibly due to the lack of good data and a clear theoretical benchmark. Obtaining executive-level or even firm-level data on corporate finance decision-making has always been more intricate than obtaining stock price data. At the same time, numerous, conflicting theories have been proposed for virtually all types of corporate finance decisions, whether on mergers, stock issuance, or dividends.

The situation in organizational economics closely resembles that of corporate finance. The lack of rich within-firm panel data sets and cleanly identified exogenous variation is a persistent stumbling block, hampering progress in our understanding of the economics of organizations. On the other hand, if systematic biases are identifiable within the firm, they may have a large and persistent impact on organizational decision-making. If workers misallocate their human capital, nobody can take advantage of their mistakes by "short-selling" their capital. Adjustment to mistakes must come from some other source than simply trading against a mistake. Moreover, biases in organizations give rise to the interesting challenges of how organizations should be designed to repair these mistakes or to exploit them, or how firms organize around them if they represent genuine regret-free preferences rather than errors.[2] Moving away from mistakes due to heuristics, a lot of psychology is involved when workers team up in an organization—social comparison, changes in identity, camaraderie, attribution and diffusion of credit and blame, and so forth. This kind of psychology has played a small role in behavioral economics in recent years, but looms large when thinking about organizations.

Our paper is divided into four parts. Each part poses a broad question and suggests some ideas. Little systematic knowledge has been cumulated on many of these topics.

[2]Frank (2005) distinguishes mistakes you regret and those you do not. Properties of preference like social comparison fall into the latter category.

For those topics, the paper should be read as a research agenda rather than a review of what has been learned.

Section 7.2 lays out the basic single-activity risk-incentive conflict model and points out psychological considerations which complicate the model. Section 7.3 notes that the simplest risk-incentive model does not particularly account for the fact that people work together in organizations and discusses the importance of group loyalty, peer effects, and the coordinating role and cognitive economics of culture. Section 7.4 is about top management and governance and special considerations that arise like CEO overconfidence. Section 7.5 asks how patterns in individual judgment and choice aggregate into organizational outcomes when organizations can repair or exploit them. Section 7.6 concludes the chapter.

7.2 Complicating the Single-Agent Risk-Incentive Model

A good place to start is to discuss how psychology complicates the simple risk-incentive model of principal–agent relations that is a workhorse in organizational economics. First we will lay out a simple agency model with one type of activity. In the standard labor economics model, workers face a prevailing wage and decide how much labor to supply at that wage (and consume the remaining hours as leisure). We assume that people like money and, by definition, dislike work and like leisure.

A useful way to critique the standard principal–agent model is to ask when its basic assumptions are violated. The goal is not to heckle the model's shortcomings (which are an inevitable byproduct of simplicity), but to offer empirical facts and build up intuition about how it could be extended in useful ways.

We start with a simple exposition of a standard agency model. Worker i chooses effort e_i, which has cost $c(e_i)$. The productivity of effort also depends on a variable called skill, s_i, so that observed output is $x_i = f(e_i, s_i) + \theta_i$, where θ_i is a random component, additive to skill- and effort-based output ("luck"), with distribution $m(\theta_i)$. (In the simplest analysis, skill is homogeneous or does not matter, so $f(e_i, s_i) = e_i$.) Firms observe output x_i and pay a wage $w(x_i)$. This could be a fixed wage, $w(x_i) = w_i$; a step function or bonus package, $w(x_i) = \{w_i \text{ for } x_i < t_i; w_i + b_i \text{ for } x_i \geqslant t_i\}$; a linear wage, $w(x_i) = w_0 + \beta x_i$, etc. In subsequent sections we will assume that $w(x_i)$ is increasing in x_i. Assume also that preferences are separable in effort disutility and utility from wages. Then the agent's expected utility is

$$\mathrm{EU}(e_i) = \int_{\theta_i} u[w(f(e_i, s_i) + \theta_i)]m(\theta_i)\,\mathrm{d}\theta_i - c(e_i), \tag{7.1}$$

where the disutility of effort, $c(e_i)$, satisfies the usual assumptions, is increasing in e_i, and convex. The principal's earnings are $\pi(f(e_i, s_i)+\theta_i) - w(f(e_i, s_i)+\theta_i)$, where $\pi(\cdot)$ is the (gross) revenue of the firm.

Including skill makes this formalism cumbersome, but allows room for behavioral influences that are not traditionally considered. We discuss several ways in which the simple model above can be complicated:

(i) workers do not know the disutility of effort.
(ii) wage preferences depend on reference points (such as previous wages, or wages of others);
(iii) workers care about the procedure that generates wages or other outcomes;
(iv) psychic income matters, and may be tied to psychological factors like perceived appreciation;
(v) financial incentives may "crowd out" intrinsic incentives or affect performance negatively;
(vi) firms may be systematically biased in judging the cause of performance (i.e., disentangling worker effort e_i from luck θ_i).

7.2.1 *Workers Do Not Know the Disutility of Effort* $c(e)$

Labor economics extends a standard assumption from consumer theory—namely, that people have complete and consistent preferences across bundles of goods—to the case of labor and leisure. But young people may have only a vague idea of what work they would like to do when deciding on their first jobs, or on college majors (or even colleges), which partly determine their career paths due to irreversibilities and path-dependence.

A way to investigate the stability of labor-market preferences is to measure how much expressed preferences can be influenced by the way that work is described, or the procedure by which preferences are elicited (e.g., bidding wages or choosing at a fixed wage). For example, Ariely et al. (2004a) asked some subjects whether they would *pay* \$2 to attend a fifteen-minute poetry reading, and asked other subjects whether they would attend if they *were paid* \$2.[3] Later, a third of those who were anchored on paying said they would attend for free, compared to only 8% attendance by those who were anchored on being paid. Of course, students may not have developed clear preferences for whether listening to poetry is labor or leisure.[4] But if a random anchor can even influence the *sign* of $c(e)$, then it is likely that stronger

[3]Later experiments link the anchor to digits in a person's social security number and make it clear that the anchor was random, so that no information about $c(e)$ is conveyed by the choice of the anchor.

[4]The ambiguity about $c(e)$ is illustrated by a joke about a couple who are going away on a vacation. They ask their teenage neighbor, Mike, if he could take care of their dog while they are away. The couple explain that their dog needs to be fed, walked, and shown lots of loving care. They ask how much this job is worth to Mike. He ponders for a minute and concludes, "Well, I guess I'd pay you \$10." What's labor for one worker may be leisure for another.

influences affect at least some labor market entry decisions about relative values of $c(e)$.[5]

If employers know about anchoring and marketing influences on $c(e)$, then they will try to generate positive anchors, convincing prospective employees that working in their firm is fun. Little is known about the long-run influence or robustness of constructed-preference effects. Anchors might wear off: after experiencing an actual poetry reading, for example, subjects may quickly develop a consistent hedonic preference which is no longer affected by the initial anchor or subsequent ones, and resembles the complete preference assumed in economic theory. Furthermore, in competitive labor markets, anchors will not affect wages, because wages are determined by the marginal revenue product of workers. But anchoring could affect the *quantity* of labor supplied, even if it does not affect the "price" (i.e., the wage) for a given amount of labor supply. In the poetry example, think of the anchor fee/wage as creating a perception about the cost of effort of listening to poetry. Fixing the wage of poetry (through firm competition), more people will listen to poetry if they came to perceive the cost as low. So anchors can affect quantities even if they do not affect prices.

7.2.2 *Wage Preferences Can Depend on Reference Points* $u(w(x_i) - r)$

In the formalism above the utility of wages, $u[w(x_i)]$, does not depend on any special point of reference. But most systems in the brain have a homeostatic dependence on a set point or point of reference (e.g., hunger depends on what you have eaten recently; sweating and shivering respond to deviations of body temperature from a set point). If reactions to income tap similar psychophysical mechanisms, then people will care a lot about their wages relative to psychologically natural benchmarks, requiring a separate component of utility $u(w(x_i) - r)$ and a theory about what r is and how it changes (see, for example, Bowman et al. 1999; Koszegi and Rabin 2004).

Bewley (Chapter 5, this volume) discusses evidence that workers compare current earnings to previous earnings and dislike wage cuts, in traditional jobs where wages are adjusted periodically, so $r_t = w_{t-1}$. (Firms anticipate this and are reluctant to cut wages.) A similar reference dependence shows up in jobs where wages and hours fluctuate daily. For example, inexperienced New York cab drivers, who can adjust their daily hours, act as if they care about a daily income "target," which leads to

[5]For example, law school applications rose for a few years after the hit TV show "LA Law" began in 1986 (see Torry 1996). The show portrayed law as fun and empowering, and law firms as sexy, vibrant places to work. It is farfetched to think students are really Bayesian-updating based on new information about the $c(e)$ for law from a television show, because most pre-law students have a lot of information in the first place, or can gather objective information, and the television show is not meant to be an accurate depiction of the life of a lawyer.

labor supply elasticities that are *negative* (Camerer et al. 1997). Experienced drivers, however, have zero elasticity, which suggests a role for learning or attrition over time. Fehr and Götte (2003) find similar results in a field experiment with bicycle messengers. While randomly determined variation in monthly wages has a strong positive effect on the number of shifts per month, it has a (weaker) negative effect on the revenues within shift. Interestingly, the messengers for whom the nonstandard negative revenue effect is strongest also exert sizable loss aversion in a laboratory experiment.

Another kind of reference dependence is two-tier wage deals, which occur when firms are struggling financially. Senior workers doing the same jobs as entry-level workers are sometimes paid a larger wage, to avoid cutting their wages, while entry-level workers are paid less to save on the wage bill. Social comparison models predict that new workers will be unhappy at being paid less than senior workers beyond the disutility from the lower wage itself.

An important feature of reference dependence is that reference points may reflect failure to adjust for purchasing power. The best-studied example is money illusion. Firms act as if workers care about nominal wages rather than inflation-adjusted real wages, in making intertemporal comparisons (see, for example, Kahneman et al. 1986; Shafir et al. 1997). Baker et al. (1994) found there were hardly any nominal year-to-year wage cuts in a financial services firm, but many real wage cuts in inflationary years. The psychological principle behind money illusion may extend to illusions in comparing purchasing power across cities (leading people to prefer higher-salary jobs in the more expensive cities to lower-paying ones in cheaper cities) and in adjusting annual salaries for work hours (leading to taking high-salary jobs with the highest hours), but we do not know of any formal studies of these illusions.

7.2.3 *Workers Care about the Procedure That Generates Wages or Other Outcomes*

A simplifying principle in economic modeling is "consequentialism" or procedure-neutrality. People care only about outcomes and their economic impact, not about the *procedure* which produced those outcomes.[6]

One procedural preference is the effect of the *source* of income. The separability of income utility and effort disutility in Equation (7.1) implies that people value money equally if they earned it through hard work (effort) and if the money arrived as a windfall. But some experimental evidence suggests that money and goods which are

[6]Weaker statements of this principle are that preferences over procedures are weak compared with preferences over outcomes (such as compensation), or that procedural preferences are simply too poorly understood to serve as a basis for managerial policy.

earned are more valuable, or at least are treated differently. Coffee mugs which were "earned" are later sold for more than ones that are randomly allocated (Loewenstein and Issacharoff 1994). A brain imaging study showed that earned money produced a stronger activation in the nucleus accumbens, a brain area associated with predicted reward, than equivalent sums of unearned money (Zink et al. 2004).

These findings suggest that the utility of wages depends positively on effort, as if $U[w(x_i), e_i]$ had a positive cross-partial derivative,

$$\frac{\partial^2 U[w(x_i), e_i]}{\partial w(\cdot)\partial e_i} > 0.$$

When workers solve for optimal effort, extra utility from an "earned-income bonus" is like an increase in performance-based pay (a bonus return to effort), except that the incentive is entirely internal. How such a "pride bonus" interacts with sorting and with optimal wage and promotion policies is an interesting open question.

A broader point is that the process by which wages and other organizational outcomes are determined (particularly terminations) may affect how people value the consequences. In organizational studies this phenomenon is called "procedural justice," and it is argued that it is important empirically (Brockner and Wiesenfeld 1996; Tyler 2001). One component of procedural justice is the desire to have a voice or participate in important decisions that affect you. Another component is consistency of procedures. People dislike not knowing the rules that are being applied to judge them. The taste for clear rules motivates bright-line rules like seniority-based firing policies, and many policies in employment law. Of course, a taste for procedural justice is often hard to distinguish from standard preferences. For example, the desire to "have a voice" could just reflect the attempt to achieve individually favorable rules. Clear rules may simply help workers to optimize their effort allocation and reduce influence costs. Moreover, we may worry that procedural concerns, even if we could identify them, are second-order compared with wages. For example, while people may complain that a coworker was unfairly fired, will they actually quit their own job in protest, or accept a wage cut to get the coworker reinstated? Experiments suggest they might. But we would need more empirical research, tied to economic models, to calibrate the taste for procedural justice by pitting it against money in experiments or field analyses.

One kind of sensitivity to organizational procedures involves control. In a standard extension of the simple agency model above, firms observe a variable, γ_i, which is correlated with the unobserved luck component θ_i. The usual presumption is that effort e_i can be controlled by worker i but the error-correlate filtering variable γ_i cannot. Optimal contracting tells us that if workers dislike variance in wages more than firms do, the firms should use what they observe (γ_i) to filter some of the variance in θ_i and reduce the risk imposed on the worker. Workers will not mind

having their wages depend on a variable out of their control if it benefits them by reducing undesirable variance in adjusted wages. This changes if workers have preferences for fairness. Sometimes the value of γ_i will lead the firm to penalize a worker when true effort e_i was high (e.g., firing a successful CEO who was much less successful than industry peers). Workers may think this is unfair.

A common example is benchmarking performance to an outside standard, like industry profits. Risk-averse workers should *like* benchmarking because it buffers them against industry-wide shocks. But evidence for direct benchmarking in executive compensation is surprisingly scarce, even though it is not hard to create simple benchmarks which would reduce compensation variance a lot. It does not appear that wages are even indirectly very sensitive to benchmarks (Antle and Smith 1996; Salanié 2003, p. 470).[7]

7.2.4 *Psychic Income*

The basic risk-incentive model divides the worker's world into efforts they dislike, and rewards they like. It is convenient to talk about wages as rewards because they are easily measured, and do not satiate. But people are motivated by many others types of nonpecuniary "psychic income" as well. For example, Stern (2004) uses multiple offers for postdoc biologists to estimate that those who accept offers at more science-oriented firms earn 25% lower wages in order to be able to publish their research and engage in the scientific community. This is consistent with a "compensating differential," in which workers trade off money wages for other aspects of the job they like (or dislike).

The existence of compensating differentials is not controversial. Behavioral economics suggests, however, that the *source* of these differentials might include emotion, social comparison, and other forces not included in typical analyses which focus on safety and education.

Ariely et al. (2004b) use a simple experimental paradigm to calibrate psychic income for what they call the "meaning" of work. Students proofread pages for a declining wage schedule (the wage for the marginal page fell) and handed finished pages to an experimenter. When their work was "ignored" (simply placed on a stack) or immediately shredded, they finished only about six pages, compared with nine pages when their pages were signed and filed away. Thus, the subjects traded off marginal wages for "meaningful" work which was kept rather than shredded.

In discussing our paper at the conference, Holmström suggested that what workers fundamentally care about is that their efforts are *appreciated*. For example,

[7] One form of benchmarking is a tournament structure, which seems to account for some features of internal labor markets with clear career paths (see, for example, Main et al. 1993).

entrepreneurs often talk about their desire to get rich as a *byproduct* of making a good product that many people want to buy. They see riches as an expression of appreciation. The fact that tentative evidence suggests income is more (neurally) valuable if it is earned by one's own effort fits with this idea too.

Of course, agents may strive for appreciation of their work for more standard reasons. In a study of "superstar" CEOs, Malmendier and Tate (2005a) found that CEOs who had won awards (such as "CEO of the Year" or "Top Manager") got a big relative pay boost. This effect goes in the opposite direction of a trade-off between appreciation and pay—the award is appreciation *and* a raise in wages. Their findings may indicate a signaling role of awards, or the ability of CEOs to use awards to justify pay increases in the face of weak governance.

Moreover, any theory rooted in the desire for effort appreciation needs to take into account the fact that paying wages is one way of showing appreciation. But appreciation may also depend on relative wages, which creates a role for social comparison in signaling appreciation. Appreciation might also be cheaply generated by organizational activities such as awards, titles, executive visits, "employee of the month" plaques, etc. The desire for feeling appreciated might also be the underpinning of the desire for procedural justice—for being listened to, holding outcomes constant—that was discussed above. Note that, if employees value appreciation, organizations who figure out how to generate effort-appreciation most cheaply will save on the wage bill and will have a competitive advantage over organizations who express appreciation only through expensive wages.

7.2.5 *Financial Incentives May "Crowd Out" Intrinsic Incentives or Affect Performance Negatively*

Psychologists' synonym for psychic income is "intrinsic motivation"—the satisfaction a worker gets from work for its own sake. An interesting phenomenon documented in psychology is the possibility that extrinsic incentives like money can "crowd out" or extinguish intrinsic motivation. For example, Lepper et al. (1973) had children coloring pictures (with no extrinsic incentive). Then they paid the children a small sum of money for each picture colored. When the wage was later removed, the children colored fewer pictures than they had in the initial phase with no payments. Their interpretation, called "overjustification," is that the children infer something about their own intrinsic motivation from the wages they are paid; when they are paid for coloring, they infer that they must not like to color very much, so when the wage is removed they quit coloring pictures.

Kreps (1997) and Benabou and Tirole (2003) consider crowding-out from an economic point of view. Since incentives often do have positive effects, the challenge

is to create a theory which permits both positive effects and negative effects (i.e., "hidden costs" or "crowding out"). In Benabou and Tirole's theory, crowding out occurs when an agent thinks the principal knows more about job difficulty or the agent's abilities than she herself does. Then the agent takes the provision of high-powered incentives as a bad sign. Their theory is much like the original idea in social psychology but more carefully specified and rich with empirical implications (chiefly, that hidden costs should only occur when workers think incentives signal bad news). Benabou and Tirole (2004) discuss the case where incentives interact with prosocial behavior.

Leaving aside crowding out, monetary incentives could also affect performance negatively if high incentives create arousal that inhibits automatic responses or creates distraction (e.g., "choking under pressure"). According to the "Yerkes–Dodson Law" (Yerkes and Dodson 1908), arousal increases performance up to a point at which further arousal degrades performance. Ariely et al. (2003) demonstrate this effect with experiments conducted in India. They allowed their subjects to earn up to six months of wages[8] in various simple games. Surprisingly, the subjects with the highest payment perform worse than those with the moderate and the low payment in almost all games. It is, of course, unclear how relevant this effect is for the economics of organizations, given actual wage-setting and reward policies. Even if the effect were prevalent, wage-setters may account for it and adjust wages appropriately. Moreover, if experience with high-stakes situations decreases arousal, the effect may only apply to inexperienced agents.

7.2.6 *Firms May Be Biased in Judging the Cause of Performance*

In the standard agency model laid out above, agents exert effort which is productive, but principals can only observe agents' output. Observed output depends on unobserved effort and unobserved "luck." The usual assumption is that principals know the relative importance of effort and error in determining output, and also anticipate the agents' reaction to knowing their effort is not clearly observed.

This simple assumption buries a lot of behavioral economics. Separating skill from luck is difficult. As a result, it is quite possible that psychological forces bias principals' judgments of whether output is due to skill or luck. Three important psychological forces suggest that separating hard work from luck will often be biased in predictable ways: hindsight bias; overattributing cause to personal traits rather than situational influences; and overconfidence.

[8]The maximum pay amounted to half of the average annual consumer expenditure (per capita) in India.

7.2.6.1 Hindsight Bias

The brain's tremendous ability to sense patterns is manifested by our tendency to rapidly rewrite our memories of the past to fit what we have learned. The problem is that revising our beliefs after an event naturally leads to misremembering how little we knew before the event. Rapid rewriting creates "hindsight bias"—the ex post recollection of the ex ante probability of an event will be biased in the direction of the event's realization. Hindsight bias is on display every day in sports commentary and news coverage. It is surely an important force in organizational life but has not been studied at all.

A simple way to think about hindsight bias in agency models is that the effort an agent exerts narrows down the distribution of some relevant variable v at time t, $\sigma_t(v)$. Suppose the agent's job is to make the best decision given the value of v, which will eventually have a public realization v^*. A principal who creates an ex post "recollection" of the ex ante variance based on the realized v^*, $\hat{E}_t[\sigma_t(v)] = E_{t+1}[\sigma_t(v) \mid v^*]$, will infer too much mass around v^*. Suppose the agent did not choose the optimal action conditional on v^*. Then the principal will (mistakenly) infer that the agent did not exert enough effort to narrow the variance $\sigma_t(v)$ (he "should have known"). If the optimal action conditional on v^* is what the agent actually chose, then the agent gets too little credit because the principal will recall the likelihood of v^* occurring as being obvious (i.e., $\hat{E}_t[\sigma_t(v)]$ will be too closely centered around v^*, and the principal will not be impressed that the agent figured out that v^* would happen).

This sketch of a model becomes more interesting if agents *anticipate* that principals will be hindsight-biased. Agents who anticipate being second-guessed will record as much ex ante evidence as they can (also known as "covering your ass"). Agents may also herd into the same decisions other agents make (see Zwiebel 1995) or seek other ways to avoid anticipated hindsight bias.

It would be very useful to have field evidence distinguishing the effects of hindsight bias from other distortions. One method is to compare behavior of owner-managed firms with similar firms where shareholders and boards can second-guess the managers. Hindsight bias should be lower and risk-taking higher among owner-managed firms.

7.2.6.2 Misattribution of Cause

Social psychologists studying "attribution theory" have found that a typical mistake in attributing cause is to credit and blame individuals more than is deserved, compared to blaming and crediting situational variables.[9] In model terms, this

[9]Interestingly, recent studies by Nisbett (2004) suggest that the "fundamental error" of overattributing cause to people, rather than situations, is distinctly Western and is reversed in Asian cultures. Applied to

means inferring more than is justified about the unobserved effort e_i from output $f(e_i, s_i) + \theta_i$ or, alternatively, underestimating the importance of luck ($\text{var}(\theta_i)$). In the context of the firm, attribution is likely to be especially important in evaluating top executives, who make far-reaching decisions with noisy and lagged feedback. They are likely to be both credited and blamed too much for forces beyond their control. Applied to a person's own judgement of his or her efforts, however, attributions tend to be self-serving. For example, Bettman and Weitz (1983) find that, when companies do well, corporate annual reports attribute the success to internal factors and skill. When companies do badly, the reports attribute blame to external factors (regulation, demand shocks, and so forth). Those findings lack, of course, the ultimate proof of distortion. Moreover, even evidently distorted information provision may simply be the rational response to investors' lack of information or to investor credulity. Bertrand and Mullainathan (2001) show that exogenous variation in profits is reflected in executives' compensation. They find that compensation of oil company executives rises with worldwide oil prices, even though oil prices are set by global demand and OPEC output and are not influenced much by one executive's action. At the same time, executives are largely shielded from blame (in the form of lower compensation) when oil prices fall. Indeed, the phenomenon of "negative shielding"—relative insensitivity of compensation to poor performance, compared to stronger sensitivity to good performance—is widely documented in executive compensation (Garvey and Milbourn, forthcoming). It remains unclear from these studies which agents—if any—are misattributing positive outcomes. One possibility is that boards mistakenly give credit for positive outcomes and CEOs rationally exploit or even induce the misattribution of board members. Alternatively, executives themselves attribute positive outcomes to their own skills or effort and bad outcomes to bad luck. If corporate governance is weak, such CEOs may then extract disproportionately higher compensation after exogenously higher profits. Finally, it is possible that no party inside the corporation exhibits misattribution error. CEOs in badly governed firms may overadjust their compensation to good outcomes and underadjust for bad outcomes, exploiting investors' naivety, lack of information, or misattribution of executive credit to skill, and blame to outside forces.

7.2.6.3 Overconfidence

Many studies show that people are often overconfident. We define overconfidence as the tendency to overestimate one's own (relative) abilities and resulting outcomes and overoptimism as overestimation of general prospects. Here we briefly mention

agency theory, this suggests that pay-for-performance will be more slowly adopted in Asian-controlled firms and the East–West difference may have something to do with CEO status and pay in the different hemispheres.

overconfidence in the basic agency model. Section 7.4, below, has more discussion of overconfidence of top managers.

What does overconfidence of workers mean in the agency model? There are many ways to model overconfidence and it is an open question what these different types of overconfidence imply in equilibrium, or in practice. One possibility is that people overestimate the output x_i they will generate. For example, they may assume a distribution $\tilde{m}(\theta_i)$ that first-order stochastically dominates the true distribution $m(\theta_i)$: observationally equivalent to overestimating luck. Or they may overestimate the marginal productivity of their effort, $\partial f(e_i, s_i)/\partial e_i$. In either case they may work less hard than if they were not overconfident. Assume for example, that their wage schedule includes a piece rate β. At the correct optimum effort with no overconfidence (and corresponding optimal piece rate), agents misperceive themselves as working too hard because they overestimate expected output and, consequently, underestimate marginal utility. As a result, overconfident agents may cut back on effort. In theory, firms will increase the piece rate β to induce more effort, but for reasonable specifications this compensatory effect will not be strong enough to push wages and effort back to the no-overconfidence point.

A different possibility is that overconfident workers think they are more skilled than they really are (i.e., they overestimate their value of s_i). Whether overconfident workers of this sort work too little or too much will depend on whether effort and skill are complementary, and on the types of contracts firms can offer. For example, if e_i and s_i are complements, the firm can offer a contract that pays a large bonus if a very high output level is achieved (similar to backloaded option packages which are common in biotechnology firms, or highly leveraged venture capital deals offered to founders hungry for capital). Workers who overestimate their skill s_i will perceive the marginal return to effort as higher than it actually is and will work hard to grasp the brass ring. If this is true, we would expect to see extraordinary bonuses being offered by firms, to exploit overconfidence, accompanied by high actual failure rates.

The opposite of overconfidence might be important in worker motivation as well. Psychiatrists discuss an "impostor syndrome" in which workers discount their own success, or attribute it to luck, and feel that others have been fooled into thinking they are talented (see, for example, Kolligian and Sternberg 1991). Self-perceived "impostors" may work harder than others if they think that they have talent but need to work even harder to achieve true success.

More generally, there are important caveats to applying the evidence of widespread overconfidence to organizations. Rather than assuming that agents uniformly overestimate their abilities, researchers need to account for heterogeneity and attempt to evaluate ex ante indicators of overconfidence. For example, some studies show that women are less overconfident than men, which might be important for personnel selection. Seventy- to eighty-year-olds are also less overconfident

than younger college students on general knowledge questions, which suggests that age teaches people what they know (see Kovalchik et al. 2005). And depressed people are typically *not* overconfident ("wiser but sadder"), which suggests overconfidence is part of mental health, and probably evolutionarily adaptive (see, for example, Taylor and Brown 1988). Overconfidence depends also on the type of task and its framing. For example, overconfidence about a particular estimate shrivels up when a portfolio of estimates are considered (see Kahneman and Lovallo (1993) on the single-case "inside view" compared to the portfolio "outside view"). Consider the 90%-confidence interval task, where subjects are instructed to give a wide interval for each numerical quantity which is only wrong 10% of the time; typically, five out of ten intervals are too narrow. Subjects typically give too narrow intervals. But when subjects are asked how many out of ten questions they got right, they typically do not say "one out of ten," they say "five out of ten"—which is empirically accurate (Sniezek and Buckley 1992). Overconfidence also shrivels up when the question is posed more narrowly, breaking an outcome into components (Dunning et al. 1989). While drivers are overconfident about overall driving ability, they are less overconfident about specific components like driving in snow, avoiding fender-bender accidents in heavy traffic, and so on. This fact suggests that if the criteria used to judge managerial success become more vague as people rise in a firm's hierarchy, overestimation of future success (and career outcomes) will rise too.

7.3 Workers as Members of Multi-Agent Firms

The basic model above focuses on a single principal–agent relation. It makes little use of the special properties of large organizations with many principals, agents, and hierarchies. Scaling the models up to large firms has a variety of implications for social comparison, group and peer effects, emotional and cognitive firm boundaries, and the coordinating role of culture and leadership in large groups.

7.3.1 *Social Comparison*

In the worker expected utility in (7.1), only the worker's own effort and wage enter into her utility. But in real organizations, people make friends and enemies, and compare themselves to others. Workers may thus sacrifice their own earnings to help their friends and harm their enemies, or to create better social comparisons. This topic, and its importance for contracting and labor economics, is discussed in some detail in Fehr and Gächter (2000).

The influence of social comparison can lead to conflict when different agents choose different comparison sets. In a study of actual teacher strikes in Pennsylvania, it was found (Babcock and Loewenstein 1997) that teachers' unions tend

to say that nearby school systems with high wages are most comparable to their own, and school boards say that nearby systems with low wages are most comparable. The gap between the two sides' comparisons is correlated with strike incidence. Experiments also show that the gap between perceptions of fairness predicts strike incidence (Babcock and Loewenstein 1997). Importantly, these perceptions are not merely strategic posturing. The gap disappears if subjects are assigned to a bargaining role *after* reading facts about the case. This implies that the self-serving perceptions of fairness arise from differences in encoding what facts are important.

One might think that social comparison is an inferior good—as people are paid more and more, concern for social comparison becomes relatively less important. But social comparison also appears to enter organizational economics even in the very high compensation of executives. Hall and Murphy (2003) note the effect of the U.S. Securities and Exchange Commission (SEC) regulatory changes in 1992 on CEO pay which required "company compensation committees to describe company pay practices and justify in detail how they determined the pay of the chief executive officer...." Presumably the SEC thought that by requiring clearer disclosure of what CEOs earned, and how earnings were linked to past performance, embarrassment would rein in high payouts unjustified by performance, and might link compensation more closely to performance. But in a world where executives hate being paid less than a peer they think is less deserving, and their perceptions are self-servingly biased, disclosing wages leads to a ratchet effect: information which tells CEOs they are being paid less than others can drive all wages up.[10] Similarly, Belliveau et al. (1996) found that CEO pay was more closely correlated with earnings of members of the board's compensation committee than to pay of other board members. One explanation is that compensation committee members compare the CEO's pay with their own.

Social comparison may also affect the boundaries of the firm. If people naturally compare their own wages to those who work in the same firm, then one of the economic effects of a merger is a shift in the comparison sets of workers in the newly merged company. Some workers will now be underpaid relative to their new colleagues. This creates turnover or influence costs, and those costs should be

[10]Hall and Murphy (2003) wrote: "Although the new rules made option grants significantly more transparent—which would lead companies to grant fewer options to the extent that options are used to 'hide' compensation—we believe that the rules on net encouraged *increased* grants" (our italics). They note that the growth in option pay was already underway in the 1980s and conclude, "the new rule constituted an implicit government 'blessing' of stock options as appropriate performance-based pay, and this implicit endorsement may have further fueled the escalation of options" (Hall and Murphy 2003, p. 62). The subsequent growth of option-based packages is also consistent with the idea that self-serving comparison fueled growth in options grants.

accounted for as part of the merger cost. For example, when General Electric (GE) bought the television network NBC, the high salaries of NBC employees angered many successful engineers, who suddenly had an immediate reason to compare their wages with the highly paid "new neighbors." Kole and Lehn (2000) make the same point about the takeover of Piedmont Aviation by U.S. Airways. The choice of the boundary of the firm, therefore, is also a choice of social comparison sets, which has economic consequences of its own.

7.3.2 *Peers and Groups*

One place in which social preferences have entered labor economics is in the concept of "peer pressure." Peer pressure is often invoked to explain why group-based incentives are surprisingly successful despite the incentive to free ride. For example, Knez and Simester (2001) note that a $65/month bonus for all Continental Airlines employees, for improved on-time-arrival rates, actually worked, even though any one employee's performance has a minuscule effect on the firm's overall on-time performance. They note the important influence of immediate pressure by peers to improve performance (see also Kandel and Lazear 1992). The experimental literature on public goods games with sanctions also shows the power of peer pressure (see, for example, Yamagishi 1988; Fehr and Gächter 2000). When subjects can punish others at a cost to themselves, they often do so, which drives contributions up. Note that not every subject punishes, but enough do so to create higher contributions.

Falk and Ichino (2003) measured peer effects on output in an experiment, exploiting the ability to assign workers to pairs randomly, which avoids the problem of endogenous peer selection that is typical in most (though not all) field data sets. They recruited subjects to stuff letters into envelopes for four hours for a fixed fee (with no performance bonus). Some subjects did the activity by themselves; others did it in the presence of a randomly assigned peer. Random pairing led to closely matched output. The absolute deviation of output within each actual peer group averaged 20.6—much smaller than 99% of simulated absolute deviations within artificial pairs composed of solo subjects. Pairing workers also raised output on average (about 10%). Pairs were more likely to escalate each others' output than to tacitly collude by goofing off. The estimated effect of a one-unit increase in output on the peer's output was 0.14 units, close to the 0.14–0.18 estimated from field data in an Italian firm (Ichino and Maggi 2000).

7.3.3 *Group Identification*

There is a substantial psychological literature on group identification which could be useful for organizational economics (see Akerlof and Kranton (2005), on the

economics of identity). Psychological evidence from "minimal group" experiments suggests that people form group affiliations rapidly, e.g., based on which room they were placed into, or according to their preference for Klee or Kandinsky paintings. Furthermore, these affiliations influence allocations in dictator and public-good-type games—they give more to in-group members and less to out-group members.[11]

Extrapolated to organizations, these findings suggest that workers will help members of their own group—*as they perceive their group and its membership.* In this view, a group is a club of people who all agree to help each other. This introduces an interesting role for management of group boundaries.

To fix ideas, suppose there is a component of effort which is productive for the firm but does not increase an employee's own measured output. Call this "helping" and think of it as a contribution to the firm-wide public good. In order to focus attention on group identification rather than multitasking (Holmström and Milgrom 1991), assume that choosing to help does not cut too deeply into a worker's other productive activities.

In the context of the organization, note first that the firm would most like its own workers to help each other, but to not help competitors' workers. But group affiliations that determine helping probably do not map neatly onto legal boundaries. Friendship, ethnicity and language, geography and architecture (which workers live or work nearby), marital status (who hang out after work together at happy hour, or during their children's play dates), and other distinctions create an archipelago of group affiliations within a firm, and across firms.

To some extent, firms can manage group creation through sorting (selecting "our kind of person" in hiring) and other organizational practices. For example, if helping customers *too much* is bad for the firm, it may use job rotation to keep its workers from becoming too helpful with favored customers. The sociology literature, however, provides numerous examples of firms misestimating group effects, such as Roy's (1952) famous study of a small-parts machine shop in Chicago. In this shop, the management designed task-specific piece rates such that for all the different jobs a certain monetary target could be reached with the same "effort." Roy's study demonstrates that workers were able to mislead management and induce piece rates, for which meeting the target was very easy. The workers imposed strong sanctions on anybody attempting to deviate and to exert more effort.

Evolutionary psychology provides an interesting way to speculate about groups and firms. Psychologists think of some kinds of groups as "essential," which means

[11] Group dynamics can also induce the opposite bias, a "black-sheep effect." Unpopular in-group members receive higher punishment than unpopular out-group members (Marques et al. 1988), especially if peer pressure is high. Another important caveat is its susceptibility to the experimental framing (Hertel and Kerr 2001).

that group members share a common essence which is immutable.[12] Gil-White (2001, p. 519), for example, argues that ethnic groups are perceived as a species. They have

> (1) a distinctive morphology... (2) many intercorrelated hidden properties... (3) within-category mating and descent-based membership, and (4) saliently labeled categories in cultural discourse.

If a firm can create a sense among workers that they are like a species—with immutable special properties that are inheritable—it might be able to hijack neural circuitry which is highly evolved to distinguish friend from foe on the basis of species and species-like ethnic characteristics, to create a deep sense of group membership and helping. For example, a body feature which cannot be changed marks a species, so one way to create a species illusion is with tattoos. In fact, Nike founder Phil Knight and many of his employees have tattoos of the Nike "swoosh" logo on their left calves as a sign of camaraderie.

Group loyalty can persist even when workers leave the firm. For example, it is typical for lawyers and consultants to eventually leave their firms and work for one of the firms' clients (whom they worked with originally). While there is a clear difference in the employee's decision rights after switching firms, there may be little difference in emotional attachment to their old group (the law firm) and their new group (the client firm). Many consulting firms keep active lists of "alumni" who have left the firm to work for their clients, and presumably use the residue of group loyalty to keep those clients and to resolve disputes.

7.3.4 *Coordination, Culture, and Leadership*

Organizational relationships are typically repeated games. As Gibbons (2004a) notes, the folk theorem gives us an embarrassment of riches about how repeated games might be played—there are *too many* equilibria. As Schelling (1960) wrote, predicting what players will do in games with many equilibria by pure mathematics is like trying to prove that a joke is funny without telling it. When there are many equilibria, social considerations like historical traditions and norms, and credibility of leaders who make announcements intended to focus attention on good equilibria, will make a difference.

---

[12] For example, the Aryan Brotherhood gang has a saying "blood in, blood out." This means you must shed someone's blood to gain admission to the group, and can only leave if your own blood is shed. U.S. Marines have a saying, "Once a Marine, always a Marine." Groups like military and police often have de facto inheritance, in the sense that special pride is placed on doing the same job as your father, grandfather, etc., and favoritism may be granted to "descendants" in hiring and promotion.

These issues are highlighted by putting incentive conflicts aside and focusing on organizations as "teams." A team has a common interest, but has to coordinate activity and information to make a commonly beneficial choice (see Radner and Van Zandt 2001). Coordination is often neglected in organizational economics, perhaps because of a belief that it is easy to achieve. Indeed, in small groups a quick e-mail may coordinate activity. But when organizations face time pressure, have rapid turnover, and scale up in size, coordination is challenging.

Experimental studies of coordination games have shown that even in very simple settings, with modest group sizes (2–16) and very clear games, coordination failure is common. Several studies have used a "weak-link" game[13]: players simultaneously choose effort levels (integers from 1–7) and earn an amount which increases in the *minimum* effort, as well as a penalty if their effort is above the minimum. The game is tricky because there is one incentive to choose high numbers, to raise the minimum, and another incentive to choose low numbers, to avoid the penalty from wasted effort. Every effort level is a different equilibrium, but choosing the highest effort is Pareto-dominant.

Many studies show that coordination failure in groups playing simultaneously without communication: they generally fail to reach the Pareto-dominant equilibrium (see, for example, Van Huyck et al. 1990; Camerer 2003, Chapter 7). Efficiency can be achieved when two people play together repeatedly, when groups "grow" slowly, starting with two and adding one player at a time (Weber 2006), or when a small bonus is publicly announced if everyone reaches the efficient choices (Knez and Camerer 1994; Brandts and Cooper 2004). Surprisingly, limited communication does not help much (cf. Cooper et al. 1994). When one "leader" is randomly chosen to talk briefly to the group (usually exhorting everyone to choose "7"), large groups still do not reach efficiency and subjects misattribute credit and blame to the leader.

These studies illustrate how coordination games could be used to study facets such as growth of firms, the attention-getting effect of bonuses, and leader credibility in very simple experimental organizations. They also suggest that efficient coordination is usually not easy, even in simple games with clear strategies and payoffs, so that coordination is likely to be important for success of large organizations, and perhaps underestimated in importance. For example, Heath and Staudenmayer (2000) had experimental subjects divide labor to create parts of a "man" from LEGO blocks; they were scored on how good their composite man looked when assembled from the

[13]Weak-link games are n-person "stag hunt games." In organizations, they are models of Leontief production processes that are sensitive to the lowest quality input, like an assembly line that is delayed by the slowest worker, or a recipe which is ruined if it has one bad ingredient (see also Kremer's (1993) "O-ring" theory of production).

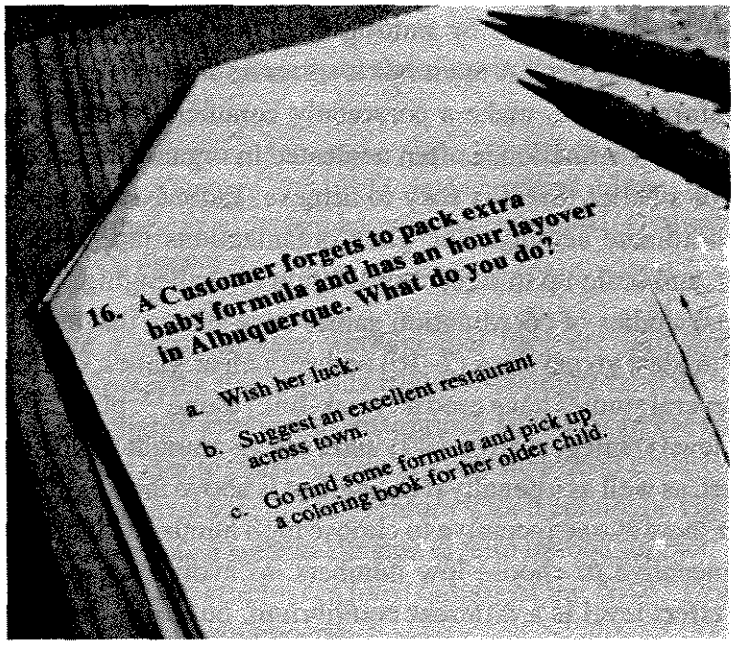

Figure 7.1. A Krepsian quiz at Southwest Airlines.

parts. Appointing a coordinator who made sure the parts being created fit together improved output, but few of the teams spontaneously did so.[14]

7.3.4.1 Culture as Focal Principles and Correlating Devices

Kreps (1990) introduced the idea of an organization's culture as a correlating device (see also Camerer and Vepsäläinen 1988). In this view, cultural rules are meta-principles which tell workers what to do in the face of unforeseen contingencies, in the spirit of Schelling's (1960) focal points in matching games, but extended to meta-principles which prescribe focal points in an ensemble of games.

The idea is illustrated by Figure 7.1, an advertisement from a Southwest Airlines in-flight magazine. This shows a hypothetical quiz given to potential employees about how to react to a customer's problem. A company's culture can be thought of as a set of general recipes which help workers and managers coordinate on the "right" answer in quizzes like this, which correspond to workplace incidents. At Southwest Airlines, the correct answer is get some baby formula for the delayed passenger and pick up a coloring book as well. The example gains force from the fact that it is based on a true story. The point of the ad is that Southwest Airlines is deeply customer-oriented (note that the term "Customer" is capitalized).

[14]Large rowing crews use a coxswain—a person who does not row, but sits at the front of the boat and synchronizes timing (and motivates the rowers). The coxswain is a factor of production who helps produce coordination, and whose value presumably outweighs his or her extra body weight (which is typically low).

The idea that culture is a set of principles for generating focal points in specific microgames played by workers and managers is not so different from the way organization theories describe culture. For example, Schein (1983) defines culture as "the *basic assumptions* and *beliefs* that are shared by members of an organization." He later notes that these "assumptions and beliefs are *learned* responses to a group's problems of *survival* in its external environment and its problems of *internal integration.*"[15] Shared understanding reduces conflict and, if the culture is visible to outsiders, also encourages efficient assortative matching between workers' attitudes and the attitudes a firm requires.

As important as culture appears to be, it has not been studied very carefully empirically. The biggest obstacle is the lack of empirical measures and exogenous variation. One stream of literature uses psychometric scales to map organizations onto some two-dimensional or multidimensional space (Cameron and Quinn 1999). These studies are not statistically sophisticated and typically treat the whole corporation as one culture, ignoring subcultures. Other studies focus on a dimension of culture, like "strength," and explore its relation to profitability or profit variation (see Kotter and Heskett 1992; Sørensen 2002). Still other studies establish the existence of distinct workplace culture (called "climate" by organizational researchers) and study its consequences for phenomena like turnover (Bartel et al. 2003; Rebitzer 2002).

The lack of a sharp formal framework and clean empirical identification has certainly held back organizational economists interested in culture but progress is being made (see, for example, Hermalin 2001). There is a small emerging literature on "codes"—the language workers use to describe events. Codes have efficiency properties, depend on shared understanding, and can be difficult to change and imitate, like other facets of culture (e.g., corporate values). Crémer et al. (2003) assume that codes are coarse partitions of a large space of events. Coarse partitions are adequate because workers need not communicate exactly which event has occurred. Wernerfelt (2004) offers a similar model.

Codes have been studied experimentally by using a picture-naming paradigm. In these experiments, agents see pictures on a screen. A "manager" must figure out a code to identify predetermined target pictures to workers. Payoffs depend on how rapidly the target pictures can be named without errors. Codes develop rapidly: subjects typically use about 120 characters to describe four picture at first (that is the length of this sentence so far), but use only about five characters per picture after 20–40 practice trials.

[15]A useful definition was mentioned by John Roberts in a December 2004 conference at MIT. He diagnoses culture by asking people, "What would a new employee need to know to function effectively in your company?"

One study shows that it is better to have one manager creating code than to rotate code-creation among the subjects (Rick et al. 2004). This suggests that having a single "founder" creating code is better—specialization is better than taking the best code ideas from different workers. However, the single-manager groups have a harder time integrating new employees; the managers are more likely to give up on teaching the new employee, who then makes a lot of costly picture-naming errors.

Two other studies explore how conflict is created when groups "merge." Weber and Camerer (2003) found that when two groups naming the same pictures merged (like a horizontal merger of similar businesses), subjects were too optimistic about how rapidly code could be created in the merged group. Feiler and Camerer (2006) replicated this result with different groups of pictures (like a diversifying merger) and where subjects were endogenously sorted into the merger group through bids for wage premia to join the merged firm. Subjects' guesses about the speed of the merger group and bids in—theoretically incentive-compatible—second-price auctions were too optimistic relative to actual earnings differentials of merged and unmerged groups. It is easy to think of other applications of this sort of code-creation paradigm, as well as expanding its scope to include development of corporate values like reciprocity and camaraderie, expressed by cooperativeness in public goods games.

7.3.4.2 The Cognitive Economics of Organizational Knowledge

Organizations are complicated. Knowledge about how to behave in complicated systems—like the world, or an organization—is often encoded in a form that is very different than the state-space partitions and Bayesian probability that dominate economic theory, and is much simpler and evolutionarily advanced. This suggests that understanding cognitive economy—what kinds of information fits the brain's natural structures for storing, remembering, and reliably repeating information—is important for understanding organizations.

Organization theorists have argued that a narrative or story is a common way of encoding information about what to do in an organization (see, for example, Martin 1982; Heath and Seidel 2004). If the punch line of the story is success or failure, a story is basically an instruction about how to behave. Narratives may be a superior encoding form from the point of view of human memory—if they are better remembered, they will not degrade as much as other kinds of information when passed down from old workers to new ones. If it is fun to tell stories—they often involve humor—they may spread more rapidly and reliably through a firm than boring instructions.[16]

[16]Memory appears to be encoded into episodic memory (time-stamped events) and semantic memory

Another element of cognitive economy in organizations is "personification." Much of human perception and judgment is organized around what we know about people and how we summarize our knowledge in the form of traits and thumbnail sketches of people.[17] It is cognitively natural to identify an organization with a charismatic founder or iconic leader, such as Walt Disney, Bill Gates, and Herb Kelleher from Southwest Airlines.

If a leader's behavior and values are thought to exemplify those of the culture, then it may be easier to answer the question, "What does my implicit employment contract require me to do?" by answering the substitute question, "What would Walt do?" (as employees at Disney often did, long after Walt's death). The personality of a corporate icon then becomes an asset that economizes on transaction costs by telling people what to do, in a way their brains can remember and understand more easily than a boring list of rules.

7.3.4.3 Leadership, Simplicity, and Mission Statements

Thinking about coordination and cognitive economy may help explain the nature of leadership and language in companies. A striking feature of corporate talk is that it is highly simplified, but catchy. "Maximize profit" is *not* an inspiring mission statement. Instead, companies have slogans like "Quality is job 1" (Ford), "Just do it" (Nike) "Do no evil" (Google), "Copy EXACTLY!" (Intel, to maximize transfer of wafer-production processes from old plants to new ones) and "Flatter the novice; reward the expert driver" (Mustang cars). The Ritz–Carlton hotel chain supposedly has a rule that up to $2,000 can be spent to solve a guest's problem. The figure $2,000 is useful for two reasons: it is a round number that's memorable; and it is a *surprisingly large* number, a "put your money where your mouth is" signal of the Ritz's commitment to luxury hotel quality.

Why are these stated goals so simple? One explanation is that meta-principles which help workers coordinate on equilibrium workplace norms have to be applied to a broad range of different microgames. The principles have to be made commonly known, and repeated over many years and between many workers without linguistic degradation. Heath and Seidel (2004) suggest some good properties of "language

(facts about the world and a distillation of episodic memories). Corporate narratives harness the power of episodic memory. Narratives often involve drama, humor, and emotion as well, which provides added "depth of processing" that enhances their memorability. From a neural point of view these may provide a kind of engineering redundancy because memory is then supported by multiple interconnected systems: Who did it? How did it feel? When did it happen? From the brain's point of view, a story may therefore be simultaneously stored in different locations which are associated, improving retrieval.

[17]Facial recognition is a very special kind of object recognition with dedicated neural circuitry (the facial fusiform area, FFA), which suggests that there may be some cognitive advantages to organizing cultural rules around behavior of a "face of a corporation."

memes": they should be simple, memorable, broadly applicable, and evoke complex emotions. Their hypotheses could easily be tested in experiments on how well various linguistic meta-principles coordinate actions in the face of garbling due to labor turnover and the passage of time.[18]

Van Den Steen (2005) tells a story about leadership which is closer to the meta-principle theory of culture. In his model, managers with "vision" have a clear preference for project types.[19] A strong vision is good because it enables efficient labor-market sorting of other managers who agree with the vision into the firm, and saves time that could be wasted by proposing projects that do not fit.

7.4 Top Managers and Corporate Finance

Behavioral considerations contribute to our understanding of corporate finance from two perspectives: first, recognizing that markets make mistakes which managers exploit; and, second, that managers make mistakes which markets do not fully correct. In this section we assume the perspective most relevant to the economics of organizations and consider the possibility that top managers make mistakes.[20] Managers are most likely to be affected in their judgment when they face decisions in which three conditions hold:

(i) decisions are not frequent and do not deliver clear feedback;
(ii) the manager does not specialize in making those decisions; and
(iii) managers are protected from market pressure and competition.

As for (i), the higher the level of a manager in the hierarchy, the less likely decisions are to be repeated enough to learn from feedback. Large-scale investments, merger agreements, or capital restructuring are relatively rare decisions, and the organizational framework of a company may not be optimized to facilitate them. It is also hard to evaluate outcomes. For example, there is no agreed-upon methodology on how to evaluate the long-run returns to mergers, even in the academic literature (Andrade et al. 2001). Regarding (ii), the level of specialization is likely to decrease the higher up in the hierarchy decisions are made. As Lazear (2004) formulates in the

[18]Experiments on slogan quality could use paradigms like the game "telephone." In telephone, people form a circle and speak a phrase very rapidly into the ear of the person to their right, who repeats it rapidly to the person on *their* right, and so on. The goal is for a slogan or cultural principle to make it around the circle with the least degradation. One could also test the durability of shared memory of a slogan by announcing it to a group, then asking group members at future times whether they remember it, and whether they think others remember it. Good cultural principles will be remembered, and expected to be remembered by others, over long spans of time.

[19]This is a stylized version of corporate vision as "a mental image of a possible and desirable future state of the organization" (Bennis and Nanus 1985).

[20]For a survey that covers the first perspective, see Baker et al. (2004).

context of entrepreneurship, the top decision maker in a firm has to be (and typically is) a jack-of-all-trades rather than a specialist. And regarding (iii), we know from many studies of corporate governance and entrenchment that CEOs often manage to protect themselves from market pressure (see, for example, Bebchuk et al. 2002). Sheltered from the discipline of the outside market, they may make significantly less efficient corporate decisions, as Bertrand and Mullainathan (2003) show, exploiting variations in anti-takeover laws.

To push the argument for biased decision-making in corporations even further, the market not only may be unable to "short-sell biased CEOs" and to thereby remedy distorted corporate decision-making, but may actually *value* managerial overconfidence. As Goel and Thakor (2000) and Gervais et al. (2003) point out, diversified shareholders, who would like managers to make risk-neutral choices and to maximize the expected value of their firms, may prefer managers to be overconfident if it balances managers' risk aversion.

If, for all of those reasons, certain managerial biases are unlikely to be eliminated by market pressures or to be remedied by the organizational design, the question is: are top managers likely to exhibit biases in the first place? Shouldn't people successful and smart enough to become the leader of a company—in particular of large companies in the S&P 500—be the very embodiment of the hyper-rational *Homo oeconomicus*? For example, would executive selection processes not ensure that more rational and socially neutral agents rise to the top? Not necessarily. As long as a specific bias is not negatively correlated with the ability to have a successful corporate career, CEOs are not necessarily less likely to deviate from rational decision-making than the average individual. Goel and Thakor (2000) also point out that managers may even be more likely to be promoted to CEOs in a tournament setting if they underestimate project risk, which leads them to take very risky projects with the largest probabilities of very large returns. In the implicit rank-order tournament among potential successors of a CEO, the biased person is thus more likely to have the highest outcome, conditional on success, and is therefore unconditionally more likely to have the highest rank.[21]

Moreover, if superior management skills are correlated with limits on rationality and self-interest then it is conceivable that top managers are even *more* judgmentally biased and social-comparison-minded than average workers. Confidence may be correlated with other kinds of managerial value—like inspiring workers or mesmerizing Wall Street analysts—so sorting may actually *select for* overconfidence rather than filter it out. Mistakes in one type of judgment may be a price one pays for

[21] See Van Den Steen (2004) for a related argument about the endogenous generation of overconfidence due to the selection of tasks in which the agent is likely to succeed.

other kinds of managerial skill. Whether this is true is an empirical matter, which cannot be resolved purely by theory.

7.4.1 *Overconfidence among Top Managers*

Overconfidence and overoptimism, which we may define as overestimation of own skills relative to others, and overestimation of one's own prospects or future outcomes[22] are two patterns that have been widely studied in the academic literature and also discussed often in the business press. One well-established stylized fact is the "better than average" effect. Svenson (1981) first demonstrated that the vast majority of subjects, if asked to assess their driving skills compared to the median person in the group, rate themselves overwhelmingly as "above average." Svenson's finding has been replicated numerous times and with respect to various IQ- or skill-related outcomes other than driving (see, for example, Larwood and Whittaker 1977; Alicke 1985). For example, when asking a sample of entrepreneurs about their chances of success, Cooper et al. (1988) found that 33% attached zero probability to failure, 81% of them between 0 and 30%.

Several additional findings suggest that high-rank executives are particularly likely to overestimate outcomes related to their own actions or skills. First, overconfidence tends to be observed especially in decision makers who have a position of control—or at least have the illusion of control (Weinstein 1980). And who, if not the person on top of the corporate (or divisional) ladder, is the person believed to be the ultimate holder of control? A CEO has the ultimate say about the big strategic decisions and is the one who decides whether or not a large-scale investment goes ahead. Such a position may induce the CEO to believe that he or she can also control the outcome—and thus underestimate the likelihood of a bad outcome (March and Shapira 1987; Langer 1975). Second, personal commitment increases the occurrence of overconfidence (Weinstein 1980). As a quick search in ExecuComp for the option and stock packages of the top four to nine executives of publicly traded companies reveals, a huge portion of their compensation—and presumably of the wealth of many of them—depends on how well the company is doing. Moreover, the value of their human capital is tightly related company returns. If the company is underperforming, this is bad news for their prospects of keeping their job and reduces their outside options. Thus, for compensation and career reasons alone, we would expect top executives to be highly committed to the outcome of their investment,

[22]The term "overoptimism" is sometimes used to refer to the overestimation of one's own abilities as well. The labeling suggested above follows the literature on self-serving attribution and allows us to distinguish the overestimation of one's own abilities (such as IQ or managerial skills) and outcomes relating to one's own personal situation from the general overestimation of exogenous outcomes (such as the growth of the U.S. economy).

merger and other corporate decisions. And, finally, a third factor contributing to the occurrence of overconfidence within the top management team is the lack of clear measures or benchmarks for good and bad decisions. Overconfidence has been found to be strongest when it is difficult to compare performance across individuals and when the reference point is abstract (Alicke et al. 1995). Given the amount of noise in stock prices, and the long horizon over which good and bad investments pay off or fail, there is little clear feedback that would shake a CEO's overconfidence.

It is thus not surprising that experimental studies have found that executives are particularly prone to display overconfidence, both in terms of the better-than-average effect and in terms of underestimating risk (Larwood and Whittaker 1977; Moore 1977). Overconfidence, in the form of the better-than-average effect, also affects the attribution of causality. Because individuals expect their behavior to produce success, they are more likely to attribute good outcomes to their actions, but bad outcomes to (bad) luck.[23]

But how important are such biases in real-world corporate decision-making? Roll (1986) was one of the first to link overconfidence to mergers. He proposed a "hubris hypothesis": in corporate takeovers the overconfidence of individual decision makers often results in bidding firms paying too much for their targets. What he had in mind is a type of winner's curse in merger fights. Roll's hubris hypothesis may be best summarized by a famous quote from the 1981 Berkshire Hathaway Annual Report:

> Many managements apparently were overexposed in impressionable childhood years to the story in which the imprisoned handsome prince is released from a toad's body by a kiss from a beautiful princess. Consequently, they are certain their managerial kiss will do wonders for the profitability of Company T[arget]. . . . We've observed many kisses but very few miracles. Nevertheless, many managerial princesses remain serenely confident about the future potency of their kisses—even after their corporate backyards are knee-deep in unresponsive toads.
>
> (Taken from Weston et al. 1998)

Confirming this hypothesis, Malmendier and Tate (2003) find that a firm is indeed significantly more likely to acquire another firm if the CEO is overconfident. In a sample of CEOs from Forbes 500 companies in the 1980s and the first half of the 1990s, Malmendier and Tate find that overconfident CEOs are about 1.7 times as likely to conduct a merger in any given year as their rational counterpart. Here, the measurement of overconfidence is not based on responses to questions, but on actual personal portfolio decisions. The authors point to the high degree of under-diversification, which high-ranking executives face due to their stock-based compensation and the investment of their human capital in their firm, and argue that

[23] Miller and Ross (1975) provide a critical review of the abundant psychology literature on self-serving biases.

rational CEOs of Forbes 500 companies should therefore try to diversify their personal portfolios whenever possible. For example, CEOs should exercise their stock options after the vesting period, if the stock price went up by enough, and should sell company stock following a regular (precommitted) schedule.[24] But some CEOs do exactly the opposite, holding options that are well in the money (far above their exercise price) and buying, rather than selling, company stock (half the CEOs who do this lose money). In fact, during recent corporate scandals, while exorbitant pay, forgiven loans, and accounting manipulations grabbed public attention, it was less noticeable that many of the apparently guilty executives truly believed in the potential of their firms and their industry. The former CEO of WorldCom, Bernard Ebbers, lost millions of dollars because his private portfolio was narrowly invested in telecommunication stocks, while doing one merger after another.[25]

Managerial overconfidence thus has the potential to explain the big puzzle implied by $2 trillion spent on acquisitions by U.S. firms between 1998 and 2001, of which $250 billion was *lost* by shareholders of the acquiring companies (Moeller et al. 2003).

The same holds for internal investment projects. While it is hard to measure whether internal investments are generally bad bets (such data are rarely available), it is easy to find examples. Roger Smith, the former CEO of General Motors, became convinced that fully robotized plants, with no human presence, are the future of automobile production. His vision led to massive firings at GM and triggered negative responses from business analysts and engineers alike because they did not believe the technology was advanced enough. Such concerns did not stop a visionary like Smith. The result was plants in which

> the robots often began dismembering each other, smashing cars, spraying paint everywhere or even fitting the wrong equipment.
>
> (*The Economist*, August 10, 1991)

The implications of overconfidence are, however, more subtle than simply "more mergers" and "more investment," as the winner's curse view of Roll would suggest. Managerial overconfidence implies a difference in opinion about the value of the firm between the (overconfident) corporate insider and the outside investor. An overconfident CEO will persistently feel undervalued by the capital market and thus be reluctant to issue equity to finance investment projects or mergers (Heaton 2002;

[24] Precommitment allows the avoidance of sending negative signals to the market.

[25] The *New York Times*, March 3, 2004, comments: "Unlike many other telecommunications kingpins of the 1990's, who cashed out hundreds of millions of dollars in inflated stock ahead of unsuspecting investors, Mr. Ebbers apparently believed that he could keep the company afloat one way or another. He kept buying WorldCom shares even as the company's house of cards began to topple."

Malmendier and Tate 2005b). It is just the flip side of traditional Myers–Majluf-type models with asymmetric information (Myers and Majluf 1984). There, the manager has inside information, and good types are reluctant to issue equity since they will be taken for bad types. Here, the overconfident manager thinks he has (positive) inside information, with the same effect on equity issuance. Overconfidence amounts to "perceived asymmetric information." As a result, overconfidence may actually lead to *less* investment or *fewer* mergers whenever the overconfident CEO needs to tap the external capital market. A CEO is expected to overinvest only if a firm has abundant internal resources or extra debt capacity. In other words, the investment decision of an overconfident CEO should be sensitive to the availability of internal cash flow, and such investment–cash-flow sensitivity should be strongest in firms with little internal resources. At the same time, the inclination to undertake bad acquisitions should have highest impact in firms with abundant internal resources.

The data confirm all of the predictions above, including some type of pecking-order financing among overconfident CEOs (Malmendier et al. 2006). Similar evidence has been found for a sample of French start-up firms (Landier and Thesmar 2004). This increasingly robust evidence on managerial overconfidence gives rise to the question of how corporate governance should respond to CEO overconfidence. A large literature in corporate finance analyzes how options and stock grants help to align the interest of the CEO (or other corporate executive) with the interest of shareholders (Murphy 1999). The possibility of executive overconfidence alters basic conclusions about these incentive mechanisms. If we worry that a CEO just may have too positive a view on the value he or she can generate, then options and stocks are not helpful. Overconfident CEOs do not need incentives to maximize the market value of the firm's equity—that is what they perceive to be doing already. In fact, options and stocks may then push them to act as if they were risk-loving and choose investment which are riskier (and lower net present value) than shareholders would prefer. This concern is of particular relevance given that the CEO overestimates the expected value of those gambles to begin with.

If CEOs are overconfident, then other incentive devices may work better than stocks and options. One example is debt (see Hart 1995, Chapter 6). An overconfident CEO in a cash-rich firm tends to undertake too many investment projects, including acquisitions. If cash is tied up for debt repayments and the capacity for senior debt is exhausted, the CEO would need to issue more risky debt or equity to finance projects. Given that overconfident CEOs perceive their firms to be undervalued by the equity market and, consequently, expected returns to their projects will be underestimated by providers of risky debt, the preexisting debt will curtail the tendency to overinvest.

More generally, a mechanism that requires the CEO to get an outsider's approval on an investment project would counterbalance CEO overconfidence. That "outsider" could be a vigilant board. Companies typically specify in their charter that

the board has to approve any investment beyond a certain dollar amount. Setting the threshold for board approval low enough seems a sensible strategy—if we believe that the board fulfills its monitoring function. The concern about upward biases in a manager's perception thus gives additional importance to group decision-making, even in the absence of informational constraints.

Of course, while CEO overconfidence can create agency costs, overconfidence may be correlated with valuable charisma, the ability to convince investors of the value of a firm, or the ability to motivate employees. Firms may want to attract overconfident leaders but also restrain their actions in some way. One method is to unbundle the leadership and resource-allocation roles of top managers, so that a confident leader creates inspired workers and optimistic investors, but does not overinvest resources.

7.4.2 *Capital Budgeting*

The last section focused on biases in top managers and whether organizational structures can constrain them. Biases in capital budgeting which occur below the CEO's sphere of influence are also important and might crop up in three ways: a tendency to divide capital equally among divisions; escalation in commitment to bad projects; and empire building (and capital suppression).

7.4.2.1 Corporate Socialism and Partition Dependence

Scharfstein and Stein (2000) show empirically that capital is allocated relatively evenly across divisions of different sizes. They call this pattern "corporate socialism," as if equal allocations are a deliberate attempt to subsidize smaller divisions at the expense of larger ones. Lamont (1997) shows the flip side of such within-company subsidization. After the 1986 oil price shock, oil companies significantly reduced their capital expenditure in non-oil subsidiaries even though the quality of non-oil investment projects did not decrease. Equality-biased allocations might also reflect a cognitive pattern called "partition dependence" or the "$1/n$ heuristic"—the tendency for allocations to be biased toward even allocations, given a particular partition of categories (Bardolet et al. 2005). For example, Fox and Rottenstreich (2003) asked subjects to guess the probabilities that the hottest day in the next week would be Sunday or another day of the week. Many subjects answered (sensibly) that there is a $\frac{1}{7}$ chance that Sunday is hottest, and a $\frac{6}{7}$ chance that the hottest day is a day other than Sunday. Others tend to give beliefs for the chance that Sunday is hottest which lie between $\frac{1}{7}$ and $\frac{1}{2}$. The latter beliefs are influenced by the arbitrary partition of the week into two subsets, into "Sunday" and "non-Sunday." A similar effect is found in both laboratory and field data of investment in 401(k) plans (Benartzi and Thaler 2001). When asked to choose between investing in a stock or bond fund,

many subjects choose to invest 50% in each. When asked to choose between two stock funds and a bond fund, many subjects choose to spread allocations equally among the three funds, which creates an aggregate investment that is more heavily weighted ($\frac{2}{3}$) to stocks.[26]

Empirical tests of partition dependence in capital budgeting could study what happens when divisions are combined or split apart, for reasons that are exogenous to changes in the demand for capital across business units (e.g., regulation or tax considerations). Anticipating Section 7.6.2 (on organizational "repairs") organizations that recognize the tendency to allocate capital equally across divisions could deliberately structure divisions to mitigate the bias. A fast-growing product line that is starving for capital might be given its own small division to be sure it gets enough capital.

7.4.2.2 Escalation of Commitment

A common theme in behavioral judgment research is that people and organizations escalate their commitments to losing projects (Bazerman and Neale 1992). Most of this literature is based on experiments in which people make ex ante choices, receive good or bad feedback, and then decide whether to continue a project. Experiments suggest that even when feedback is bad, people continue to invest, throwing "good money" after "bad."

Despite the clear experimental evidence of escalation to commitment to bad projects, there is little convincing field evidence of this pattern. Case studies of projects which failed in retrospect have been related to escalation of commitment (see, for example, Staw and Ross 1993). But these may simply be examples of hindsight bias about especially bad projects. Other field studies analyze whether new managers who take over projects pull the plug on them. But the tendency to do so might just reflect differences in prior beliefs rather than escalation by the managers who initiated the bad projects (see, for example, Barsade et al. 1997; Weber and Camerer 1999). The ideal situation for an empirical analysis would thus be

(i) data on similar types of decisions (such as plant closings) for a set of firms,
(ii) where in a strict subset the decision maker changed for exogenous reasons (like the "sudden deaths" of CEOs in Johnson et al. 1985), and
(iii) where the new decision maker would have originally implemented the investment project as well, had he already held the position.

The constraint to (i) is the poor quality of data on decision-making inside the firm and on the success of individual projects. As for (ii), changes in managerial position

[26] Huberman and Jiang (2004) present evidence rejecting the $1/n$ heuristic on the individual level, though it is present in the aggregate.

that are exogenously based on schedule or promotion (rather than promotions which endogenously follow project failure) are most promising. And condition (iii) could be remedied by analyzing decisions with objective criteria which leave little room for differing priors.

Even if data satisfying criteria (i)–(iii) were available, we would need to go one step further and account for endogenous organizational response. Organizational design may limit bad escalation. Budget constraints and milestones are in place precisely to avoid escalation. The bigger the risk of bad escalation, the more extensive should be the implementation of such mechanisms. In fact, they may sometimes produce an opposite bad effect: good projects are terminated prematurely. Heath (1995) points out that sometimes people de-escalate *too quickly*, when a project's losses hit a budget constraint.

A promising place to study escalation and organizational responses are firms with large portfolios of small-scale investments, as in corporate R&D, or consumer product companies which expect a high rate of product failure. These firms need to routinely pull the plug on bad projects. Understanding how they cope with high rates of failure could be useful in understanding how to prevent escalation.

7.4.2.3 Empire Building and Capital Suppression

The tendency of top managers to "build empires"—to create divisions which are too large in an economic sense, benefiting managers at the expense of shareholders—is often discussed as an agency cost within the standard, rational framework. Behavioral economics offers a reinterpretation: when managers grow their firms at the detriment of shareholders, they may do so not (only) because of their desire to reign a larger empire and in full awareness of the resulting value destruction. Rather, they may (also) overestimate how much value will be created under their leadership, as discussed in Section 7.2.6.3 on overconfidence. Or they may feel competitive towards top managers of other firms, without full awareness of the resulting value destruction.

The same psychology, then, could help a better understanding of the reverse tendency—suppression of capital—in business units that are speculative but growing rapidly and have huge potential. As an example of capital suppression, Holmström and Kaplan (2001) note that it is hard to imagine how internal capital allocation could have created the kind of fast-growing technology-based firms that grew rapidly through venture capital financing and initial public offerings (IPOs) in the late 1990s. They suggest that "if Netscape, eBay or Amazon had been invented inside a big company, their potential value would probably have been overlooked" (Holmström and Kaplan 2001, pp. 138–39). They further suggest the barriers to "intrapreneurship" are due to weak internal incentives: "Even if some degree of

value had been seen, it would have been difficult or impossible to give management a strong incentive to maximize the value inherent in these ideas."

But what prevents large firms like IBM, Xerox, or Microsoft from simply reproducing the capital allocation mechanisms used by venture capitalists and then selling off its successes to IPO investors (or growing them internally) just as venture capitalists do? Given the interest of larger firms in replicating high-powered incentives for intrapreneurship, it is natural to ask whether there are other barriers that prevent them from allowing very rapid growth of promising business. One barrier might be the psychology of managers of conventional businesses who compete with the rapidly growing new business units. Loss-aversion and social comparison suggest that managers in traditional businesses will be uncomfortable about seeing lots of cash poured into businesses which could outgrow their own.[27] The managerial psychology which creates empire building probably also encourages suppression of growing empires within the firm.

7.5 Implications for Corporate Governance

Our discussion so far considered biases on the side of the managers. We touched upon adjustments in organizational design and corporate governance mechanisms to deal with them. Obviously, limits on rationality and expression of social preferences affect only managers but also the "corporate governors" e.g., board members. The influence of governance forces examination of the conjectures made in the last section. While overconfidence about the firm's projects may be less likely in external board members than in CEOs, partition dependence or escalation of commitment easily affects board decisions as well.

Another very important issue is the selection of board members. Do CEOs look for the most able and knowledgeable people? Or do they install yes-men? The fundamental problem in evaluating board quality is that the literature has yet to come up with a reliable measure of high-quality boards. The analysis of governance rules such as corporate bylaws in Gompers et al. (2003) is certainly a big step forward. But sample limitations (there are still few panel data sets with a long enough time series to analyze changes in corporate governance within firms) and endogenous firm response limit the broader use and long-term reliability. The behavioral perspective may help the quest for better measures of "who speaks up in the boardroom." Monetary incentives may not be all that matters. The ability to speak the language of the CEO and/or the other board members, or even feelings of moral obligation

[27] "The hierarchical investment approval process that is characteristic of internal capital markets is another impediment to innovation within firms.... By design, the large corporation is not set up for revolutionary inventions" (Holmström and Kaplan 2001, p. 139).

or social image concerns, may be far better determinants of directors monitoring quality.

7.6 Organizational Reactions: Sorting, Repairs, and Exploitation

A special contribution of *organizational* economics to behavioral economics will come from highlighting the role of selection and endogenous firm response to non-standard preferences. Given a pool of workers with various skills, cognitive limits and social preferences, organizational output will depend on how workers are sorted, how organizations "repair" mistakes, and whether worker mistakes can be exploited for the good of the firm.

7.6.1 *Sorting*

The starting point is heterogeneity. People are different in skills and tastes. From a behavioral economics point of view, heterogeneity means people might differ in their limits on rational calculation, willpower, and greed, and their familiarity with labor–leisure trade-offs. The impact of these behavioral biases will depend on which kind of people sort into which organizational environment. Lazear et al. (2005) demonstrate how sorting can reverse the impact of behavioral preferences in a simple experimental setting. In a standard dictator game, where subjects decide how much to share with an anonymously matched person, about 75% of people share some positive amount and 25% keep everything for themselves. These proportions reverse if subjects are allowed to sort and decide whether to play the dictator game or to "opt out" and keep the full amount for themselves (without the receiver finding out about the dictator game). Only 30% end up sharing. Interestingly, making the dictator game financially more attractive attracts first and foremost the "sharks," who keep all or most of the amount available to dictators for themselves. The flip side of this result is the selection of agents with "desired" behavioral traits. Firms may design hiring practices to identify "team players" with a sense of fairness. Promotion criteria may select workers who are subject to strong in-groups effects, inducing effort beyond monetary incentives.

Similar reasoning applies to cognitive limits. An organization will be more productive if it matches people who make mistakes with jobs in which the types of mistakes they make are not costly. This means that the organization's overall performance could be partly or largely immunized against mistakes by assortative matching.[28] Put differently, the output of organizational calculation could be more

[28]In a magazine interview, Gary Becker opined that division of labor "strongly attenuates if not eliminates any effects" caused by bounded rationality (see Stewart 2005). He conjectures that " 'it doesn't matter if 90 percent of people can't do the complex analysis required to calculate probabilities. The

rational than the averaged decisions of its individuals (or even than the smartest individual). Not much is known, however, about which types of cognitive mistakes are actually eliminated by organizational sorting and aggregation.

7.6.2 *Repairs*

Beyond the selection process, organizations can attempt to "repair" mistakes of individuals and to overcomes biases with organizational practices. Heath et al. (1998) give many illustrative examples. One example is from Microsoft. Software engineers did not believe the many complaints coming from 1-800 customer service phone lines about the difficulty of using software, or disparaged the frustrated users as stupid. To make customer complaints vivid, Microsoft used a room with a one-way mirror. Through the mirror, engineers could watch customers—who looked like normal, reasonably intelligent folks, not Luddite morons—struggling with software. The room replaced statistical information about phone line complaints with a smaller sample of more vivid information, exploiting the way in which the brain automatically privileges firsthand visual information over abstract symbols. The sheepish engineers went back to the drawing boards to try to make their software easier to use for the people on the other side of the mirror.

Another example was conveyed to us by Chris Mayer. In their study of loss-aversion in housing, Genosove and Mayer (2001) found that the nominal purchase price condominium owners in Boston, MA, had paid for their houses influenced their listing prices and the prices at which they sold their houses. In talking about his findings before some investment bankers, one person in the audience said he knew about this effect, and that his firm had taken steps to limit it. His firm combats loss-aversion by forcing a trader to periodically switch his "position" (the portfolio of assets the trader bought and is blamed or credited for) with the position of another trader in the firm. Switching ensures that traders do not make bad trades because of loss-aversion and emotional attachment to their past actions without affecting the firm's net trading position. If attachment to one's trading position is bad, the switch creates a purely behavioral scope economy, which permits firms with many traders to outcompete firms with fewer traders.

Deadlines are another kind of organizational repair. When people procrastinate, methods for enforcing deadlines will help them get work done (O'Donoghue and Rabin 2001). If firms are able to enforce deadlines better than individuals can, the firm has a cost advantage over individual workers who are self-employed. Deadline

10 percent of people who can will end up in jobs where it's required'—such as dealing blackjack, Becker's example, or managing mutual funds." (Actually, blackjack dealers' strategies are completely constrained by casino rules and therefore require no skill at probability judgment.)

enforcement can come in the hard form of firing threats, or in the soft form of feeling bad about letting coworkers down, or the in-between form of risking coworkers' wrath if a missed deadline on a joint project leads to mistaken blame.

While we are lacking convincing field data on organizational repairs, these examples suggest the following recipe for studying them. Find a mistake most people make in individual decisions; ask what organizational practice would limit its effect; then investigate whether the practice works. (Ideally, convince a firm to randomly vary the organizational design across similar departments!) Or reverse engineer the recipe: think of a curious organizational practice; ask what economic (Gibbons 2004b) or behavioral problem in individual choice it might be a solution to.

Repairs can also come from the market rather than from an organizational practice. Commenting on the slow speed of industry consolidation, Jensen (1993, pp. 847) wrote:

> In industry after industry with excess capacity, managers fail to recognize that they themselves must downsize; instead they leave the exit to others while they continue to invest.

The reluctance to downsize might stem from emotional reluctance of executives to fire people (though self-interest surely plays a role[29]). An aversion to firings and layoffs is an organizational manifestation of loss-aversion. For most managers this is a kind of effort in the simple agency model sketched above—an activity which they dislike doing but which is good for shareholders.

One solution to the problem of necessary downsizing is sorting—hire managers who do not find firing people effortful (like "Chainsaw Al" Dunlop and "Neutron Jack" Welch). Jensen suggested that the emergence of the leveraged buyout (LBO) organizational form has been a useful mechanism for downsizing (and other changes in governance would help too). An important feature of the LBO is that control shifts to new investors (the LBO partnership), who typically have no emotional attachment to the firm's workers. This element of psychological transfer, and willingness to accept a "loss" from the firm's previous policy, may be central to the success of this kind of restructuring. Of course, the LBO form also offers huge financial incentives to restructure, and to do so rapidly (because of the large debt service burdens). But maybe these large rewards just reflect how emotionally painful downsizing is for most managers.

[29]Most studies show that compensation is more strongly tied to firm size than to any other variable. Thus, laying off workers generally reduces the size of an executive's compensation. The suggestion here is that layoffs are effortful beyond these incentive effects.

7.6.3 *Exploiting Worker Biases*

A growing literature on the "behavioral economics of industrial organizations" deals with the interaction of individuals and firms (DellaVigna and Malmendier 2004; Eliaz and Spiegler 2004; Gabaix and Laibson 2005; Heidhues and Koszegi 2005). The baseline assumption in this literature is that consumers make systematic mistakes and firms are rational actors. This is reasonable because experience, specialization, larger resources, sorting, and market competition all work in the direction of enabling organizations to train and choose managers who are expert at exploiting mistakes by consumers—and by their own workers (Glaeser 2004). Rational firms can exploit the biases of individual consumers, for instance by charging high flat rates for services or goods consumers will consume less than they plan to—think of your gym attendance!—by designing contracts with automatic renewal if consumers tend to procrastinate cancelation, or by issuing stock when investor sentiment is high (DellaVigna and Malmendier 2003; Baker and Wurgler 2000).

Similarly, the individual employee in an organization may display biases from which the organization can extract rents. For example, Oyer and Schaefer (2005) analyze why firms give options to medium- or lower-level employees. Given the relatively small size of these grants and the risk premia associated with them, options are unlikely to be an efficient way to provide incentives for these employees. Option grants may, however, allow firms to attract and retain employees. The analysis of Oyer and Schaefer suggests that the benefits of option compensation are particularly high if employees are overly optimistic about the future prospects of the company. Making options part of the compensation package allows firms to identify those (potential) employees who have particularly strong beliefs about the future prospects of the company and to reap the rents from compensating them with overvalued options rather than cash. Such a compensation strategy is profitable if these employees are also more productive, for instance, due to their enthusiasm about the company.[30]

Other elements of organizational design have similar properties. Lazear and Rosen (1981) have shown conditions under which tournaments dominate other forms of

[30]As pointed out by Bergman and Jenter (2004), overoptimism about the prospects of the company by itself is not sufficient to explain options as part of the employment contract. Individuals with such beliefs can be compensated equally cheaply in cash—their reservation wage will be lower due to the overvaluation of company stock and options which they can buy with that cash. Or, they may not even work for the company and just buy the stock or options of the company. These constraints, however, do not bite if markets are incomplete. For example, the typical ten-year employee option on a public company's stock is not traded on a public exchange, and even shorter-duration options are not traded for private companies. Employees may also not be able to value options based on observed market prices. In these cases, options may be part of a profit-maximizing compensation contract design, even if employees with optimistic beliefs about a company do not have higher abilities.

labor contracts, especially with sorting among high-ability and low-ability workers. Their argument is strengthened (and more robust) if we introduce heterogeneity in self-confidence among workers. If some workers are overly convinced of their own abilities (and their overconfidence is not negatively correlated with ability), tournaments become a particularly cheap form of compensation. Overconfident individuals will overestimate the expected value of entering the tournament outcome and will be willing to accept lower current compensation.

Option-based compensation and rank-order based promotion are two examples of how firms can reduce their compensation costs at the expense of employees who are too optimistic about the firm's prospects. These examples assume implicitly that the organization itself—or the top management teams of an organization—is not affected by biases observed in individuals further down the hierarchy or outside the organization, but may remedy or even exploit them.

7.7 Conclusion

This chapter is about both *applying* behavioral economics to organizations, and *enriching* behavioral economics by asking questions that are specific to the roles of workers and managers in organizations. This chapter is less of a review of what is known than it is a research agenda, with more questions than answers. This conclusion recaps some of the common themes.

The first question is how to complicate the basic agency model. One complication is that agents' preferences (i.e., effort disutility and reservation wages) may be sensitive to how jobs are described or how wage offers and job choices are elicited. Limits on greed and self-interest imply that people often like to help their friends and harm their enemies, and dislike unequal treatment. Precise models of social preference that have been carefully honed on experimental data from simple games would need to be plugged into organizational economics applications. In the laboratory, preferences are often sensitive to reference points, such as previous wages (due perhaps to habit formation) and wages of other workers. There is evidence of such effects from labor practices (e.g., two-tier wage deals in financially distressed firms) and executive compensation. Workers also seem to care about fair or just procedures for determining outcomes. The strength of these preferences in field data and organizational design responses has, however, not been studied thoroughly. Psychological influences on judgments of causality (in agency-theory terms, inferring agent effort from output) are also important. Hindsight bias, overattribution of cause to workers rather than luck, asymmetric attribution (taking credit and exporting blame), diffusion of blame across a large group, and overconfidence about skill may all play important roles, and these have not been studied much in organizational contexts.

Section 7.2 moves from the individual to the organizational level. Organizing requires informal rules for coordinating action. Following Kreps (1990), we take the view that corporate culture is a set of shared meta-focal principles for resolving coordination problems in the face of the cognitive inability to fully specify an employment contract. Having a sharp economic concept like this to work with could provide a basis for useful empirical work on culture. The way the brain organizes information ("cognitive economy") may also explain why culture is often transmitted in the form of stories, personification, and simple slogans (Heath and Seidel 2004). The value of having cultural principles that can be commonly understood, and are immune to errors in language translation, gives a role for leadership and managerial statements that are simple (like corporate mission statements).

In large business organizations (especially U.S. firms), the judgment of top managers (e.g., CEOs) is especially influential. Section 7.3 discusses the fact that top managers are particularly prone to specific biases, such as overconfidence, given their past successes and the noisy feedback on their large-scale decisions. Then they may make mistakes, especially if weak corporate governance does not discipline mistakes. Overconfident managers, for example, will rely too much on internal cash flow when (over)investing, and enter into value-reducing mergers.

If organizations are aware of the mistakes workers and managers make, they can sort around them (via hiring practices and internal matching of workers into jobs where their mistakes are least costly), they can design "organizational repairs," or they can exploit mistakes by workers (e.g., offering backloaded incentives to optimistic workers who are sure they will succeed). Thinking about the optimal design problem in the face of worker error, as suggested in Section 4.4, is both positive economics and potentially normatively useful for managers.

We end the book with an important caveat. Psychology and economics share a basic methodological individualism: psychologists are interested in people, and economists are interested in how economic outcomes arise from the interaction of people (and institutions). Many organizational theorists, however, work at a middle level sometimes called "meso," in which the organization is the unit of analysis. Meso concepts like organizational routines and learning are not rooted in individual behavioral concepts. It is possible that organizational economics could be tied to meso-organizational (or macro-organizational) constructs, but this is not an active area of research.

References

Akerlof, G., and R. Kranton. 2005. Identity and the economics of organizations. *Journal of Economic Perspectives* 19:9–32.

Alicke, M. D. 1985. Global self-evaluation as determined by the desirability and controllability of trait adjectives. *Journal of Personality and Social Psychology* 49:1621–30.

Alicke, M. D., M. L. Klotz, D. L. Breitenbecher, T. J. Yurak, and D. Vredenburg. 1995. Personal contact, individuation, and the better-than-average effect. *Journal of Personality and Social Psychology* 68:804–25.

Andrade, G., M. Mitchell, and E. Stafford. 2001. New evidence and perspectives on mergers. *Journal of Economic Perspectives* 15:103–20.

Antle, R., and A. Smith. 1996. An empirical investigation into the relative performance evaluation of corporate executives. *Journal of Accounting and Economics* 24:1–40.

Ariely, D., U. Gneezy, G. Loewenstein, and N. Mazar. 2003. Large stakes and big mistakes. Federal Reserve Bank of Boston Working Paper 05-11.

Ariely, D., G. Loewenstein, and D. Prelec. 2004a. Tom Sawyer and the myth of fundamental value. MIT Working Paper. (Available at http://sds.hss.cmu.edu/faculty/Loewenstein/downloads/Sawyersubmitted.pdf.)

Ariely, D., E. Kamenica, and D. Prelec. 2004b. Man's search for meaning: the case of Legos. Harvard Business School Working Paper. (Available at http://www.people.fas.harvard.edu/-kamenica/papers/legosFinal.pdf.)

Babcock, L., and G. Loewenstein. 1997. Explaining bargaining impasse: the role of self-serving biases. *Journal of Economic Perspectives* 11:109–26.

Babcock, L., X. Wang, and G. Loewenstein. 1996. Choosing the wrong pond: social comparisons in negotiations that reflect a self-serving bias. *Quarterly Journal of Economics* 111:1–19.

Baker, G., M. Gibbs, and B. Holmström. 1994. The wage policy of a firm. *Quarterly Journal of Economics* 109:921–55.

Baker, M., and J. Wurgler. 2000. The equity share in new issues and aggregate stock returns. *Journal of Finance* 55:2219–57.

Baker, M., R. Ruback, and J. Wurgler. 2004. Behavioral corporate finance: a survey. Harvard Business School Working Paper. (Available at http://www.people.hbs.edu/mbaker/cv/papers/behavioralcorporatefinance.pdf.)

Bardolet, D., C. Fox, and D. Lovallo. 2005. Partition dependence in organizational capital allocation: an experimental investigation. UCLA Anderson School Working Paper.

Barsade, S., K. Koput, and B. Staw. 1997. Escalation and the credit window: a longitudinal study of bank executives' recognition and write-off of problem loans. *Journal of Applied Psychology* 82:130–42.

Bartel, A., R. Freeman, C. Ichniowski, and M. M. Kleiner. 2003. Can a work organization have an attitude problem? The impact of workplaces on employee attitudes and economic outcomes. NBER Working Paper 9987.

Bazerman, M. H., and Neale, M. A. 1992. Nonrational escalation of commitment in negotiation. *European Management Journal* 10:163–68.

Bebchuk, L., J. Fried, and D. Walker. 2002. Managerial power and rent extraction in the design of executive compensation. *University of Chicago Law Review* 69:751–846.

Belliveau, M. A., C. O'Reilly, and J. Wade. 1996. Social capital at the top: effects of social similarity and status on CEO compensation. *Academy of Management Journal* 39:1568–93.

Benabou, R., and J. Tirole. 2003. Intrinsic and extrinsic motivation. *Review of Economic Studies* 70:489–520.

Benabou, R., and J. Tirole. 2004. Incentives and prosocial behavior. Princeton University Working Paper. (Available at http://www.wws.princeton.edu/rbenabou/w11535.pdf.)

Benartzi, S., and R. Thaler. 2001. Naïve diversification strategies in defined contribution saving plans. *American Economic Review* 91:79–98.

Bennis, W., and D. Nanus. 1985. *Leaders: The Strategies for Taking Charge*. New York: Harper and Row.

Bergman, N., and D. Jenter. 2004. Employee sentiment and stock option compensation. MIT Sloan Research Paper 4504-04.

Bertrand, M., and S. Mullainathan. 2001. Are CEOs rewarded for luck? The ones without principals are. *Quarterly Journal of Economics* 116:901–32.

———. 2003. Enjoying the quiet life? Corporate governance and managerial preferences. *Journal of Political Economy* 111:1043–75.

Bettman, J. R., and B. A. Weitz. 1983. Attributions in the board room: causal reasoning in corporate annual reports. *Administrative Science Quarterly* 28:165–83.

Bowman, D., D. Minehard, and M. Rabin. 1999. Loss aversion in a consumption–savings model. *Journal of Economic Behavior and Organization* 38:155–78.

Brandts, J., and D. Cooper. 2004. A change would do you good.... An experimental study on how to overcome coordination failure in organizations. Weatherhead Business School Working Paper. (Available at http://pareto.uab.es/wp/2004/60604.pdf.)

Brockner, J., and B. Wiesenfeld. 1996. An integrative framework for explaining reactions to decisions: interactive effects of outcomes and procedures. *Psychological Bulletin* 120:189–208.

Camerer, C. F. 2003. *Behavioral Game Theory*. New York: Russell Sage Foundation Press and Princeton University Press.

Camerer, C. F., and G. Loewenstein. 2004. Behavioral economics: past, present, future. In *Advances in Behavioral Economics* (ed. C. Camerer, G. Loewenstein, and M. Rabin). New York: Russell Sage Foundation Press and Princeton University Press.

Camerer, C. F., and A. Vepsäläinen. 1988. The efficiency of cultural contracting. *Strategic Management Journal* 9:77–94.

Camerer, C. F., L. Babcock, G. Loewenstein, and R. Thaler. 1997. Labor supply of New York City cab drivers: one day at a time. *Quarterly Journal of Economics* 112:407–42.

Cameron, K. S., and R. E. Quinn. 1999. *Diagnosing and Changing Organizational Culture*. Reading, MA: Addison-Wesley.

Cooper, A. C., C. Y. Woo, and W. C. Dunkelberg. 1988. Entrepreneurs' perceived chances for success. *Journal of Business Venturing* 3:97–108.

Cooper, R., D. DeJong, B. Forsythe, and T. Ross. 1994. Alternative institutions for resolving coordination problems: experimental evidence on forward induction and preplay communication. In *Problems of Coordination in Economic Activity* (ed. J. Friedman). Kluwer.

Crémer, J., L. Garicano, and A. Prat. 2003. Language and the theory of the firm. University of Chicago Graduate School of Business Working Paper. (Available at http://gsbwww.uchicago.edu/fac/luis.garicano/research/index_files/codes.pdf.)

DellaVigna, S., and U. Malmendier. 2003. Overestimating self-control: evidence from the health club industry. Stanford GSB Research Paper 1800.

———. 2004. Contract design and self-control: theory and evidence. *Quarterly Journal of Economics* 119:353–402.

Dunning, D., J. A. Meyerowitz, and A. D. Holzberg. 1989. Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology* 57:1082–90.

Eliaz, K., and R. Spiegler. 2004. Contracting with diversely naive agents. CEPR Discussion Paper 4573.

Falk, A., and A. Ichino. 2003. Clean evidence on peer pressure. IZA Discussion Paper 732. (Available at ftp://repec.iza.org/RePEc/Discussionpaper/dp732.pdf.)

Fehr, E., and S. Gächter. 2000. Cooperation and punishment in public goods experiments. *American Economic Review* 90:980–94.

——. 2000. Fairness and retaliation: the economics of reciprocity. *Journal of Economic Perspectives* 14:159–81.

Fehr, E., and L. Götte. 2003. Do workers work more when wages are high? Evidence from a randomized field experiment. IEW Working Paper 125.

Feiler, L., and C. F. Camerer. 2006. Code creation in endogenous merger experiments. Caltech Working Paper. (Available at http://www.hss.caltech.edu/~camerer/codecreation_feb2006.pdf.)

Fox, C. R., and Y. Rottenstreich. 2003. Partition priming in judgment under uncertainty. *Psychological Science* 14:195–200.

Frank, R. 2005. Departures from rational choice: with and without regret. In *The Law and Economics of Irrational Behavior* (ed. F. Parisi and V. Smith), pp. 13–36. Stanford University Press.

Gabaix, X., and D. Laibson. 2005. Shrouded attributes, consumer myopia, and information suppression in competitive markets. NBER Working Paper 11755.

Garvey, G., and T. Milbourn. Forthcoming. Asymmetric benchmarking in compensation: executives are paid for good luck but not punished for bad. *Journal of Financial Economics*.

Genosove, D., and Ch. Mayer. 2001. Loss aversion and seller behavior: evidence from the housing market. *Quarterly Journal of Economics* 116:1233–60.

Gervais, S., J. B. Heaton, and T. Odean. 2003. Overconfidence, investment policy and stock options. Duke University Working Paper. (Available at http://faculty.haas.berkeley.edu/odean/papers/Managers/GervaisHeatonOdean0703.pdf.)

Gibbons, R. 2004a. Four formal(izable) theories of the firm? MIT Department of Economics Working Paper 04-34.

——. 2004b. What is economic sociology and should any economists care? *Journal of Economic Perspectives* 19:3–7.

Gil-White, F. J. 2001. Are ethnic groups biological "species" to the human brain? Essentialism in our cognition of some social categories. *Current Anthropology* 42:515–54.

Glaeser, E. 2004. Psychology and the market. *American Economic Review* 94:408–13.

Goel, A. M., and A. V. Thakor. 2000. Rationality, overconfidence, and leadership. University of Michigan Working Paper.

Gompers, P. A., J. L. Ishii, and A. Metrick. 2003. Corporate governance and equity prices. *Quarterly Journal of Economics* 118:107–55.

Hall, B., and K. J. Murphy. 2003. The trouble with stock options. *Journal of Economic Perspectives*. 17:49–70.

Hart, O. 1995. *Firms, Contracts, and Financial Structure*. Oxford University Press.

Heath, C. 1995. Escalation and de-escalation of commitment in response to sunk costs: the role of budgeting in mental accounting. *Organizational Behavior and Human Decision Processes* 62:38–54.

Heath, C., and V. Seidel. 2004. Language as a coordinating mechanism: how linguistic memes help direct appropriate action. Stanford University Working Paper. (Available at http://www.si.umich.edu/ICOS/Linguisticmemes4.2.pdf.)

Heath, C., and N. Staudenmayer. 2000. *Coordination Neglect: How Lay Theories of Organizing Complicate Coordination in Organizations* (ed. B. M. Staw and R. I. Sutton), Research in Organizational Behavior Series, Volume 22, pp. 153–91. Greenwich, CT: JAI Press.

Heath, C., R. P. Larrick, and J. Klayman. 1998. Cognitive repairs: how organizations compensate for the shortcomings of individual learners. *Research in Organizational Behavior* 20:1–37.

Heaton, J. B. 2002. Managerial optimism and corporate finance. *Financial Management* 31:33–45.

Heidhues, P., and B. Koszegi. 2005. The impact of consumer loss aversion on pricing. CEPR Working Paper 4849.

Hermalin, B. 2001. Economics and corporate culture. In *The International Handbook of Organizational Culture and Climate* (ed. C. L. Cooper, S. Cartwright, and P. C. Earley). Wiley.

Hertel, G., and N. Kerr. 2001. Priming in-group favoritism: the impact of normative scripts in the minimal group paradigm. *Journal of Experimental Social Psychology* 37:316–24.

Holmström, B., and S. Kaplan. 2001. Corporate governance and merger activity in the United States: making sense of the 1980s and 1990s. *Journal of Economic Perspectives* 15(2):121–44.

Holmström, B., and P. Milgrom. 1991. Multi-task principal–agent analyses: incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization* (Special Issue) 7:24–52.

Huberman, G., and W. Jiang. 2004. The $1/N$ heuristic in 401(k) plans. EFA 2004 Maastricht Meetings Paper 2036. Mimeo.

Ichino, A., and G. Maggi. 2000. Work environment and individual background: explaining regional shirking differentials in a large Italian firm. *Quarterly Journal of Economics* 115:1057–90.

Jensen, M. 1993. Modern industrial revolution, exit, and the failure of internal control systems. *Journal of Finance* 48:830–80.

Johnson, B., R. Magee, N. Nagarajan, and H. Newman. 1985. An analysis of the stock price reaction to sudden executive deaths: implications for the management labor market. *Journal of Accounting and Economics* 7:151–74.

Kahneman, D., and D. Lovallo. 1993. Timid choices and bold forecasts: a cognitive perspective on risk taking. *Management Science* 39:17–32.

Kahneman, D., J. Knetsch and R. H. Thaler. 1986. Fairness and the assumptions of economics. *Journal of Business* 59:285–301.

Kandel, E., and E. P. Lazear. 1992. Peer pressure and partnerships. *Journal of Political Economy* 100:801–18.

Knez, M., and C. F. Camerer. 1994. Creating "expectational assets" in the laboratory: "weakest-link" coordination games. *Strategic Management Journal* 15:101–19.

Knez, M., and D. Simester. 2001. Firm-wide incentives and mutual monitoring at Continental Airlines. *Journal of Labor Economics* 19:743–72.

Kole, S., and K. Lehn. 2000. Workforce integration and the dissipation of value in mergers: the case of USAir's acquisition of Piedmont Aviation. In *Mergers and Productivity* (ed. S. Kaplan), pp. 239–79. NBER and University of Chicago Press.

Kolligian Jr., J., and R. J. Sternberg. 1991. Perceived fraudulence in young adults: is there an "imposter syndrome"? *Journal of Personality Assessment* 56:308–26.

Koszegi, B., and M. Rabin. 2004. A model of reference-dependent preferences. EconWPA Working Paper 0407001.

Kotter, J., and J. L. Heskett. 1992. *Corporate Culture and Performance*. New York: Free Press.

Kovalchik, S., C. F. Camerer, D. M. Grether, C. R. Plott, and J. M. Allman. 2005. Aging and decision making: a comparison between neurologically healthy elderly and young individuals. *Journal of Economic Behavior and Organization* 8:79–94.

Kremer, M. 1993. The O-ring theory of economic development. *Quarterly Journal of Economics* 108:551–76.

Kreps, D. M. 1990. Corporate culture and economic theory. In *Perspectives on Positive Political Economy* (ed. J. Alt and K. Shepsle). Cambridge University Press.

——. 1997. Intrinsic motivation and extrinsic incentives. *American Economic Review* 87:359–64.

Landier, A., and D. Thesmar. 2004. Financial contracting with optimistic entrepreneurs: theory and evidence. CEPR Discussion Paper 3971.

Lamont, O. 1997. Cash flow and investment: evidence from internal capital markets. *Journal of Finance* 52:83–109.

Langer, E. 1975. The illusion of control. *Journal of Personality and Social Psychology* 32:311–28.

Larwood, L., and W. Whittaker. 1977. Managerial myopia: self-serving biases in organizational planning. *Journal of Applied Psychology* 62:194–99.

Lazear, E. 2004. Balanced skills and entrepreneurship. *American Economic Review* 94:208–11.

Lazear, E., and S. Rosen. 1981. Rank–order tournaments as optimum labor contracts. *Journal of Political Economy* 89:841–65.

Lazear, E., U. Malmendier, and R. Weber. 2005. Sorting in experiments with application to social preferences. NBER Working Paper 12041.

Lepper, M. R., D. Greene, and R. E. Nisbett. 1973. Undermining children's intrinsic interest with extrinsic reward: a test of the "overjustification" hypothesis. *Journal of Personality and Social Psychology* 28:129–37.

Loewenstein, G., and S. Issacharoff. 1994. Source-dependence in the valuation of objects. *Journal of Behavioral Decision Making* 7:157–68.

Main, B., C. O'Reilly, and J. Wade. 1993. Top executive pay: tournament or teamwork? *Journal of Labor Economics* 11:606–28.

Malmendier, U., and G. Tate. 2003. Who makes acquisitions? CEO overconfidence and the market's reaction. Stanford Research Paper 1798.

——. 2005a. Superstar CEOs. 7th Annual Texas Finance Festival Paper. (Available at http://ssrn.com/abstract=709861.)

Malmendier, U., and G. Tate. 2005b. CEO overconfidence and corporate investment. *Journal of Finance* 60:2660–2700.

Malmendier, U., G. Tate, and J. Yan. 2006. Corporate financial policies with overconfident managers. (Available at http://faculty-gsb.stanford.edu/malmendier/personal_page/Papers/cs050206.pdf.)

March, J., and Z. Shapira. 1987. Managerial perspectives on risk and risk taking. *Management Science* 33:1404–19.

Marques, J. M., V. Y. Yzerbyt, and J.-P. Leyens. 1988. The "Black Sheep Effect": extremity of judgements towards ingroup members as a function of group identification. *European Journal of Social Psychology* 18:1–16.

Martin, J. 1982. Stories and scripts in organizational settings. In *Cognitive Social Psychology* (ed. A. Hastorf and A. Isen), pp. 255–305. Elsevier.

Miller, D., and M. Ross. 1975. Self-serving biases in the attribution of causality: fact or fiction? *Psychological Bulletin* 82:213–25.

Moeller, S., F. Schlingemann, and R. Stultz. 2003. Do shareholders of acquiring firms gain from acquisitions? *Research Technology Management* 46:62–66.

Moore, P. G. 1977. The manager's struggle with uncertainty. *Journal of the Royal Statistical Society*A 149:129–65.

Mullainathan, S., and R. H. Thaler. 2001. Behavioral economics. In *International Encyclopedia of Social Sciences*, 1st edn, pp. 1094–1100. New York: Pergamon.

Murphy, K. J. 1999. Executive compensation. In *Handbook of Labor Economics* (ed. O. Ashenfelter and D. Card), Volume 3B, Chapter 38. Amsterdam: North-Holland.

Myers, S., and N. Majluf. 1984. Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics* 13:187–221.

Nisbett, R. E. 2004. *The Geography of Thought: How Asians and Westerners Think Differently... and Why*. New York: Free Press.

O'Donoghue, T. D., and M. Rabin. 2001. Choice and procrastination. *Quarterly Journal of Economics* 116:121–60.

Oyer, P., and S. Schaefer. 2005. Why do some firms give stock options to all employees? An empirical examination of alternative theories. *Journal of Financial Economics* 76:99–133.

——. 1998. Psychology and economics. *Journal of Economic Literature* 36:11–46.

Radner, R., and T. Van Zandt. 2001. Real-time decentralized information processing and returns to scale. *Economic Theory* 17:545–75.

Rebitzer, J. 2002. Monitoring, motivation and management: the determinants of opportunistic behavior in a field experiment. *American Economic Review* 92:840–64.

Roll, R. 1986. The hubris hypothesis of corporate takeovers. *Journal of Business* 59:197–217.

Roy, D. 1952. Quota restriction and goldbricking in a machine shop. *American Journal of Sociology* 57:427–42.

Salanié, B. 2003. Testing contract theory. *CESifo Economic Studies* 49:461–77.

Scharfstein, D. S., and J. C. Stein. 2000. The dark side of internal capital markets: divisional rent seeking and inefficient investment. *Journal of Finance* 55:2537–65.

Schein, E. H. 1983. The role of the founder in creating organizational culture. *Organizational Dynamics* 12:13–28.

Schelling, Th. 1960. *The Strategy of Conflict*. Harvard University Press.

Shafir, E., P. Diamond, and A. Tversky. 1997. Money illusion. *Quarterly Journal of Economics* 112:341–74.

Sniezek, J., and T. Buckley. 1992. Confidence depends on the level of aggregation. *Journal of Behavioral Decision Making* 4:263–72.

Sørensen, J. 2002. The strength of corporate culture and the reliability of firm performance. *Administrative Science Quarterly* 47:70–91.

Staw, B., and J. Ross. 1993. Organizational escalation and exit: the case of the Shoreham nuclear power plant. *Academy of Management Journal* 36:701–32.

Stern, S. 2004. Do scientists pay to be scientists? *Management Science* 50:835–53.

Stewart, S. A. 2005. Can behavioral economics save us from ourselves? University of Chicago Magazine. Feature article, February, Volume 97, Issue 3.

Svenson, O. 1981. Are we all less risky and more skillful than our fellow drivers? *Acta Psychologica* 47:143–48.

Taylor, S. E., and J. D. Brown. 1988. Illusion and well-being: a social psychological perspective on mental health. *Psychological Bulletin* 103:193–210.

Torry, S. 1996. The paper chase slows down, here and nationally. *Washington Post*, June 10, 1996. (Available at http://www.washingtonpost.com/wp-adv/classifieds/careerpost/library/lesslaw.htm.)

Tyler, T. R. 2001. Cooperation in organizations. In *Social Identity Processes in Organizational Contexts* (ed. M. A. Hogg and D. J. Terry), pp. 149–66. Philadelphia, PA: Psychology Press.

Van Den Steen, E. 2004. Rational overoptimism (and other biases). *American Economic Review* 94:1141–51.

——. 2005. Organizational beliefs and managerial vision. *Journal of Law, Economics and Organization* 21:256–83.

Van Huyck, J. B., R. C. Battalio, and R. O. Beil. 1990. Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review* 80:234–48.

Weber, R. A. 2006. Managing growth to achieve efficient coordination in large groups. *American Economic Review* 96:114–26.

Weber, R. A., and C. F. Camerer. 2003. Cultural conflict and merger failure: an experimental approach. *Management Science* 49:400–15.

——. 1999. The econometrics and behavioral economics of escalation of commitment: a re-examination of Staw and Hoang's NBA data. *Journal of Economic Behavior and Organization* 39:59–82.

Rick, S., R. A. Weber, and C. F. Camerer. 2004. The effects of organizational structure and codes on the performance of laboratory "firms". Carnegie Mellon University Working Paper. (Available at http://ssrn.com/abstract=644346.)

Weinstein, N. 1980. Unrealistic optimism about future life events. *Journal of Personality and Social Psychology* 39:806–20.

Wernerfelt, B. 2004. Organizational languages. *Journal of Economics & Management Strategy* 13:461–472.

Weston, J. F., K. Chung, and J. Sui. 1998. *Takeovers, Restructuring and Corporate Governance*. Upper Saddle River, NJ: Prentice Hall.

Yamagishi, T. 1988. The provision of a sanctioning system in the United States and in Japan. *Social Psychology Quarterly* 51:264–70.

Yerkes, R. M., and J. D. Dodson. 1908. The relation of strength of stimulus to rapidity of habit-formation. *Journal of Comparative Neurology and Psychology* 18:459–82.

Zink, C. F., G. Pagnoni, M. E. Martin-Skurski, J. C. Chappelow, and G. S. Berns. 2004. Human striatal response to monetary reward depends on saliency. *Neuron* 42:509–17.

Zwiebel, J. 1995. Corporate conservatism. *Journal of Political Economy* 103:1–25.

COMMENT BY MICHAEL D. COHEN

Beyond Bounded Rationality [31]

This chapter views organizational phenomena through the lens of behavioral economics. It generates many productive insights. In the course of doing so it shifts between two different conceptual frameworks in a way that reflects an important contest of perspectives playing out at the moment across the social sciences.

The careful reader will notice that two quite different organizing perspectives are employed at different points in the article, each reflecting one side of the contention that is in progress. The first perspective employed might be called "deviations from baseline rationality," while the second organizing principle might be labeled "the economic implications of how the brain works." One sees examples of the former in their discussions of topics such as hindsight bias and partition dependence. Examples of the latter occur in their discussions of the "cognitive economies" achieved by narratives and personification and of the set-point dependencies that may shape the valuation of wages.

In my view, this oscillation among frameworks embodied in the essay accurately mirrors the state of the field. Across the social sciences, the working assumptions about the human mind that have relied on metaphors of computation are contending with metaphors of biological process. The "deviations" perspective exemplifies the older approach, with roots stretching back to Herbert Simon's bounded rationality and his models of heuristic human problem solving that rest in turn on his symbol processing hypotheses. Although Simon and his followers engaged in some bitter contests with orthodox economists, both camps worked with a view centered on a strongly cognitive approach to human decision-making. Both presumed that much of what was important about choice and problem solving was accessible for research

[31] The essay by Colin Camerer and Ulrike Malmendier on "Behavioral Economics of Organizations" included in the present volume reflects not only the stimulating presentation at the conference, but also the lively discussion that followed, and it embodies further valuable work done subsequently by the authors. Similarly, this commentary contains material from remarks made at the conference, but also responds both to the following discussions and to the revisions made in the Camerer–Malmender article. By now both documents are well-removed from being transcripts of the original sessions. I have, however, retained the title of the conference remarks, which was chosen with the aim of being somewhat provocative. Despite many changes, that remains my goal, as I believe that a period of important change is underway in this field, and I hope to accelerate it.

via introspection and "thinking aloud" protocols. Both cleanly factored emotional aspects of action, allowing them to enter models of choice or decision only through concepts such as value, utility, goal, or aspiration level. In this approach, models of economic action based on empirical observation are framed as conforming with, or systematically deviating from, behavior expected under conventional assumptions of rationality.

The perspective that might be called "economic implications of how the brain works" is newer. It is being propelled into the foreground by a number of cultural and technical forces, but most notably by the very rapid gains in psychological measurement capability that are discussed below. In this approach the starting point is models of mental processing as understood in psychological research. Economic action is taken to be generated by individuals who are characterized by the memory processes, habitual dispositions, and emotional response patterns found in psychological studies. Economic analysis then tries to elucidate the consequences of such actions.

The defining feature of the "deviations from baseline rationality" approach is that it is organized around observed patterns of economic action that fall short of some plausible standard of efficiency. People are found to exhibit a consistent tendency to act in some way that "leaves money (or utility) on the table." For example, they underestimate the probabilities of events that did not occur (examples of this type are collected as "hindsight bias"). Indeed, from this viewpoint, observed deviations are often characterized as "mistakes." The term implies that psychological processes are being assessed for their adequacy in meeting the demands of economic rationality.

In the second approach, it is the problem of economic research to accurately model action generated by normal psychological processes. Though psychological accuracy would seem to be a valuable goal, the risk of this approach is that psychologically rich models may lose contact with the powerful theoretical machinery developed over the last century that allows economists to infer outcomes for groups from assumptions about individuals.

A research survey such as the Camerer–Malmendier essay tries to look both backward over what is known, and forward toward what should be researched. Since so much of what is known was accumulated in the "deviations" perspective, it is appropriate that it play a large role. However, since the gains are developing so rapidly in contemporary psychology and the new tools becoming available are so powerful, it is natural that much of the forward thinking about what research is promising should bring in the perspective of "how the brain works."[32]

[32]Indeed, though it is not mentioned here, Camerer himself is doing intriguing work of this kind, as in his recent collaboration on brain imaging of trust development (King-Casas et al. 2005).

My remarks are therefore divided into two clusters. The first consists of observations on the strengths and limitations of the deviations framework. The second discusses examples that may help to bring out the potential—and the challenges—for economics of taking "how the brain works" as a starting point for analysis of organizational economics.

The "Deviations" Perspective

There are some important advantages to the deviations framework as a way of organizing results from behavioral economics for application to the economics of organization. First and foremost, the underlying behavioral economics literature has often used this framework, and so transfer to application in the economics of organization is most direct when it is maintained.

But additional advantages are also obtained. For example, the deviations approach often suggests extensions of existing formal apparatus. Camerer and Malmendier provide a nice example of this with their addition of a skill term to a more standard equation for effort level in an agency model. Though it is a complication, the extension is not unmanageable, and the authors show that the term may allow theory to accommodate a range of observed deviations from simpler representations. The deviation approach is often able to suggest promising modification of existing formalisms. Indeed the whole apparatus of risk aversion, now taken for granted as standard theory, can be viewed as embodying just such a modification.

The deviations approach likewise suggests areas of direct potential benefit to management or policy. Because people are known from field and laboratory observation to have systematic tendencies to miss potential gains, pointing out a general pattern like hindsight bias can be directly useful to practitioners in avoiding costly errors. Moreover, the naming conventions of the deviations approach result in quite memorable labeling of the principles, with the implication that students can be fairly effectively taught to watch out for problems like "the fundamental attribution error."

And finally, the deviations approach suggests to those who study the economics of organization research questions that are likely to be productive. Camerer and Malmendier illustrate this advantage nicely with their account of known individual deviations that may motivate organizations to develop corrective devices. For example, they suggest that budget constraints and project milestones may be organizational devices that function to limit individual tendencies to escalate bad commitments.

There are, of course, some costs to the deviations framework, along with the benefits.

Earlier I suggested the advantage that a deviations approach facilitates adjusting current theory to observed patterns of "mistake". But the quite distorted psychology this can build into theory may itself become a rhetorical albatross. As the richer model

of mind now developing in psychology gains credibility in the general intellectual community, that shift may markedly constrain the freedom of economics to operate credibly in the eyes of policy makers and managers while retaining the strongly stylized representations of economic action that might be the most convenient for formalization. Consumers of economic advice may begin to ask where human habits, skills, identifications, and emotions enter into the analysis.

And within economics itself the strategy of modifying existing theory to accommodate behaviorally observed deviations may gradually undermine the elegant simplicity that has been one of its hallmark attractions. The danger is that theory can become encumbered with a mass of Ptolemaic epicycles and the discipline can be left behind after an eventual Copernican simplification.

Still, we have seen that there are considerable advantages to the deviations approach and therefore good reasons to continue using it, despite some problems. Now we can turn to some examples suggestive of the potential—and challenges—of the alternative. I offer my views on the second perspective, the approach that starts from "how the brain works," not as an economist, and also not as a psychologist, but as an organizational theorist who tries to pay close attention to psychology and economics.

At bottom, the relationship to psychology has been a major feature distinguishing organization theory from organizational economics. Roughly, one might say that many organizational theorists have been closer to the second approach described here, feeling obliged to take psychological (and social psychological) results on their own terms, and build up organization theory with those results as starting points, while economists have asked—and rightly so—how psychological results may be assimilated to the needs and theories of economics. In consequence, organization theory has far less formalized theoretical development than does organizational economics, though it addresses a broader array of significant issues.[33]

The Perspective Starting from "How the Brain Works"

My method for illustrating the possibilities now emerging for the second approach will be to consider a few examples of striking developments in psychology and their possible significance for economists' views of organization. I have focused in

[33]I think this distinction raises important choices for economics: the economics of organization that meets the needs of economics' other fields (finance, health, labor, macro, . . .) may not need to take into account the full richness of contemporary psychological research, or the full complexity of what we see going on in real organizing. Or it may. . . . But "real economists" have far more wisdom about this question than I do, and so I will devote myself to extending the portrait of what I believe is going on, leaving for those within economics the problems of judging how much of the development needs to be assimilated, and of how that can be done without losing the inferential power of current theoretical formalizations.

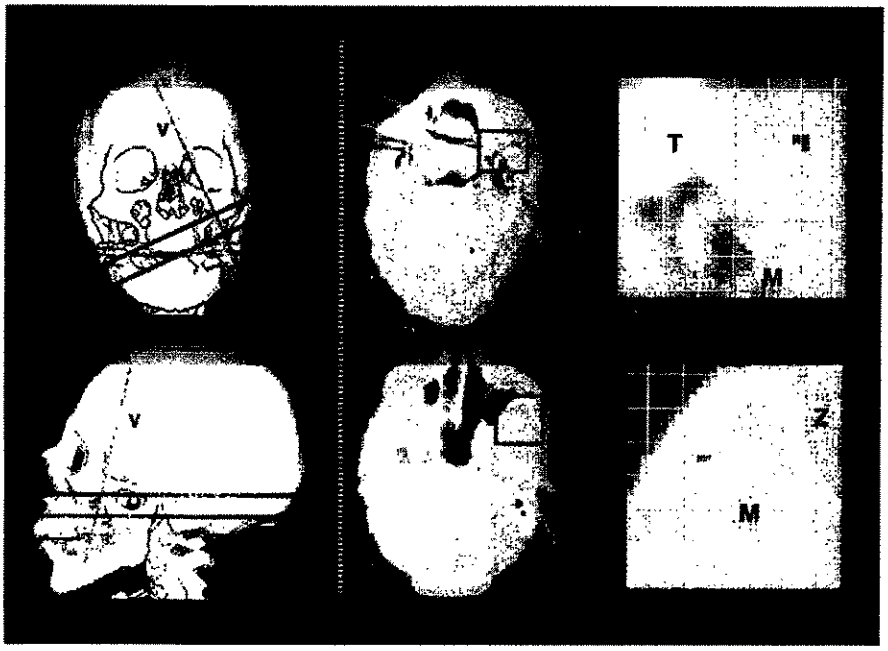

Figure 7.2. Images for analysis of skull of Phineas Gage. Reprinted with permission from Damasio et al. (1994). "The return of Phineas Gage: Clues about the brain from the skull of a famous patient." *Science*, Volume 264, p. 1102. Copyright [1994] AAAS. Reproduced courtesy of Dr Hanna Damasio, the Dana and David Dornsife Cognitive Neuroscience Imaging Center and Brain and Creativity Institute, University of Southern California.

particular on examples involving emotions, since these seem to present some of the deepest challenges.

The first thing to say is that psychology itself is being transformed—as so often happens to a science—in the wake of stunning new possibilities of measurement. The most dramatic and widely noted are various forms of brain imaging—such as positron emission tomography (PET) and functional magnetic resonance imaging (fMRI)—but there are others as well, including improved precision of lesion studies in humans and in animal models; improved chemical detection methods (a few dollars for a simple, reliable saliva assay of important chemicals such as testosterone, cortisol, or progesterone); knockout gene studies in animal models; improved placement of, and recording from, micro-electrodes, now reaching the single brain-cell level; and many more.

The consequences of this measurement gain in psychology are many. Among them is an important increase in our ability to exhibit well-measured counterparts of the previously "soft" observations of clinicians. Thus "hardened," they become more difficult to set aside as we study decision and choice.

A case that dramatizes this change is that of Phineas Gage, a young man of excellent reputation and work habits who was the victim of a construction accident

in 1848 that shot a metal rod cleanly through his head. Gage made a full physical recovery from the accident. His intelligence was unaffected. But his physician documented dramatic changes in his character in the dozen years he lived thereafter. Gage lost his ability to conform to social conventions or to honor his own commitments. He became an unemployable and foul-mouthed wanderer. The physician archived Gage's skull after he died, and 155 years later researchers were able to correlate the areas between the entry and exit holes with what is now known from many other studies about lesions affecting the same areas of the brain, the right ventro-medial prefrontal cortex.

Some images from a paper reporting the re-analysis of Gage's skull are shown in Figure 7.2.

And here is the authors' interpretation, linking their analysis to the cases of many modern patients with similar behavioral disturbances:

> For patients with damage comparable to Gage's, their ability to make rational decisions in personal and social matters is invariably compromised and so is their processing of emotion. On the contrary, their ability to tackle the logic of an abstract problem, to perform calculations, and to call up appropriate knowledge and attend to it remains intact. The establishment of such a pattern has led to the hypothesis that emotion and its underlying neural machinery participate in decision-making within the social domain and has raised the possibility that the participation depends on the ventromedial frontal region. This region is reciprocally connected with the subcortical nuclei that control basic biological regulation, emotional processing, and social cognition and behavior, for instance, in amygdala and hypothalamus.
>
> (Damasio et al. 1994)

The case serves to demonstrate two points:

(i) softer clinical observations involving emotions, feelings and social responsibilities can now be aligned with much more precise results from imaging techniques, and lesion studies; and

(ii) fundamental abilities of memory and calculation that have received so much of our attention in work on choice (or decision) are not sufficient to produce minimally effective economic activity.

In normal humans these rational capabilities participate together with emotional responses in a complex choreography that generates thought and action, and that learns from ensuing events—or sometimes does not. In the special cases of injury and disease that have been carefully observed, the choreography breaks down and the breakdown gives us a window onto the fundamental—and often subtle—role of emotion in normal choice activity in social settings.

Another line of research that provides a stimulating example for thinking about organizational settings is work on "mirror neurons." This example is taken from a

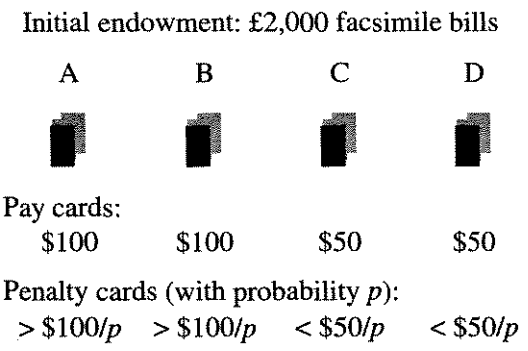

Figure 7.3. Setup of card-drawing decks.

review by Rizzolatti et al. (2001) of work on such cells. Mirror neurons are active both in motor acts, such as grasping an apple, and in observing the same motor act being carried out by another. Many techniques are involved, but among the most striking is firing data at the individual neuron level recorded from monkeys watching food being placed on a platter and then taking the food off the platter themselves.

The monitored neurons fired both on observing food placement by the experimenter's hand and in the execution of the motor act by the monkey's hand. But when the monkey observed the food being placed using pliers, there was no response from its mirror neurons.

This mirror system is implicated in understanding the actions of others by partially activating the same systems within ourselves that would be involved in carrying out the observed act. And I should add that there is considerable evidence that the mirror neuron systems connect to areas responsible for emotions, such as the amygdala. So, for example, a congenital nerve malformation gives rise to Mobius syndrome. Those suffering it have difficulty moving their facial muscles. These patients also have difficulty learning to understand the emotions expressed in the faces of others.

Results of this kind suggest that emotional responses to experiences of others may be a fundamental predisposition of our normal biological "equipment," that our understanding and evaluation of actions of others arises to some extent through rather detailed mapping onto our own experience, which our body develops special circuitry to carry out. Further work along these lines will presumably have implications for issues such as group identification and altruism.

As a last example, consider an experiment conducted by Bechara and colleagues (1997). In it, subjects, who were either normal or patients with damage to the ventromedial prefrontal cortex, drew cards from one of four decks, beginning with a stock of $2,000 in facsimile bills, and instructed to maximize winnings.

The decks, as indicated in Figure 7.3, contained payoff cards mixed with low-frequency (probability p) penalty cards. In decks A and B, although the payoffs were larger ($100 instead of $50), the penalties were large enough to imply that a

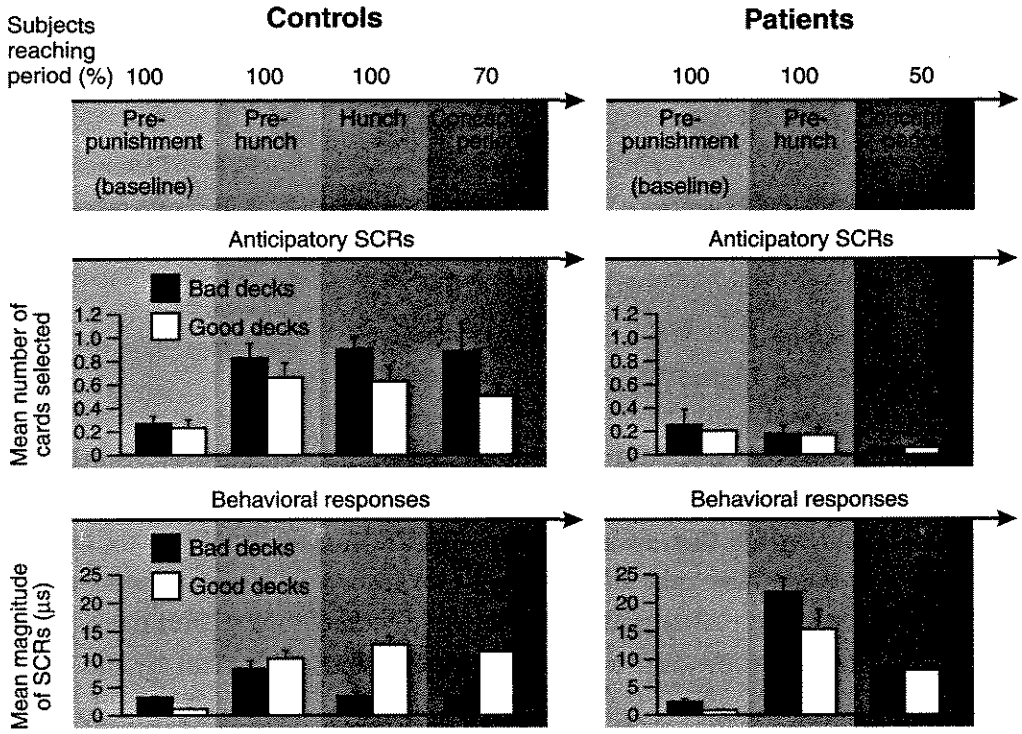

Figure 7.4. Results of card drawing experiment. Reprinted with permission from Bechara et al. 1997. "Deciding advantageously before knowing the advantageous strategy." *Science*, Volume 275, p. 1293. Copyright [1997] AAAS.

draw from the deck was an expected value loss. In decks C and D, with payoffs half the size, the penalties were small enough that they implied expected value gain. No end of play was pre-announced. Play was stopped after 100 trials.

First notice that across the top the data are distinguished for controls on the left and patients on the right. Below those labels this band indicates data separation into phases: prepunishment (before any penalty cards are drawn), pre-hunch phase (after some penalties, but before participants say "I like this deck"), hunch phase (expressed liking *and* anticipatory skin conductance response (SCR)), and conceptual phase (articulating something like "penalties outweigh benefits in $100 decks, but not in $50s").

Skin conductance reveals the subjects' emotional responses as their hands move over the cards. There are a number of things of note:

(i) there is no SCR for any subjects in prepunishment;

(ii) SCR rises for controls in the pre-hunch phase, but stays down for patients, who, in fact, never show significant SCR;

(iii) for controls, the initial preference for bad decks gives way in the pre-hunch

phase and becomes very weak in the hunch phase (in fact some control subjects never got the theory but steadily played the C and D decks);

(iv) half the patients were able to state the theory, but they still played A and B more than they played C and D!

Normal subjects are responding emotionally to projected actions even before they can voice a hunch, while patients only have baseline responses, never express hunches, and do not execute the theory even if they do state it correctly. The data thus provide another indication that in the brain a complex process is required to relate evaluations of prior experiences to future actions. That process may affect action before it affects cognition or even without affecting cognition. Conversely, the results for the patients show that a cognitive grasp of the situation may be attained without affecting action.

The examples provided are intended to illustrate how much difference it will make to economic analysis of organization—and other topics in economics—if the shift continues from "deviations" to "how the brain works." The magnitude of the difference can be seen in the concluding remark of the Camerer–Malmendier review:

> Meso concepts like organizational routines and learning are not rooted in individual behavioral concepts. It is possible that organizational economics could be tied to meso-organizational (or macro-organizational) constructs, but this is not an active area of research.

From the deviations perspective on "individual behavioral concepts" this blunt summary is correct. Yet to an organizational theorist it looks odd, since the evolutionary economics community has produced hundreds of papers about firms stemming from the treatment of routine in Nelson and Winter (1982). Moreover, the performance properties of routine action have been linked to "how the brain works" (Cohen and Bacdayan 1994)—in this case to the distinction between procedural and declarative memory (Squire and Kandel 1999). This stark difference suggests that the visibility of research on routine action in organizations, and of many other important aspects of the day-to-day performance of firms, will depend crucially on which of the two contending perspectives future economic researchers have in mind.

References

Bechara, A., H. Damasio, D. Tranel, and A. Damasio. 1997. Deciding advantageously before knowing the advantageous strategy. *Science* 275:1293–95.

Cohen, M., and P. Bacdayan. 1994. Organizational routines are stored as procedural memory: evidence from a laboratory study. *Organization Science* 5:554–68.

Damasio, H., T. Graboski, R. Frank, A. Galaburda, and A. Damasio. 1994. The return of Phineas Gage: Clues about the brain from the skull of a famous patient. *Science* 264: 1102–5.

King-Casas, B., D. Tomlin, C. Anen, C. F. Camerer, S. R. Quartz, and P. R. Montague. 2005. Getting to know you: reputation and trust in two-person economic exchange. *Science* 308:78–83.

Nelson, R., and S. Winter. 1982. *An Evolutionary Theory of Economic Change*. Belknap Press of Harvard University Press.

Rizzolatti, G., L. Fogassi, and V. Gallese. 2001. Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Neuroscience* 2:662–70.

Squire, L., and E. R. Kandel. 1999. *Memory from Mind to Molecules*. New York: Scientific American Library.