

Organizational Design: A behavioral economics approach

Colin F. Camerer, Caltech

Version Monday, October 18, 2004

Reader note: TBA denotes material “To Be Added”

Chapter 1: Thinking like a behavioral economist

This book is about basic theories and facts underlying how companies organize workers to be more productive. An important part of this process is how top managers and boards govern those companies to accomplish various goals.

The ideas in this book are a new blend of economics, psychology and sociology.

The central approach is economic. Economic theories usually assume that people know what they want and make the best choices given their desires and information. (In technical language, we say “preferences are complete”—that is, each choice can be assigned a numerical “utility” and people are making the best choice, maximizing utility.)

Extended to organizations, maximization means that workers pick jobs which provide them the most happiness and financial reward. Managers anticipate what workers will do, and create companies which provide the most economic wealth to shareholders. Sometimes these good organizations are created by design; other times they come about through trial-and-error, or imitation-- one good design succeeds and is copied by other firms.

Psychology enters as a way of modifying the economic theories to respect the fact that people have natural limits on their abilities to make difficult calculations and resist temptation. People are optimistic about the future and often underestimate how long it will take to solve a hard problem. People also care emotionally about people other than themselves (e.g., they feel envy, guilt and moral obligation), which affects how they behave at work. Judgments about how valuable, numerous, or likely things are likely to be are also constrained by psychophysical properties and the nature of brain mechanisms.

Sociology reminds us that people are “socialized”: What we want is often influenced by what others have or want; how we behave is influenced by informal norms of proper behavior, as well as by formal laws and rules; and people are linked to others in

social networks which influence what they know and whose opinions and needs they care most about.

“Behavioral economics” is a relatively new approach to economics which takes into account facts and ideas from psychology and sociology to modify economic theories. The idea is to use facts and ideas from neighboring social sciences to improve economics while respecting the two stylistic principles which made economics successful—namely (i), to use mathematics to express ideas clearly and generate insight, and (ii) to explain phenomena in naturally-occurring situations, and give managers and governments advice.

This book is quite unusual by introducing behavioral economics concepts into ideas about organizational economics that have become sharply refined in the last couple of decades.

1. Why organizational rules and incentives matter

The essence of organization is the idea that if you take the same group of people with the same amount of money, skill, and time, the way their work is organized can make a substantial difference. The last couple of decades of thinking suggest that three part of the organization are crucial: **Decisions; incentives, and evaluation.**

Decisions refer to the assignment of who has the authority to make decisions (often called “decision rights”). The authority to make decisions shows us most sharply when there is a dispute— when one side wants A and another wants B, what actually happens? In many legal and political systems, for example, there is a clear system of decision rights (typically with checks and balances). In the US, for example, the President has the clear authority to nominate Federal judges. However, they must be confirmed by the United States Senate.

There is often an important distinction between nominal authority and real authority (sometimes called “formal” and “informal”). For example, in many families you would think the parents make decisions. But a screaming toddler often has the real authority because whenever the toddler screams loudly the parents do what she wants.

Incentives are the way in which evaluations determine pay and anything else that workers value—promotions, job titles, larger offices, parking spots and so forth. Keep in mind that money is a powerful incentive because everybody wants more of it, and it is a crisp numerical measure. But other incentives matter too (and pay may even erode or “crowd out” other kinds of incentives, like pride in a job well done, or the fun of being part of a team).

Evaluation refers to the system by which performance is evaluated. In well-functioning firms there is usually a formal component and an informal component. A formal component might be an annual review or measurement of how much a worker produced. The informal component is typically opinions, in the form of written letters or a meeting at which opinions are expressed.

Good organizations understand that these three components have to work together. The idea is to give the right to make **decisions** to people who have an **incentive** to make good decisions, and **evaluate** them so they are rewarded for good decisions and penalized for bad ones.

The last paragraph sounds so bland and obvious that you might wonder why you need a whole book, or course, in this topic. The short answer is that getting the assignment of decision rights, incentives, and evaluation— three different features-- to all fit together at the same time, and to change them in response to economic and cultural changes, is not so easy. It is a little like juggling three balls: A juggler has to constantly be passing two balls between her hands *rapidly*, while the third ball is in the air.

For example, a typical mistake in balancing decisions, incentives and evaluation occurs in multitasking, when a job requires somebody to perform more than one task (as is typical). Then incentives which reward one activity, but not another, often lead to people overproducing the incentivized task and not doing enough of the task they aren't paid for.

Another fact about human nature that makes organization difficult is that people are different. Some people like responsibility in their job and like having decision rights. Others get stressed by difficult decisions and prefer that others have decision rights.

Different people are motivated by different incentives. In this book, the default meaning of “incentive” is money. However, *do not* make the mistake of thinking that money is the only incentive people care about, or is the strongest. In wartime, people risk their lives for combat pay, but also for national pride, the camaraderie of engaging in a history-making enterprise, or out of a principled sense of justice. In some cultures, shame is such a powerful incentive that people commit suicide rather than face shame, or kill their own relatives (typically sisters) if they have brought shame upon the family for actions that other cultures would completely forgive (such as being raped, or marrying for love).

Entrepreneurs often describe their goal as providing a fulfilling place to work, which creates a wonderful product that makes people's lives better. These entrepreneurs see financial earnings as the market's way of rewarding them for making a great product, and for taking a large risk to do so.

Because people are different, the right combination of decision rights, incentives, and evaluation will also depend on the people who work in the organization. This raises an important challenge of *sorting*—finding the right people for the right jobs.

Organizational meltdowns

One way to appreciate why organization matters is to study cases where organizations failed, even when they are composed of talented, energetic people. (We will see many throughout this book, and we will study successes as well.) Often these stories remind us of the difficulty of coordinating decision rights, incentives, and evaluation, especially when people care about different goals.

For example, many adventure books describe how physically fit, mentally tough adventurers band together to accomplish some goal. In Shooting the Boh (1992), Tracey Johnston describes a white-water rafting trip to a very difficult river in Borneo which white men had never rafted before. Their expedition endures unbearable hardship in the form of tropical jungle diseases and the physical challenge of rafting a steep river surrounded by huge rocks on both sides through many of its swiftest rapids. In Running the Amazon, Joe Kane describes a kayaking trip from the very headwaters of the Amazon river, on a mountaintop in Peru, through the entire Amazon to the Atlantic Ocean. (His book ends with one word—"salt"—, what he tastes on his fingers as he dips them in the ocean, after paddling in the fresh water river for months.)

Both adventures end on bitter notes. (Keep in mind that these are the *successful* adventures-- everybody survived, and the trips were worth writing books about. Presumably the unsuccessful expeditions faced even bigger organizational challenges.)

Johnston's Borneo group bickers over physical hardship. Stress is caused by the fact that ill-prepared travelers end up borrowing routine items from others, like dry socks, which become extremely valuable in the wet rainforest. The heroic river guide, who single-handedly helps the group overcome many rafting and logistical obstacles (e.g., finding food) ends up romantically involved with a beautiful ex-model who is on the trip. Others come to think that he favors the model at their expense, which causes jealousy and tension.

Kane's Amazon trip includes an Olympic-level paddler, and a businessman who raised money to have part of the adventure filmed. The trip goes slower than planned and the group begins to run out of money. The Olympian is frustrated by the slow pace of the others. The businessman, who is the slowest paddler, holds the group back but feels no guilt because his fundraising (for the planned film) financed much of the trip. Only Kane and one other paddler finish the entire trip.

Adventure trips like these are special organizations in two respects: First, they are temporary or "instant" organizations created once, to accomplish a special challenge. This means that people often go into the trip not knowing much about the goals and skills of others. It also means that the threat of not doing business with somebody in the future if they don't live up to expectations, which is often a strong incentive in long-lived businesses, has little power.

Second, the expeditions require people to travel together, and are physically demanding. These trips are like a chain whose strength depends on its "weakest link". If one person complains a lot, eats too much food, sleeps too late, or paddles too slowly, the entire group is demoralized, hungry, or slowed down.

Organizations like these are especially vulnerable to design mistakes. One design mistake reflected in these adventures is pure optimism: The adventurers underestimated how difficult their tasks were. This is common in business and government as well: Large building projects, for example, usually take twice as long as planned and cost twice as much. [CITE TBA]

Another design mistake is failing to recognize how differently people are motivated. For example, the Borneo trip's guide is certainly motivated by pay and pride.

But on this particular trip, he was also motivated by personal feelings toward the ex-model. Either the guide could not put these feelings aside— and no evaluation system was in place to discipline him—or, more likely he felt he could fulfill his obligation to the group while still enjoying his romantic time with one of the people in the group.

In simple economic terms, the issue with the river guide was the extent of his working hours—during downtimes, is he still “on the clock” and obligated to deal with the group’s endless suffering, or is he “off the clock” and entitled to rest and relaxation of his own. The group thought they had the decision rights to allocate the guide’s spare time, and the guide thought *he* had the decision rights.

GEEK ROULETTE STORY HERE TBA



A “eudaemonic” shoe

There is a common theme in the stories about the geeks playing roulette, and the adventure travelers. In both cases, groups are relatively good at solving enormously complicated *technical* problems which seem almost impossible. The geeks could guess where a roulette ball is likely to land, and the rafters could navigate daunting physical challenges posed by dangerous river rapids and huge distances. But in the end, the *social* challenge of organizing their activity as a group is what limited their achievements. These stories remind us of the remarkable capacities of the human brain and body, but also the important challenge of organizing ourselves to reach a common goal. This challenge is what motivates the ideas in this book and their application to running business organizations.

Fraud: Burning down the house

A useful way to appreciate the importance of balancing decisions, incentives and evaluation is in studying large-scale frauds and scandals, which often have a huge

financial cost to firms (and entire industries, through reputational spillovers) far beyond the gains to the people who perpetuated them.¹

Let's see a few examples and then think about what lessons can be learned from them.

Daiwa: How to lose a billion dollars...slowly

Toshihide Iguichi went to work for Daiwa Bank in 1977. Daiwa was a large bank with \$200 billion in assets and \$8 billion in reserves (as of 1996) and a growing trading operation. For seven years, Iguichi worked in the back office of Daiwa's growing New York government bond trading operation, keeping track of the enormous amount of paperwork involved in reconciling accounts of a trading operation. Iguichi became a trader in 1984—authorized to buy and sell government bonds for customers, and for the Bank's own account—but Iguichi also kept his job supervising the “custody” department (which kept track of who owned the bonds). Daiwa's custody account was administered via a “sub-custody” account at Bankers Trust. Daiwa and its customers kept track of activity in their accounts through reports from Bankers Trust. The reports were filtered through Iguichi.

Early in his trading career, Iguichi lost a few hundred thousand dollars trading bonds. He decided to sell other bonds in the Bankers Trust account to cover his losses, but falsify the reports so nobody who saw the reports realized the bonds had been sold. As he lost more and more money trading, Iguichi falsified more and more Bankers Trust reports to disguise transactions he had made to cover the trading losses. The false reporting persisted because Daiwa's internal auditors never checked Iguichi's false reports against Bankers Trust's own records. Presumably they trusted Iguichi because he had worked in back-office operations for many years before beginning to trade.

Iguichi's fraud lasted not one year, or two...it lasted 11 years. During that time, sold \$377 million of customers' securities, and \$733 million of Daiwa's securities—a total of \$1.1 billion. To cover these fraudulent trades, he forged 30,000 trading slips, an average of 10 each day.

In fact, Iguichi was never actually *caught*. He confessed. On July 13, 1995, he sent a 30-page letter to the president of Daiwa in Japan explaining what he'd done. He said he confessed because “after 11 years of fruitless efforts to recover losses, my life was simply filled with guilt, fear and deception”.² He got tired of covering up his bad trades, and of waiting to be found out.

The aftermath of Iguichi's massive, long-lasting fraud is just as remarkable. Daiwa waited a month to tell the Japanese Ministry of Finance about the losses, on August 8. Despite a clear legal requirement to report the fraud to U.S. regulators, Daiwa actually waited another month to tell the U.S. Federal Reserve Board. (The finance minister, Masayoshi Takemura, later denied that his Ministry had failed in its reporting duties, made a veiled apology to the U.S. Treasury Secretary, then later denied apologizing.) During September, Daiwa higher-ups took Iguichi's advice to keep the

¹The website at <http://www.erisk.com/Learning/CaseStudies.asp> has some good case studies, two of which are excerpted here.

² http://www.time.com/time/magazine/1997/int/970210/interview.i_didnt_set.html interesting interview

losses secret until “appropriate measures” could be taken to ensure stability after the news broke. Iguichi was told to pretend he was on vacation when he was actually hiding out in the apartment of a Daiwa manager in New York trying to reconstruct the history of his long-lasting fraud.

Daiwa finally told the Feds what had happened, in mid-September. Unhappy about being left in the dark, FBI agents interviewed Iguichi on September 23. Shortly after that, Iguichi was arrested and the bank itself was indicted on 24 counts including conspiracy, fraud, and failure to disclose federal crimes.

It then became apparent that Daiwa’s New York trading operation had a long history of rogue operation which was invisible to Daiwa’s top managers back in Japan. The bank had operated an unauthorized trading area for seven years, from 1986 to 1993; they even disguised as a trading room temporarily as a storage room when regulators stopped by. Iguichi also said that from 1984 to 1987, other New York traders had created major losses which were concealed from regulators by shifting the losses to overseas affiliates.

In November, 1995, the Fed ordered Daiwa to quit doing business in the U.S. A month later, it sold its U.S. assets and 15 offices to Sumitomo Bank for \$3.3 billion. In February, 1996, Daiwa agreed to a \$340 million fine to settle the indicted charges—a record for a criminal case. Many senior executives resigned or took early retirement. Iguichi himself was sentenced in December 1996 to four years in prison and a \$2.6 million fine. The final sting was an order from a Japanese court requiring 11 current and former board members to pay the bank \$775 million in damages, as compensation to the bank’s shareholders and as a penalty for the managers’ failure to report the incident promptly.

An important cost from this kind of fraud is an “externality”—the loss of reputation of people and institutions associated in people’s minds with Daiwa. Even if the fraud was the work of a single rogue (and unskilled) trader, if the capital markets think it is the tip of an iceberg, and other frauds are going on undetected, that can harm capital market credibility. Indeed, before the scandal, Japanese banks were charged a .1% premium above the London Interbank Offer Rate (LIBOR, an international benchmark for interest rates). Afterwards the premium went up to .25%. Daiwa’s corporate mistake therefore tainted other Japanese corporations and cost them huge sums in extra borrowing costs. One analyst said the problem was that the Japanese Ministry of Finance was “severely compromised” as a regulator, making it “difficult to be reassured that this is not going to happen to someone else.”³ The markets therefore perceived Daiwa’s fraud as something that could be going on in other banks, that would not be sharply regulated by the Finance Ministry, which led them to charge an extra premium just in case.

What can we learn from the Daiwa case? The basic one is that Daiwa did not have decision rights, incentives, and evaluation balanced properly. Iguichi had the right to make trades independently, but was also in charge of the back-office reporting of account gains and losses, which provide numbers to evaluate the performance of his trades and those of others. Iguichi was in charge of supervising himself.

³ <http://www.asiaweek.com/asiaweek/95/1027/biz2.html>

Keep in mind that people who are in charge of their own evaluations do not *always* cheat. The sorting process, assigning traders to small offices where they run their own back-office operations, could work fine if you can find enough traders who are both talented traders and honest back-office worker bees. But as one source wrote [find source], “Bankers who hire money hungry geniuses should not always express surprise and amazement when some of them turn around with brilliant, creative, and illegal means of making money”.

Furthermore, Daiwa’s actions upon learning of Iguichi’s fraud made matters worse, and perhaps reflected a culture of tolerating or covering up scandal which licensed Iguichi to start covering up his own losses in the first place. Daiwa managers, and the Finance Minister, seemed incapable of taking quick action or admitting fault, until a lot of diplomatic and economic damage was done.

A behavioral economics question is: What motivated Iguichi to commit his crimes? (Understanding his motivation might help us prevent situations like these or know how to sort for employees who are both talented and unlikely to commit fraud.) In a 1997 interview conducted in prison, Iguichi said his early actions did not feel like a crime. “To me,” he said, “it was only a violation of internal rules.” He then explains that his actions were driven by a deep desire to erase trading losses—a pattern psychologists call the “break-even effect”.⁴ As Iguichi described it: “I think all traders have a tendency to fall into the same trap. You always have a way of recovering the loss. As long as that possibility is there, you either admit your loss and lose face and your job, or you wait a little – a month or two months, however long it takes”.

In fact, most traders *do not* fall into this trap, or certainly not for as long as Iguichi did. The ability to resist “chasing losses” by taking bigger risks, or covering up the losses with a back-office fraud, and making peace with losses, may be one of the traits that distinguishes successful traders from those with perennial losing records, like Iguichi.

Barings: TBA

⁴ Loss-aversion discussion here. TBA



Figure: Nick Leeson (seated at right), the man who destroyed Barings Bank, talking to the press after his release from prison in Singapore in July, 1999.

Most of these frauds have three ingredients: Poorly-designed rules which enable the fraud to go on without being revealed; high stakes which make the fraud worthwhile; and a moral willingness or psychological motive to commit the fraud. Often, the guilty party blames others as if he just sees the fraud as winning at a game.

An interesting question is: Are these frauds common or rare?⁵ Assume that the *marginal cost* of monitoring rises dramatically with the probability p of detecting a fraud—mathematically, that is the same as saying that it is expensive to track down the most wily and clever fraud (or the marginal cost of raising the probability $c'(p) \rightarrow \infty$ as $p \rightarrow 1$)— then there will be *some* frauds even if firms are optimizing.⁶

⁵ CFC NOTE TO SELF Joe Jett case

<http://www.smeal.psu.edu/faculty/huddart/Courses/BA521/Jett/JettNYTimes97.shtml>

(note starbuck quote) http://money.cnn.com/1999/12/29/investing/century_greed/

⁶ This is why economists think policies which are intended to reduce the number of violations of some rule to zero, are noneconomic pipe dreams. The *rhetoric* of zero-tolerance may prove politically useful, or may motivate people to get closer to that goal, but it is not the result of optimization if the marginal cost of detection skyrockets when you try to catch every last user or terrorist.

An interesting question is how widespread such frauds are and whether firms are taking all the right steps to prevent them. No monitoring, enforcement, or regulatory mechanism is perfect. In economic language, the marginal cost of reducing fraud to zero is just too high, so even when firms balance costs and benefits optimally, there will be an optimal amount of fraud.⁷ It is difficult to show that firms are *systematically* misaligning decisions, incentives and evaluation without a full model showing the costs and benefits of mistakes.

TBA

2. Building blocks of economic analysis

Before proceeding, let's spend some time on the most basic concepts of economics. These form a foundation for eventually thinking about workers in companies.

Marginal analysis

In economics, we start with the presumption that people will perform an activity until the *marginal* benefit of continuing is equal to the marginal cost of continuing. This way of thinking is loosely called “marginal analysis”.

Marginal analysis is also important at the market level: If prices are determined by the intersection of supply and demand, then prices are not determined by the person who values a good most, or the seller who can make it most cheaply: Instead, prices are determined by the *marginal* buyer and seller.

A famous example is water and diamonds. Suppose there was a massive shortage of fresh water on the earth. Because water is necessary for life, the price would skyrocket. (In fact, after natural disasters like earthquakes and floods, water and other valuable staples, and useful items like batteries, candles and shovels, often do rise sharply in price.) So why is water so cheap? The answer is that water plentiful—sometimes you can even collect your own in a bucket for free, when it falls from the sky. Developed economies can purify and distribute clean water on a massive enough scale to spread the fixed costs across a large market. Thus, the marginal cost of supplying one more gallon of water is quite low. Since the water business is usually competitive (and the largest part

⁷ An analogy which can help us think about the optimal amount of fraud is “shrinkage”—the polite term for employee theft and customer shoplifting—in retail stores. A typical shrinkage rate is about .5% of inventory. In some businesses, like groceries, profit margins are only 1% of sales, so shrinkage may do a lot of damage to the bottom line. But reducing shrinkage to zero is just too costly. The lowest rates among retail chains are only on the order of .1%.

is regulated as a public utility), prices are low too.⁸ So even though water is necessary for life, it is not very valuable “at the margin”.

An opposite case is diamonds. While diamonds may be a girl’s best friend, no girl ever died if she went without wearing them for a few days, as she would if she had no water. So why are diamonds so expensive. The answer is that the supply of diamonds is a classic monopoly. A large portion of the world’s diamonds are controlled by the DeBeers corporation, which carefully controls the supply to keep prices high. Furthermore, diamonds are rare and harder to mine than water is to pump and purify.

Constrained utility maximization

The central modeling tools in the economics of consumer behavior are the concept of utility, budget constraint, and (combining the two) maximization of utility subject to constraint. These ideas form the mathematical underpinning of a “demand curve”—a function that tells you the number of units of demand for a product $D(P)$ if the product’s price is P .

A product or activity’s “utility” is a number which reflects how much people seem to anticipate liking the product when they make choices. (Some nuance to this definition will be offered below.) There is much confusion about the concept of utility. Here’s how you should think about it: Suppose you sat down and watched a bride and groom and a wedding planner, planning a wedding. Suppose each wedding “menu” has a certain number and type of flowers, a venue where the wedding is held, and a food menu for dinner. As the wedding planner leafs through a thick book listing these (flowers, venue, menu) “bundles”, the bride says “I like this one more than that one” many, many times. We call this list of statements of what she likes a “preference relation”, often written mathematically in the notation $A \geq B$ (read “A is weakly preferred to B”, or $A \sim B$ if A and B are equally preferred). Suppose that her preferences are “complete”, so that for any pair of (flowers, venue, menu) triples she can decide that she likes one, or that she is truly indifferent between them. Suppose further, that her preferences are transitive—if she likes A more than B, and B more than C, she also likes A more than C.

If the bride’s preference relation is complete and transitive, then we can keep track of what she likes by assigning numbers to each triple of (flowers, venue, menu). Suppose there are N such triples. Just assign N points to the one she likes most, $N-1$ points to the one she likes second-most, and so forth. These point assignments are “utilities” which are **revealed** by her expressed preferences. The utility numbers just equate the statement “I would rather have $A=(\text{roses, Peninsula Hotel, seafood})$ than $B=(\text{tulips, Malibu beachfront, filet mignon})$ ” with the statement “A has a higher utility

⁸ Keep in mind that prices depend on cost and how competitive a market is. Suppose there was a water monopolist who controlled all the water and could charge whatever they would like. Because water is such a necessity (we say its price elasticity—the percentage change in demand relative to the percentage change in price—is low) the monopoly price would be high. (Generally, when a good is price-inelastic, a monopoly’s profit-maximizing price is high). Thus, water is cheap in developed economies because it can be mass-produced cheaply, and also because the supply side of the market is usually competitive (or regulated to create low prices).

number than B”, or “ $u(A) > u(B)$ ” (where $u(x)$ is the function which assigns a number to x).

Notice that we are not insisting that the bride actually consciously has numbers in mind when she expresses preferences (although people may do so in unusual cases). But if she can always make up her mind between two bundles, and does so transitively, she is acting *just as if* she is would act by consulting a list of numbers, and picking the bundle with the higher number.

The important part of this utility-maximization process for economics is the concept of **tradeoff**: Assuming the bride has a limited budget, she must decide whether she is willing to spend less on flowers in order to have more hors d'oeuvres served at the reception. We call the amount of one good which is worth giving up to get more of another the **marginal rates of substitution** between the two goods.

In most choices people actually face, there are two other components which are missing from the simple sketch above: **Risk**, and **time**.

Usually we are not sure what the consequences will be when we make a choice—the choice is **risky**. For example, if the wedding is planned for a Malibu beachfront, there is a small chance that it will rain and the wedding must be moved inside. So the bundle (tulips, Malibu beachfront, filet mignon) is really a probabilistic series of different realized bundles. In traditional economics we usually assume that people value these probabilistic “gambles” by taking a probability-weighted sum of the utilities of the possible consequences. That is, if a gamble G has n different possible outcomes, denoted x_i , and each has probability p_i then $u(G) = \sum_{i=1}^n p_i u(x_i)$. This is called the “expected utility” rule, and it will play a substantial role in thinking about employment contracts where a worker’s income can be high or low because of random factors out of their control (so that valuing different jobs for their income potential is like choosing a gamble).

Most practical choices also yield a stream of goods or activities which unfold over **time**. This raises a question of how people weight the utility of good and activities they get at different points. The standard theory is called “discounted utility”. The idea is that people weight utilities which they receive one time period from now by a discount factor, δ . Utilities received t periods in the future are weighted by δ^t . Therefore, the discounted utility of a stream of outcomes received at points $t=0, 1, 2, \dots, T$ in the future is $u(x_0, x_1, x_2, \dots, x_T) = \sum_{t=0}^T \delta^t u(x_t)$.

Behavioral economics: Limits on rationality, temptation and self-interest

The simple model of utility maximization above, on which almost all modern economic theories rest, has proved to be a very powerful tool for explaining many features of our economy, for giving people advice, and for engineering new rules and systems. However, because the economic model is a mathematical simplification of a complicated cognitive process, a more careful study of human psychology complicates the simple picture of utility maximization (and the extensions to including risk and time delays).

“Behavioral economics” is an approach to economics which tries to use empirical regularities and ideas in psychology (and to some extent other fields) to improve economics mathematically, to make it more empirically accurate and more useful (Camerer, Loewenstein and Rabin, 2004). Here I will just give a quick sketch of some of the basic ideas that will be put to use later in this book.

The basics... TBA

Some day, our understanding of these limits will be replaced by a detailed understanding of the brain processes which create human behavior. (The subfield that strives to do this is called “neuroeconomics”.⁹)

Property rights

A “property right” is the right to use or sell property. It is often helpful to distinguish “use rights”—the right to use the property—from selling rights (or “alienable” rights). In America, you have a use right to your kidney but not a selling right (it’s illegal for you to sell it, though you can donate it). If you rent an apartment, you have a use right but not a selling right.¹⁰

Some property rights are enforced by social convention. For example, in Chicago it snows a lot in the winter, but many people do not have garages and park their cars on the street. As a result, people often spend laborious hours shoveling snow away from their car after a big snowfall, so that they can drive the car to work. It is common after digging out your car to put a chair or traffic cone into “your” empty space (sometimes with a “do not park here” sign). These markers are understood by Chicagoans to establish a property right to the empty space. The idea is that when a person arrives home from work, “their” space is still available.

The system incentivizes people to dig spaces out. Those people who do not mark their spaces with a chair or cone are “donating” the space to whomever needs it. The social norm works amazingly well, despite the fact that it is not backed by legal enforcement. If you removed a cone and took somebody’s parking space, the police could not issue you a ticket for parking illegally or “parking space theft”. (However, if the police were called they would probably berate you for “stealing” the space and suggest you park elsewhere.) The system is also enforced by vandalism— if you park in somebody’s spot, there is a chance your car would be scratched by a key or even have its tires slashed.

In developed economies, clear property rights are taken for granted and are enforced by the “rule of law”— a system of clearly specified legal rules, policing, access to legal representation, and fair judges. Poor enforcement of property rights can lead to

⁹ See Camerer, Loewenstein and Prelec (2004a,b).

¹⁰ Can you think of an object that has a selling right but not a use right?

underinvestment, because firms are afraid their machinery will be confiscated or stolen, and low consumer demand for valuable goods others can easily steal.

The damage done to an economy by weak enforcement of property rights can be seen most clearly in less developed economies, in occasional lapses in developed economies (like looting during urban riots when police are overwhelmed), or during dramatic economic transitions.

For example, some small-scale societies practice what some anthropologists call “tolerated theft”. Tolerated theft is a norm of sharing community property, in which the property “owner” is not obliged to offer, but the “owner” cannot turn down a request to share. For example, suppose you spend a couple of hours picking a large basket of berries in the heat, and head down the road toward our village. If I see you and grab a handful of berries out of your basket, and you don’t stop me, then you’ve tolerated theft. Tolerated theft means that your private property isn’t just yours. This practice, in turn, reduces the private incentive to gather anything valuable, if it takes work to do so (like hunting, fishing, or gathering). Not surprisingly, small-scale societies with tolerated theft are often poor. In some of these societies, one of the most valuable goods is a pair of pants (which, thankfully, cannot be shared). Pants with large pockets are particularly treasured. Why? In a society with tolerated theft, pants with pockets *are like a safe deposit box*—they are a place in which goods can be hidden and protected from theft.

Poaching on property rights in a way that is akin to tolerated theft can occur at much higher economic scales too. In 2004, Nigeria was the world’s seventh-largest oil exporter. There, Alhaji Mujahid Dokubo-Asari openly admits that his homegrown militia of up to 2,000 armed volunteers siphon off hundreds of barrels of crude oil from pipelines operated by Royal Dutch Shell, refine it, and sell it cheaply to locals (LATimes, 2004). His group do not see themselves as stealing somebody else’s oil. They say they are simply reclaiming some of the oil that should belong to the people, but which was given away by a corrupt elite who he accuses of “eating the money” from the oil. “As far as we are concerned,” Dokubo-Asari says, “it is the Nigerian state that is stealing oil”. He says he is being pursued by the state for criticizing the government about electoral ballot fraud. The government says Dokubo-Asari is wanted for gangland killings involving another outlaw, Ateke Tom; but Dokubo-Asari says Tom was hired by the government to kill *him*. A Shell-commissioned report concluded that thieves annually took \$1.5-4 billion of crude oil illegally (though less interested estimates are much lower). When he is asked how much crude oil his militia siphons from the pipelines, Dokubo-Asari says “As much as we can. It’s free.”

Economic transition in formerly-Communist countries in Eastern Europe, and in the former U.S.S.R., has been erratic and spotty. One problem is that poor property rights enforcement gives firms an incentive to invest heavily in private security (as a substitute for publicly-supplied security in the form of policing and legal enforcement of contract; e.g., Grief and Kandel, 1995). This gives a competitive advantage to firms that are Mafia-connected; as a result, many observers are shocked by the deep influence of various organized crime mafias on legitimate Russian business.

Another problem in economic transition is “deprivatization”—when governments change the rules after a privatization. After the collapse of the USSR, about 70,000 state-

run companies were privatized (i.e., sold by the state to a private owner or syndicate). Some of the new owners were investors from Europe and the United States. For example, the LBO firm Kohlberg Kravis Roberts teamed up with an enterprise called the US-Russia Investment fund and bought a majority interest in the Lomonosov Porcelain Factory in St. Petersburg in 1998 (Business Week, 1999). But in October 1999, a Russian court stripped the American owners of their majority interest by declaring the company's initial privatization several years earlier had been illegal. The problem is that judgments of illegality often hang on the tinnest of legal threads or judicial whims—like a missing piece of paper in a document file, or a failure to meet a requirement by an obscure deadline. The newly-deprivatized company was rumored to be sold to its former Soviet-era managers who were due to be “downsized” when the Americans took over.

Figure ? below shows the powerful effect of property right protection on economic growth during transition (Hoff and Stiglitz, 2004). The y-axis shows the ratio of GDP in the year 2000 to 1989 GDP. The x-axis is the percentage of survey respondents in different countries who *disagreed* with the statement “I am confident that the legal system will uphold my contract and property rights in business disputes”. A high percentage (right hand side) means a lot of respondents *do not* think their property rights are protected. (This measure is also correlated with other measures of the “rule of law”, such as a 0-10 index constructed by the Wall Street Journal.)

Each circle in Figure ? is a different post-Communist society. In all those countries which grew in GDP (the y-axis ratio is above 1, Poland, Slovenia, Hungary, and Slovakia) less than 40% of respondents were insecure about their property rights. In the countries with the most insecurity (highest disagreement percentage, which is furthest to the right, such as Ukraine, Macedonia, and Russia) GDP actually shrank over the 11 years by half or more.




FIGURE 1. GROWTH AND PROPERTY RIGHTS INSECURITY IN 20 TRANSITION ECONOMIES

Note: The index of insecurity is the fraction of respondents to the survey in 1999 who disagreed with the statement: “I am confident that the legal system will uphold my contract and property rights in business disputes.”

Sources: Survey data: <http://www.worldbank.org/wbi/governance/beepsinteractive.htm>, and for documentation, Joel S. Hellman et al. (2003). GDP data: EBRD (2002).

Land resettlement in Zimbabwe [TBA]

Complications with property rights can also arise when there are *too many* layers of property rights which are not clearly prioritized. A nightmare version of this problem was seen after the devastating 1994 earthquake that struck Kobe, Japan. Landlords, landowners, tenants and subtenants all had various legal rights over decisions about property they owned or rented. The Japanese decided that everyone with any rights over property had to agree to a rebuilding place before construction could begin. But the thicket of property rights made the process very slow. One city block had 303 different renters, lessees, and subletters. Two years after the quake, despite a \$30 billion rebuilding fund, only 4 of 100 “construction teams” had actually begun doing any new building (Wall Street Journal, 1996).

SIDEBAR: Disputed property rights and law: The Barry Bonds 73rd Home Run Ball¹¹

¹¹ Ian Krajbich wrote the background material for this section.

On October 7th 2001, Barry Bonds [pictured below] hit his 73rd home run of the year into the right field arcade of PacBell Park in San Francisco, setting the Major League Baseball (MLB) record for most home runs in a single season. The MLB single season home run record is arguably the most prized record in sports; as a result, the record breaking ball was highly sought after by sports enthusiasts (and was guessed to be worth a million dollars).



Left: San Francisco Giant Barry Bonds. Right: Patrick Hayashi, photographed by Josh Kepler after picking up the disputed Barry Bonds 73rd home run ball.

As the baseball soared over the right field fence, dozens of fans scrambled and fought for a chance to catch the ball. The first fan to get his glove on the ball was Alex Popov. But Popov was quickly “buried face down on the ground under several layers of people...[and was] grabbed, hit and kicked”¹².

Another fan, Patrick Hayashi [pictured], was also “involuntarily forced to the ground... While on the ground he saw the loose ball. He picked it up, rose to his feet, and put it in his pocket”. Hayashi motioned to a photographer, Josh Keppel, to point the camera to him, which Keppel eventually did when someone else in the crowd asked him to. Then Hayashi “held the ball in the air for others to see. Someone made a motion for the ball and Mr. Hayashi put it back in his glove.”

Who owned this valuable ball? This case illustrates how property rights can be under dispute, even in a developed economy with clear legal rules. As the Superior Court verdict noted, “While there is a degree of ambiguity built into the term possession, that ambiguity exists for a purpose. Courts are often called upon to resolve conflicting claims of possession in the contexts of commercial disputes... Because each industry has different customs and practices, a single definition of possession cannot be applied to different industries without creating havoc”. Indeed, different legal rules of possession have emerged differently in different businesses. For example, in hunting and fishing, “the hunter acquires possession upon the act of wounding the animal not the eventual capture”.

¹² This quote, and all others in this description (unless noted otherwise), is from the December 18, 2000 Superior Court verdict at <http://www.court tv.com/trials/baseball/docs/BondsVerdict.PDF> . See also http://www.court tv.com/trials/baseball/ballsold_ctv.html

On Oct 24th, Popov sued Hayashi for ownership of the ball, claiming that it had been stolen from him. Popov and Hayashi both thought they were the rightful owners and would prevail in the case.

According to the law, prior to the time the ball was hit, it was owned by Major League Baseball. At the time it was hit, the ball became ‘intentionally abandoned property’. The first person who came in possession of the ball owned it.

But who *possessed* the ball? Popov was the first to *touch* the home run ball, but he lost it in the crowd. The Court defined possession like so: “The central tenant [sic¹³] of Gray’s Rule is that the actor must retain control of the ball after *incidental* contact with people and things” [note *emphasis*] and thereby concluded that Popov had not possessed the ball.

Hayashi argued that Popov did not possess the ball because he did not have “complete control” of the ball. As a law professor testified, “the first person to pick up a loose ball and secure it becomes its possessor.” That person was Hayashi.

But wait. The contact between Popov and the fans nearby was clearly not “*incidental*”. It is not as if Popov bumped into a fan and dropped the ball, which would then be legally possessed by whomever picked it up. The fans mobbed Popov. A major purpose of law is to protect its citizens and to discourage reckless and dangerous behavior. Denying Popov possession of the ball because of the illegal actions of the fans who mobbed him would set a dangerous precedent by encouraging violent behavior at ballgames.

As the Court ruled: “The reason we do not know whether Mr. Popov would have retained control of the ball is not because of incidental contact. It is because he was attacked. His efforts to establish possession were interrupted by the collective assault of a band of wrongdoers . . . Mr. Popov should have had the opportunity to try to complete his catch unimpeded by unlawful activity. To hold otherwise would be to allow the result in this case to be dictated by violence. That will not happen.”

The judge acknowledged that Popov’s chance to control and possess the ball was undermined by the unruly San Francisco mob. At the same time, the judge noted that awarding the ball to Popov would be unfair to Hayashi because Popov could not prove that he *would* have caught the ball even if the mob weren’t surrounding him. Popov’s fumble established a “pre-possessory interest” in the ball. But, in the court’s words, that interest “does not establish a full right to possession that is protected from a subsequent legitimate claim” (like Hayashi’s).

Taking into account both interests—neither should win, but neither should lose--, the judge ruled that both men had equal rights to the ball and should evenly split the proceeds from selling it. When the ball was actually auctioned off, it sold for only \$450,000, roughly half of its originally estimated value.

After the sale, Hayashi and Popov were both disappointed with the sale price. Popov hinted that his legal debt was in the hundreds of thousands of dollars, because he had paid his lawyer by the hour. (Whoops.¹⁴) Hayashi had cut a deal with his lawyers for

¹³ The Court meant “tenet”.

¹⁴ Suppose you hire someone to paint your house and you pay by the hour. Your neighbor hires somebody and pays a lump sum for the whole paint job. Whose house do you think gets painted more rapidly? More carefully?

a contingency fee based on a percentage of the ball's sale price, so he came out ahead compared to Popov.

In 1960 Ronald Coase (later awarded the Nobel Prize in 1991) wrote a remarkable paper suggesting that if two parties can easily bargain to an agreement, which party has a property right doesn't matter for social efficiency—because the party who values the property right most will bargain to get it.¹⁵ The Barry Bonds disputed ball-possession case should be a perfect example of the Coase theorem in action: Hayashi had the ball in his possession. If it really was more valuable for Popov—either he wanted it more, or thought he could sell it for more-- then Popov should have paid Hayashi to turn it over. But if Hayashi wanted it more, or thought it was more valuable, there would be no price Popov could offer. Either way, the person who valued the ball most should have ended up with it— the only question is whether was whether Popov would have to pay.

Looking back, neither Hayashi nor Popov valued the ball just to keep it in a glass case. They wanted to sell it. They should have agreed to sell it and split the money. They would have saved themselves enormous lawyer's fees, which probably chewed up half or more of the ball's \$450,000 value at auction.

This case illustrates why the Coase theorem can fail. The crucial assumption in the Coase theorem is that people can cheaply reach an agreement by bargaining. Then it doesn't matter who has the property right (“ownership”) to begin with; after bargaining, the person who values it most will end up with it.

The problem, of course, was that Popov and Hayashi each thought they had the morally rightful claim to the ball, and also thought they would prevail in court and get all the money. The Coase theorem can fail if bargaining fails, because both sides think they are more likely to win than they really are, so they are willing to tolerate a costly delay in bargaining or risk an expensive litigation process. Psychologists call this kind of collectively-inconsistent belief “self-serving bias”.

END SIDEBAR

Specialization, gains from trade, comparative advantage

People are good at different kinds of work. In an ideal economic system, people specialize in what they are best at and enjoy. (Those two things can be different of course.) Specialization is the engine of economic gain because it is easy to show that if everyone can specialize in what they are best at (as judged by the tastes of consumers), and trade occurs freely, the system makes the largest quantity of products everyone wants. Specialization is essentially the same as a **division of labor** in which tasks are assigned to those people who are best able to do them.

When specialization is allowed to flourish, economic systems benefit when people have diverse skills as workers, and tastes as consumers, because the value of dividing labor is higher. An example that illustrates this principle is the television show, “The Amazing Race”. In the show, pairs of people are required to complete a series of

¹⁵ Of course, the property right assignment matters for who gets the most from the deal, because if the property right is assigned to the lower-value owner, then that person can get paid by the higher-value party. If the higher-value party gets the property right, then they don't have to pay.

physical and adventure tasks— like riding a board across a 15-foot waterfall, or finding an envelope with a clue in a dark cave filled with bats-- while racing around the world. The teams which complete the tasks fastest win prizes. Some tasks require brainpower, others require brawn, and still others require bravery. Generally, there are two big challenges for the teams: (1) Coordinating their activity, so they don't waste time bickering or arguing about what to do; and (2) having at least one person in the team who is particularly good at each task.

One season's show featured a pair of identical twins. Is being in a team with your twin a good or bad idea, if you want to win "The Amazing Race"? For coordinating with your teammate, working with your twin is probably good— twins share all their genes, and probably think and act alike on many dimensions. On the other hand, a pair of twins will have less diversity in abilities, and are poorly equipped to have at least one teammate who is a specialist in each type of task. In fact, the twins—two athletic men— placed in about the middle of the competition, perhaps because their skills were *too similar* so the benefits of a division of labor and specialization were too small.

The benefits of specialization can be illustrated by a joke: In Heaven, all the chefs are French, the car mechanics are German, the bankers are Swiss, the police are British, and the lovers are Italian. In Hell, the chefs are British, the car mechanics are French, the bankers are Italian, the police are German, and the lovers are Swiss.

The advantage of specialization is driven by **comparative advantage**. A worker's *absolute advantage* is her productive capacity in a certain activity. A worker's *comparative advantage* is her relative absolute advantage—that is, her absolute advantage relative to other workers. The point is that even workers who are not good at any activity in absolute terms have some comparative advantage because there is always an activity they are relatively less bad at. For example, a father making sushi is faster than his 10-year old at making proper sushi rice, cutting nori seaweed, buying sushi-grade fish, making sushi rolls, and setting the table for dinner. But the 10-year old is not so bad at setting the table, and is hopeless at everything else. So the 10-year old's comparative advantage is in setting the table.

It is useful to think about comparative advantage in terms of opportunity cost. An activity's **opportunity cost** is the value of the opportunity that is lost by doing that activity versus another. If a friend gives you Lakers playoff tickets as a treat, does going to the game cost you anything? At first blush, the answer is "No"—the tickets were free.¹⁶ But suppose that you could have sold the tickets for \$500 to a ticket agent, or "scalped them" in the active market that assembles before each game just outside Staples Center in Los Angeles. Then by going to the game you gave up the opportunity to have

¹⁶ You're thinking like an economist if you immediately thought that the tickets might have been costly because you could owe the friend a favor in the future (i.e., the tickets create a debt which has a future cost), or you got them because you were being repaid from a previous debt (so that accepting the tickets forecloses the opportunity to collect on the debt in another form).

\$500 instead. The tickets “cost” you \$500 in an economic sense (though not in accounting sense).

In behavioral economics, we often find that people underweigh opportunity costs (Thaler, 1985). The ticket example is a classic one, in fact— many people *would not* pay \$500 to go to the game, but *would* go and forego the opportunity to earn \$500. It seems that people don’t see giving up \$500 they could easily have as psychologically equivalent to pulling \$500 out of an ATM. Thinking about opportunity costs does not seem to be natural: It takes the ability to reason counterfactually, to figure out something that did not happen (but could easily have), and give the counterfactual the same cognitive status as what did happen.

An example that illustrates the link between opportunity cost and comparative advantage is the fictional small town in which there is just one lawyer, who is also the fastest typist among the 100 people in town who can type at all. What should the lawyer do with her scarce workday? If she types, the most typing will be produced. But if she only types, the opportunity cost is high because the fast-typing lawyer is the only lawyer; no lawyering will get done. The correct answer is that the lawyer should practice only law and have somebody else type. The best activity for each worker to perform is the one that has the lowest opportunity cost.

Partial and general equilibrium: Unintended consequences

A very important concept in economics is “partial” versus “general” equilibrium. A general equilibrium analysis takes account of *all* of the likely reactions and counterreactions of *all* agents in the economy. A partial equilibrium analysis assumes— usually incorrectly—that some agents do not react after a change, or are not behaving optimally.

Why would you do a partial equilibrium analysis rather than a general one? One answer is, “For simplicity”. Computing what happens in general equilibrium, when all agents are optimizing, is often enormously difficult. And ignoring the actions of some agents is often be a reasonable approximation to the result of a perfect analysis that includes actions by all agents (just as ignoring Pluto’s gravitational pull on the Earth is not a bad approximation).

For example, suppose you are interested in the effects of a serious earthquake in Los Angeles on the world economy. The most obvious economic effect is commercial and residential property damage in Los Angeles. But an earthquake might also cause people to move away from Los Angeles. A general equilibrium analysis will include this relocation effect and takes into account its impact on the cities Los Angelenos move to. Suppose Los Angelenos move to Las Vegas, which causes a spike in housing prices. Suppose higher property taxes that are collected by the city are spent on infrastructure, like expanding the airport (which is always busy). That lowers the cost of traveling to Vegas to gamble, which lowers demand for Native American casinos in northern California. So one effect of an earthquake in Los Angeles could be a recession in Native American casinos which are hundreds of miles away.

General equilibrium analysis is like the zany machines designed by Rube Goldberg [find picture]. At each step of the analysis you ask “And then what?”—given the temporary result of the analysis, do any agents have an incentive to change what they will do? The analysis ends when no agents have an incentive to change; the economy is said to be “in equilibrium”.

The distinction between partial and general equilibrium is important because a general equilibrium analysis will often reveal surprising effects or “**unintended consequences**” of a policy which a partial equilibrium analysis misses. (We will see some of these unintended consequences below in studying the effects of changes in incentive systems—sometimes they create cheating, or produce too much of one kind of worker effort at the expense of another.)

A common unintended consequence is that a change in the economy which is bad for many people actually benefits some. For example, the end of the Japanese stock market bubble in the late 1990’s led to a severe recession which led to a rise in burglaries in big cities Tokyo and Osaka. Where’s the silver lining in this gloomy cloud? Ask Makoto Iida, the founder of a leading security systems company Secom. Iida said “I feel very good now” because sales of his \$80/month intruder-detector systems rose 20% a year through the mid-90’s as the crime rate rose.

Incidence of taxes is another example of how the general equilibrium effects of an economic change might be quite different than they appear at first blush. When a tax is applied, who ends up actually paying the tax depends on supply, demand and competitiveness. For example, suppose you impose a profit tax on corporations. Suppose profits were determined by competitive forces before the profit tax, so that shareholders who supplied capital were earning a fair rate of return on investing in corporations. Then after the tax is imposed the net (post-tax) corporate profit will be lower; and by inference, it will be less than shareholders will accept because the new profit rate will be lower than the higher pre-tax rate. To make up the effect of the tax to shareholders, firms will raise price (or cut costs). Raising prices will usually cut demand, to an extent which depends on the price elasticity. If profit taxes are imposed across the board (on all firms), then prices will generally increase. So *consumers* may end up paying higher prices because of a tax on *corporate profit*.

A few years ago, a luxury tax was imposed on the sales of fancy cars and boats that only wealthy people buy. The tax was 10% of the price of cars and boats above some threshold (which implicitly defined the threshold of “luxury”). A small political debate ensued after people realized that the tax would reduce demand for boats, which were often built in small shipbuilding towns like New Bedford, Massachusetts which were struggling economically even before the proposed luxury tax. So what appears to be a tax on the rich had an unintended consequence of raising unemployment of some un-rich craftsmen.

Lorenz effect TBA

Supply and demand

Supply emergence.

Drug cartels (Al Qaeda).

Centralized planning vs decentralized markets TBA

Old (now resolved?) debate: Can centrally-planned economies operate more efficiently than laissez-faire decentralized ones? Most evidence suggests that central planning is worse, *as judged by consumers*. Possible exceptions: Targeted industries in “Seven Tigers” in Asia (Korea in technology, Japan).

Central planning: An agency orders state-owned factories to produce a certain amount of goods-- say **bread** (perhaps informed by surveys of public demand, perhaps by fiat or whim). A price is fixed by the agency. Goods are produced and sold at the fixed price.

Apparent advantages: *Price stability*. Governments can guarantee that prices do not fluctuate. But! If there are fluctuations in demand how are goods rationed (e.g. if demand goes up)? Typical answers: Nepotism (friends get bread first); queues (lines); bribes. *Full employment*. Government employs everybody. *Low prices*. Can be set low (i.e., factories run at a loss) through cross-subsidy, taxation, or inflation.

Planners may get poor signals about whether to make more or less bread, and perhaps poor signals about quality. Factory managers have no direct (reward) incentive to earn more profit. Workers have no direct incentive to work hard.

Decentralization. Nobody in the government cares about bread. Firms are allowed to spring up and bake bread and sell it however they want. If they make too much they lose money. If the bread is popular and lines form, they make more or raise the price. If bread machines become more efficient workers lose their jobs; they hopefully find more productive employment elsewhere (as personal trainers?)

Central planning vs decentralization (bread, claiming races, private vs publicly funded exploration, BSZ p 56)

Another example: **Claiming races for horses.**

The public likes competitive races (bet more when odds are close). How to arrange them?

Obvious, not-so-great idea: Somebody at the racetrack knows all the horses and asks around, and puts them in races together.

Good idea: Announce a race for \$20,000 purse and \$20,000 claiming price. If horses are entered they can be claimed (bought) by anybody for \$20k. Horses worth more than \$20k won't enter. Bad ones won't either (too expensive to race in a hopeless race).

Public vs privately organized adventure expeditions (North Pole, Northwest Passage 1818-1909):

	Public (n=35)	private (n=56)
Crew deaths	8.0%	6.2%
# of ships (# lost)	1.63 (.53)	1.15 (.24)
scurvy	46.7%	13.2%

Why? Public had worse leaders, adapted slowly to new info

(see BSZ p 56)

Financial markets

Orange juice

What is equity? /debt? Market efficiency. Sketch of evidence. Good news and bad news.

Important for incentive conflict

3. Models of human nature MORE TBA

A crucial ingredient for organizing people is a working hypothesis about human nature. Do people only care about themselves? Or do they care about doing right, or helping their friends, or something beyond narrow self-interest? Are there different *types* of people in these categories; and if so, how can we identify them and match them with the right kinds of jobs?

George Stigler quote from Fehr.

Experimental evidence suggests two important modifications to the view that self-interest is reliable and basic:

First, not everybody is self-interested all the time. In fact, in experiments only a minority of people will act purely self-interestedly when they know it will harm somebody else, even when they can get away with it. TBA TRUST EXPERIMENTS

Second, people are different (or at least, in any given situation people will not all act in the same way). Heterogeneity matters because it can interact with economic structures in interesting ways.

A beautiful illustration of this is STORY

Thaler story.

No. Experimental evidence.

Several models have emerged of “social utility”. These models try to add variables that represent emotions or other forces to the basic model of self-interested used in game theory and labor economics. The idea is to calibrate parameters which represent weights on the variables and use them to see whether the *same* basic model of social utility can explain patterns in *many* different games.

One model was created by Fehr and Schmidt (1999). Call the vector of payoffs created in a game $X \equiv (x_1, x_2, \dots, x_n)$, where x_i is the payoff to player i (and there are n players). They assume the utility of player i for the payoff vector X is

$$(1) \quad u_i(X) = x_i - \alpha / (n-1) \sum_{k=1}^n x [x_k - x_i]_0 - \beta / (n-1) \sum_{k=1}^n [x_i - x_k]_0$$

where the function $[x]_0$ denotes the maximum of x and 0 (i.e., if $x < 0$ then the function is set to zero). The first term (after the own-payoff x_i) represents “envy”, how a player’s utility is reduced by knowing that other players are getting *more* than she is. The second term represents “guilt”, how much a player’s utility is reduced by knowing that other players are getting *less* than she is. The parameters α and β represent the “strength” of envy and guilt relative to self-interest.

Experimental evidence from various types of games suggests that α and β are both positive and $\alpha > \beta$ (i.e., envy is stronger than guilt; in fact $\beta=0$ is often a good approximation and reduces the theory to one with one free parameter). In fact, some observations suggest that people like being ahead of others, so that $\beta < 0$ and people get *pleasure* from having higher income status than others.

One problem with the Fehr-Schmidt formulation in (1) is that it does not incorporate a conception of intention...MORE FROM BOOK

ADD ROBERTO

TO FINISH suggests there are three types of people: Conditional cooperators; purely self-interested agents; and “saints” who always cooperate, regardless of whether others do.

But rationalization...point is that there are differences and possibly structural influences.

“Love thy neighbor but keep your guard up/You can’t trust nobody when they hard up”—hip hop lyric.

“Trust, but verify”- Sam Walton.

“Traits or states”: Nice people, or nice times?

Attribution theory conclusion

Situational influence (cannibalism example)

How alike are people? (probably more alike than we think)

The behavioral economics view:

Natural limits on computation, perception, willpower and greed. Complex contracts are not fully digested.

Overconfidence and self-serving bias.

Money as symbolic (exchange of \$ with Mrs. Wong, rejecting ultimatum offers, “fine is a price” etc).

Some case studies:

Pizza delivery guy <http://tipthepizzaguy.com/stories/story17.htm> wild stories!! He would take pizzas that couldn’t be delivered (which as a matter of policy the drivers were

allowed to return, and eat or give away) and keep them in his car. Occasionally somebody would flag him down and ask if he had a spare pizza; he'd sell them the old undelivered pizza and pocket the cash.

The same driver would offer to pick up anything else somebody wanted ("beer, cigarettes[sic], McDonalds fries, tampons, whatever...") for a 100% delivery surcharge.

Moral hazard: Employee bartered cheap pizza (due to employee discount) for services.

- I used to pay for my oil changes with pizza. I think we paid 25% of retail for pizzas, so an XL one topper carryout was like \$3. Equivalent to the oil change guy, \$20. Price of an oil change, \$20. I got it for \$3.
- I attempted to pay for car repair with pizzas, but found myself being this guy's "pizza \$@&#!" for two months. Car repair is too expensive to pay for with pizzas. Pony up the \$. ;-)

4. Game theory and incentives

Basics of game theory:

A "game" is a frivolous term social scientists use for a mathematical description of a strategic interaction. The elements of a game are: Players, strategies, information, order of moves, outcomes, and utilities. Outcomes are determined by the strategies players choose, and the information they have. While we will often talk about financial outcomes as a canonical case (especially in thinking about companies), it is crucial to remember that an outcome is usually a physical event in the world—getting promoted, bankruptcy, earning profits, an apology and money in a legal settlement, and so forth. We assume that people have numerical *utilities* over outcomes. Usually we only need to assume that people can rank which outcomes they like best (i.e., utility is "ordinal"—1st, 2nd, 3rd, and so on).

Game theory is a promising mathematical language to partly unify natural and social sciences because games can be played at many levels of analysis—from genes to nations. Here are some examples: Tennis players deciding whether to serve to the left or right side of the court; the only bakery in town offering a discounted price on pastries just before it closes; employees deciding how hard to work when the boss is away; an Arab rug seller deciding how quickly to lower his price when haggling with a tourist; rival drug firms investing in a race to reach patent; an e-commerce auction company learning which features to add to its website by trial-and-error; real estate developers guessing when a downtrodden urban neighborhood will spring back to life; San Francisco commuters deciding which way to work will be quickest when the Bay Bridge is closed; Lamelara men in Indonesia deciding whether to join the day's whale hunt, and how to

divide the whale if they catch one; airline workers hustling to get a plane away from the gate on time; MBA's deciding what their degree will signal to prospective employers (and whether quitting after the first year of their two-year program to join a dot-com startup signals guts or stupidity); a man framing a memento from when he first met his wife, as a gift on their first official date a year later (they're happily married now!); and people bidding for art or oil leases, or for knickknacks on eBay.

There are many books on game theory. The goal in the next few pages is to give the briefest sketch, and some notation, to equip you to grasp the essentials of what follows in the rest of the book. If you do not have some other background in game theory, and are serious about understanding the experimental results which are described later, you should read other books.¹⁷

First, some notation. Usually we talk about a finite number (n) of players, indexed by i . Player i 's strategy is denoted s_i . A vector of strategies, one for each player is denoted $S = \{s_1, s_2, \dots, s_n\}$. The part of this vector which removes player i 's strategy (i.e., contains every other player's strategy) is denoted s_{-i} . The utility of player i 's payoff from playing s_i when others are playing s_{-i} is denoted $u_i(s_i, s_{-i})$.

We start with the principle(s) of **dominance**. A strategy d_i is a **strictly dominant** strategy if it is a strict best response to any feasible strategy that the others might play. That is, $u_i(d_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_i \neq d_i$ and for all s_{-i} . (Sometimes a strategy d_i only dominates one or more other strategies s_i ; then the phrase "for all $s_i \neq d_i$ " in the last sentence is replaced by "for *some* $s_i \neq d_i$ ". Sometimes one strategy is better than another for *some* choices by other players, but is not always strictly better and is never worse. Such a strategy d_i is called **weakly dominant**

STOPPED HERE for all s_{-i} . d and is never worse, $u_i(d_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_i \neq d_i$ and for all s_{-i} . $u_i(d_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_i \neq d_i$ and for all s_{-i} .

Strict dominance is a useful idea because there is no reason why a player should *not* choose a strictly dominant strategy; choosing a different strategy just guarantees that there will be some regret, regardless of what other players do.

```
\begin{equation*}
u_{i}(s_{i}^{\ast},s_{-i})>u_{i}(s_{i}^{\prime},s_{-i})\text{\ \ }\forall\text{\ \ }
s_{-i},s_{i}^{\prime}\neq s_{i}^{\ast}
\end{equation*}
\end{definition}
```

The strategy s_{i}^{\prime} is dominated if there exists s_{i}^{\prime} $\in S_{i}$ such that

```
\begin{equation*}
```

¹⁷ A good introductory book (low on math) is Dixit and Skeath (1999). More mathematical books include Rasmusen (1994) and Osborne and Rubinstein (1995). Gintis (1999) includes fresh material on evolutionary theory and experimental data, and tons of problems. The heavy tomes which are used in graduate classes at places like Caltech include Fudenberg and Tirole (1991).

$$u_i(s_i^{\prime}, s_{-i}) > u_i(s_i, s_{-i}) \text{ for all } s_{-i}$$

The strategy s_i^{\prime} is weakly dominated if there exists s_i such that

$$u_i(s_i^{\prime}, s_{-i}) \geq u_i(s_i, s_{-i}) \text{ for all } s_{-i}$$

$$u_i(s_i^{\prime}, s_{-i}) > u_i(s_i, s_{-i}) \text{ for at least one } s_{-i}$$

`\begin{example}`

Consider the simple normal-form game below. In a normal-form (a/k/a strategic-form, or matrix) game players are presumed to move simultaneously so there is no need to express the order of their moves in a graphical tree (or extensive-form). Each cell shows a pair of payoffs. The left payoff is for the row player (1) and the right is for the column player (2). The payoffs are *utilities* for consequences. That is, in the original game the consequences may be money, pride, reproduction by genes, territory in wars, company profits, pleasure or pain. A key assumption is that players can express their satisfaction with these outcomes on a numerical utility scale. The scale must at least be ordinal-- i.e., they would rather have an outcome with utility 2 than with utility 1-- and when expected utility calculations are made the scale must be cardinal (i.e., getting 2 is as good as a coin flip between 3 and 1).

	Player 2		
	L	M	R
Player 1	U	D	U
	1,0	0,1	1,2
	0,3	2,0	0,1

For player 2, strategy R is strictly dominated by M (because M gives a higher payoff if player 1 chooses U, 2 instead of 1, and a higher payoff if player 2 chooses D, 1 instead of 0). Deleting strategy R (i.e., assuming a rational player 2 will never play it) makes D strictly dominated by U. But if player 1 plays U, then player 2 should play M. Therefore, the iterated-dominance equilibrium is (U, M).

`\end{example}`

Dominance is important because if utility payoffs are correctly specified (one need only get their *order* right) and players care only about their own utility, there is no good reason to violate strict dominance. One

step of iterated dominance is a judgment by one player that the other player will not make a dumb mistake. This often tells a player what she herself should do. In the example, player 1 might consider choosing D because of the chance of earning the 2 payoff in the lower right (D,R) cell. But will she ever earn that payoff? Only if player 2 does something which is dominated. If player 1 assumes player 2 won't do that, she can rule out R and her hope of earning 2 disappears. Then she should obviously choose U.

`\begin{example}`

(Battle of the Sexes) The game

`\begin{tabular}{cccc}`

`& & \multicolumn{2}{c}{Player 2} \\`

`& & L & M \\ \cline{3-4}`

Player 1 & U & \multicolumn{1}{|c|}{2,1} & \multicolumn{1}{|c|}{0,0} \\

`\cline{3-4}`

& D & \multicolumn{1}{|c|}{0,0} & \multicolumn{1}{|c|}{1,2} \\ \cline{3-4}

`\end{tabular}`

is not dominance solvable. Neither strategy is dominant (or dominated) for either player because there is no one strategy that is always best. Put differently, each strategy **might** be best depending on what you think the other person will do.

`\end{example}`

`\subsection{Nash Equilibrium}`

`\begin{definition}`

The strategy profile $s^* = (s_1^*, s_{-1}^*)$ is a Nash equilibrium if each player's strategy is a best response to the other players' strategies. That is, no player has incentive to deviate, if no other player will deviate. (If players find themselves in equilibrium, there is no reason to move away.)

`\begin{equation*}`

$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i', s_{-i}^*)$

`\text{ } \forall s_i' \in S_i`

`\end{equation*}`

`\end{definition}`

Note that if a strategy profile is an iterated-(strict) dominance equilibrium, then it is a Nash equilibrium. This is not true of equilibria which are created by iterated application of **weak** dominance.

`\begin{example}`

(Battle of the Sexes) Solving for pure strategy Nash equilibrium:

`\begin{tabular}{cccc}`

```

& & \multicolumn{2}{c}{Player 2} \\
& & L & R \\ \cline{3-4}
Player 1 & U & \multicolumn{1}{c}{2,1} & \multicolumn{1}{c}{0,0} \\ \cline{3-4}
& D & \multicolumn{1}{c}{0,0} & \multicolumn{1}{c}{1,2} \\ \cline{3-4}
\end{tabular}

```

If Player 1 plays U, 2's best response is L. If Player 1 plays D, 2's best response is R. If Player 2 plays L, 1's best response is U and if 2 plays R, 1's best response is D. Therefore, U is a best response to L and L is a best response to U. Likewise, D is a best response to R and R is a best response to D: Pure strategy NE are (U,L) and (D,R).

\end{example}

\subsection{Mixed Strategies}

A mixed strategy for player i is a probability distribution over all the strategies in S_i

\begin{example}

(Battle of the Sexes) Solving for mixed strategy Nash equilibrium:

```

\begin{tabular}{cccc}
& & \multicolumn{2}{c}{Player 2} \\
& & L & R \\ \cline{3-4}
Player 1 & U & \multicolumn{1}{c}{2,1} & \multicolumn{1}{c}{0,0} \\ \cline{3-4}
& D & \multicolumn{1}{c}{0,0} & \multicolumn{1}{c}{1,2} \\ \cline{3-4}
\end{tabular}

```

Suppose Player 1 plays U with probability p and D with probability $1-p$ and player 2 plays L with probability q and R with probability $1-q$.

Then the expected value to 2 from playing L is

\begin{equation*}

$$1p+0(1-p)$$

\end{equation*}

and the expected value to 2 from playing R is

\begin{equation*}

$$0p+2(1-p)$$

\end{equation*}

Player 2 is indifferent iff

\begin{equation*}

$$1p+0(1-p)=0p+2(1-p)$$

\end{equation*}

or

```

\begin{equation*}
p=\frac{2}{3}
\end{equation*}

```

The expected value to 1 from playing U is

```

\begin{equation*}
2q+0(1-q)
\end{equation*}

```

and the expected value from playing D is

```

\begin{equation*}
0q+1(1-q)
\end{equation*}

```

Player 1 is indifferent iff

```

\begin{equation*}
2q+0(1-q)=0q+1(1-q)
\end{equation*}

```

or

```

\begin{equation*}
q=\frac{1}{3}
\end{equation*}

```

As a result, a pair of (weak) best responses constitutes a mixed strategy equilibrium:

```

$\left( (\frac{2}{3}U, \frac{1}{3}D), (\frac{1}{3}L, \frac{2}{3}R) \right)$
\end{example}

```

Mixed-strategy equilibrium is a curious concept. Introducing mixed strategies makes the space of payoffs convex (i.e., for any two points in

the space, all points in between are in the space too), which is necessary

to guarantee existence of a Nash equilibrium (in finite games).

Guaranteed

existence is a beautiful thing and is part of what makes game theory productive: For *any* (finite) game you write down, you can be sure to

find an equilibrium. This means that a policy analyst or scientist trying to

predict what will happen will *always* have something concrete to say.

However, the behavioral interpretation of mixing strategies is dubious. By

definition, a player only desires to mix when she is indifferent among pure

strategies, which means she does not (strictly) desire to mix with particular probabilities; she just doesn't care what she does.

Furthermore,

one player's equilibrium mixture probabilities depend *only* on the

other player's payoffs which is odd. A modern interpretation of mixed-strategy equilibrium (called "purification") is that one player might

appear to be mixing but is actually choosing a pure strategy conditional on some hunch variable they privately observe. Mathematically, this works the same way-- as long as each player's *\emph{belief}* about the other players' choice matches the predicted probabilities, the mixed equilibrium is a mutual best-response point. Chapter 2 gives more detail.

Denote player I's strategy choice by s_I
 Strategy s_d is **dominated** for player I if
 I.e., $u_i(s_1, s_2, \dots, s_d, \dots, s_n) \leq u_i(s_1, s_2, \dots, s_{i^*}, \dots, s_n)$
 for all (s_1, s_2, \dots, s_n) and for (at least) one s_{i^*} .

Nash equilibrium $u_i(s_1^*, s_2^*, \dots, s_i^*, \dots, s_n^*) \geq u_i(s_1^*, s_2^*, \dots, s_i, \dots, s_n^*)$
 For all s_i
 (a mutually consistent point where A's beliefs match B's optimal choice and vice versa)

Mixed strategy equilibrium: Probabilistic mixtures of strategies are the only "mixed" strategies that form a (weak) equilibrium.

Coordination

Many interesting games have multiple equilibria which have different payoff consequences. These games capture the problem of "coordination".

An important game in this category is "battle of the sexes" (BOS), shown in Table?. Two people, Pat and Chris can choose a big fish or a little fish. (This version is modeled on anthropological situation in which two people in a small-scale society are trying to decide how to divide two differently-valued goods.) Both prefers to divide the two fish between themselves, so that one gets the big fish and one gets the little fish, but each prefers to get the big fish. If the players mismatch, and both choose the same-sized fish, they get zero. (If the payoffs from both sides grabbing for the little fish are between 1 and 3, then this a game of "chicken".¹⁸).

The BOS is a perfect game to illustrate "mixed motives": Both players want to agree on one of the two divisions on the diagonal of the table (since either division if preferred to getting zero), but each prefers a different division.

Table: Battle of the sexes (BOS) game

		Pat's choice	
		Little fish	Big fish
Chris's choice	Big fish	3,1	0,0
	Little fish	0,0	1,3

¹⁸ MORE HERE

Another interesting coordination game is called “stag hunt”, or the “assurance game” (see Table ?). In this game players can play it safe and earn 1 for sure, by hunting for rabbit. Or they can join a stag hunt, and earn a larger amount x (with $x > 1$) if the other person also hunts stag. However, hunting stag is risky because if the other players hunts rabbit the stag-hunter gets 0. So hunting stag is a bet that the other player is likely to hunt stag too. It is called an “assurance game” because hunting stag requires an assurance that another player will help out.

[TBA picture from *Time* magazine]

Table: Stag hunt game ($x > 1$)

		Pat's choice	
		Stag	rabbit
Chris's choice	stag	x, x	0, 1
	rabbit	1, 1	1, 1

Communication in games

Prisoners' dilemma		C	D
	C	2, 2	0, 3
	D	3, 0	1, 1

Hotelling location games:

N firms locate on a line $[0, 1]$ irreversibly, in order.

Customers evenly distributed.

Customers go to the nearest store.

Where do customer-maximizing firms locate?

$N=2$ case. $N=3$ case.

Line versus circle

Betting games:

1 and 2 are each privately informed about a "set of states"

1 will know "It's A or B" (A B) or "It's C or D".

2 will know "It's A" or "It's B or C" or "It's D".

	A	B	C	D
1's payoffs	<u>+32</u>	<u>-28</u>	<u>+20</u>	<u>-16</u>
2's payoffs	<u>-32</u>	<u>+28</u>	<u>-20</u>	<u>+16</u>

e.g. A="earnings will be great"

B= "earnings will be good"

C= "earnings will be bad"

D="earnings will be terrible"

1 knows whether earnings will be (good or great) or (bad or terrible)
2 knows whether earnings will be surprising-- either great, (good or bad), or terrible.

When should you bet?

Game theory and organizational contracting

Before shaking hands: *Hidden information*

Until the 1970's or so, game theory was only capable of addressing strategic thinking when both players have the same information, which is quite unrealistic in most organizational contexts. Workers generally know how hard they have been working and usually have an incentive to hide what they know from employers. Other times the opposite is true—employers may know something the employees don't because they have more experience, like how dangerous high-rise construction work is or how frequently biotech projects lead to valuable products.

As is usually true in scientific modeling, game theorists were well aware that people often have different information—and people with less information may know that there is something they don't know, but don't know what that something is-- but they didn't know how to model such situations. Then in 1967 John Harsanyi showed how, for which he shared the Noble Prize in 1994.


Harsanyi suggested studying these situations by assuming that each player has some private information, summarized by their “type”. The key step which makes the modeling workable is that the chance that each player is of a particular type is “commonly known”—which means both players know the chances, both know that others know the chances, and so on.

This, too, is unrealistic but Harsanyi's framework provided a useful start.

The ability to mathematically model situations with differences in information rapidly drew attention to the influence of two different kinds of informational asymmetries in organizational life. One is called “adverse selection” or “hidden information”. The term “adverse selection” comes from the insurance business. Consider car insurance. Insurance companies can ask a few basic questions before deciding which drivers to insure—such as age, gender, accident history, zip code, how many miles the driver usually drives, and the type of car the driver owns. But these variables do not perfectly predict whether a person is a safe driver who is likely to get in accidents and file insurance claims. Remember that in a competitive market, insurance companies will offer insurance at a price which equals the expected cost of accidents (plus a premium to pay for administrative overhead, and profit margin). Drivers will choose the policies with the lowest prices.

Now suppose the drivers themselves know *something* about how safe a driver they are, which the insurance company doesn't know. Then, in theory, the drivers who realize they are unsafe drivers are more likely to buy insurance, if the price is less than what they know the true expected cost to be. Thus, the self-selection of buyers who buy insurance goes against the company's interest—we say they are being “adversely selected against”.

The possibility that the kinds of people who select to undertake an action depends on their hidden information happens throughout life and is central to modern economic thinking. (Figure ? below shows how it might apply to savvy fish.) In labor markets, selection is sometimes called “sorting” or “assortative matching”. For example, when companies switch to a system where pay is more closely linked to measured output or performance, performance often goes up. Part of the effect is probably due to the fact that people work harder because they now earn more because of their extra effort. But part of the effect could also be due to the fact that the higher bonus draws very hard-working employees into the firm. (At the national level, this is familiar in the stereotype of hard-working immigrants who often literally risk their lives to migrate to a country where hard work can earn them more money.)



Of course, companies are often well aware of the adverse selection problem. There are basically two techniques which help companies minimize adverse selection by helping companies: Measure, or screen.

Measurement means trying to uncover the hidden information which potential buyers have about themselves. Insurance companies use “experience rating”, which means charging more for people who have had more accidents, or refusing to sell insurance at any price for very high-risk groups. A related technique is to offer lower rates to people who are members of a group, since the insurance company can gather reliable statistics if the group is large, and simply being a member of a group may be correlated with driving safety. For example, young drivers are notoriously accident-prone and insurance companies dislike insuring them. Student with good grades in school can often get discounts on their insurance rates, because insurance companies think (probably based on large samples of accident rates) that students who are conscientious and sharp enough to get good grades are also likely to be safer drivers than average.

Screening refers to creating a set of different insurance policies so that drivers with different self-perceived ability choose different contracts. That way, the contracts “screen out” the bad drivers from low-cost policies (which are priced to make profits only if expected accident costs are low). A good example is *deductibles*. A deductible is the amount that is deducted from the accident cost before the insurance company pays the claim; or put differently, it is the amount the driver “copays” if there is an accident.

Drivers who believe they are really safe drivers will not mind a policy with a large deductible (\$1000 is typical for a large deductible) since they will not expect to pay that amount very often. Bad drivers will prefer a lower deductible-- or no deductible at all—because they think there is a high chance they will have an accident and have to pay the deductible.

The difference in insurance premiums for different deductible policies is enormous. No-deductible policies are wildly expensive. [put some numbers here]. This is consistent with the theory that deductibles are helping the insurance company screen drivers by ability: The safe ones accept high deductibles and pay much lower premiums; the riskier ones are happy with lower deductibles even though their premiums are much higher.

"Lemons" problem.

A firm is worth v to seller, uniform $[0,100]$

It's worth $1.5v$ to buyer. How much should buyer pay to maximize expected gain?

Winner's curse

After shaking hands: *Hidden actions*

Another problem of limited information is called “hidden action” or “moral hazard”. Hidden action means that one party has an incentive to take an action to benefit themselves at the expense of the other. An extreme form is flat-out fraud and cheating. Oliver Williamson uses the term “opportunism”, which he defines as “self-interest seeking with guile”.

Hidden action often shows up in ordinary transactions. In many African countries, bags of rice are sold by weight. Stones are sometimes added to the bottom to increase the weight of the bag, and covered up with rice so the stones are not visible. Savvy buyers know to dig their hands into the rice bag looking for stones before they buy.

Hidden action is common in one-time transactions, especially if one side has paid in advance for a product or service. Once I hired a contractor to paint my house. We signed a contract and I paid a deposit. The verbal agreement was for 2-4 painters to come the next day and finish the job over a couple of days. The next morning only one painter came; the job took a total of 5 days. In situations like this, the seller of the service has the buyer over the barrel—once the house is partly painted, the buyer typically does not want to cancel the contract and fire the sellers.

In professional service firms (such as law) where clients pay for billable hours of service performed by professionals, overbilling can be common. In fact, it is possible that an equilibrium emerges in which clients *expect* to be overbilled, but also expect to haggle

for a reduction in their total legal bill when a case or project is finished. If so, the firm may tolerate overbilling because if they billed honestly the client might still expect a reduction.

Cardboard

Sometimes markets flourish in one place, rather than another, because geographical differences or other forces influence the extent of hidden action. For example, a lot of used cardboard boxes are collected and sold in Los Angeles, to be used again or recycled. Why Los Angeles? Cardboard is not more plentiful in Los Angeles than in another large city, like Houston or New York.

The reason why the Los Angeles market flourishes has to do with the nuances of how the market operates. It takes too much time to count the sheets of cardboard in a large stack; it is easier to sell cardboard by weight. But selling cardboard by weight creates the potential for hidden action. Cardboard is porous and becomes much heavier when it absorbs water. This property of cardboard makes it hard for buyers to tell whether they are overpaying because the cardboard they are buying is wet, or because sellers watered the cardboard to increase its weight and sell it for more money. But Los Angeles is basically a desert city; it only rains about 10 days a year. The used cardboard market flourishes in Los Angeles because the air is dry so that variations in cardboard weight from water are low. This fact reassures buyers who are paying by the pound and makes it relatively easy to see if cardboard has been artificially watered to increase its weight.

Pizza barter

Pizza stores that have delivery to your home or office have Pizza delivery
Pizza delivery guy <http://tipthepizzaguy.com/stories/story17.htm> wild stories!! He would take pizzas that couldn't be delivered (which as a matter of policy the drivers were allowed to return, and eat or give away) and keep them in his car. Occasionally somebody would flag him down and ask if he had a spare pizza; he'd sell them the old undelivered pizza and pocket the cash.

The same driver would offer to pick up anything else somebody wanted ("beer, cigarettes[sic], McDonalds fries, tampons, whatever...") for a 100% delivery surcharge.

Moral hazard: Employee bartered cheap pizza (due to employee discount) for services.

- I used to pay for my oil changes with pizza. I think we paid 25% of retail for pizzas, so an XL one topper carryout was like \$3. Equivalent to the oil change guy, \$20. Price of an oil change, \$20. I got it for \$3.

- I attempted to pay for car repair with pizzas, but found myself being this guy's "pizza \$@&#!" for two months. Car repair is too expensive to pay for with pizzas. Pony up the \$. ;-)

(Near-)death from above: Air traffic control

One of the most remarkable studies of moral hazard is a study of air traffic controllers, by Staten and Umbeck (1982).

Air traffic controllers work under great stress, to guide planes safely into major airports. Because of job burn-out, the Federal Aviation Administration (FAA) has an elaborate set of rules for detecting worker disability and remedying it. Like all government employees, controllers are covered by the Department of Labor's Office of Workers' Compensation Programs (OWCP), starting in 1974. The Federal Employees' Compensation Act (FECA) provides disabled workers with either a fixed award or a percentage of salary paid for the duration of an injury. If the OWCP verifies that an injury is disabling, and job related, controllers with dependents can earn 75% of their base salary tax-free (which can higher than their after-tax pay for highly-paid controllers with many years on the job).

Starting in September 1974, several amendments to the FECA act increased the incentive for burned-out controllers to collect disability. One amendment was to allow testimony of clinical psychologists, rather than M.D.'s, as primary medical evidence of a disability. Another amendment allowed controllers to pick the private doctor of their choice to supply medical testimony, rather than being examined by their airport's flight surgeon. These two amendments allowed controllers to "shop" for a doctor willing to attest to their disability. These amendments came at a time when the OWCP, coincidentally, abolished its investigative staff (who previously scrutinized dubious claims) and assigned investigation to busy claims examiners with heavy caseloads, without increasing the number of examiners proportionally to the increased caseload. Furthermore, in May 1972 Congress authorized funds for a "Second Career Training Program", to provide up to two years of pay for job training, as well as paying 100% of the base rate of pay, for controllers who had worked for more than five years.

A controller who files for disability with the intention of being deemed medically unfit to continue is said to "punch out" of the job. The 1974 changes in the OWCP, the reduction in claims investigation, and the Second Career incentive created a "perfect storm" situation that is ripe for moral hazard by controllers who want to punch out and get fully paid for two years while training for a new job.

Table ?

Disorder	Pre-Second Career Disorder incidence	Post-Second Career Disorder incidence	Percentage change
Respiratory	1.9	1.5	-21
Muscles	.5	.4	-20
Ear,nose,throat	6.7	8.0	+19
Abdominal	16.7	20.4	+22
Eye	5.6	7.4	+32

Bones and joints	2.3	4.6	+39
Cardiovascular	22.1	32.5	+47
Neuropsychiatric	10.9	27.2	+150

Table ? shows the increase in medical incidence (per thousand person years of service) between the periods before and after the Second Career provision. Notice that reports in all categories rose (except for rare muscle and respiratory cases). The increase was by far the largest in the “neuropsychiatric” category, which almost tripled.

The number of controllers who were disqualified from duty for medical reasons, who were not otherwise eligible for early retirement (older than 50, or with 25 or more years on the job) also rose sharply during the mid-1970’s, from 143 in 1973 (before the OWCP changes) to about 400 a year after 1975. This increase, and in the medical incidence increases shown in Table ?, are rough clues that at least some controllers were appealing for medical disability in higher numbers when the incentive to do so rose, because it was easier to shop for a doctor to testify to disability, and because of the generous Second Career provision (two years of full pay during retraining).

The most striking evidence, however, comes from specific incidents. In evaluating claims of psychological burnout (as opposed to medical claims, which are easier to verify), OWCP claims examiners “have been instructed to look for specific stressful incidents on the job that were symptoms of or contributed to a disability”. In fact, there are two prominent categories of “specific stressful incidents” the FAA carefully tracks—“system errors”, and the ominously-named “near mid-air collisions” (NAMCs). At air terminals, the minimum physical separation between incoming planes is three horizontal miles, or 1,000 vertical feet. (Keep in mind that jets going 600 miles an hour cover three miles of separation in 18 seconds, and descend before landing at about 1,000 feet in about a minute.) A NMAC occurs when aircraft are forced to take evasive action, or come less than 500 feet apart (a distance they can cover in half a second).

Staten and Umback did a statistical regression of the number of system errors in each quarter of the year, for each of 10 FAA regions, from 1973 to 1976, a period that covers the time before and after the OWCP amendments. The biggest determinant of the number of system errors is the amount of air traffic—there were .012 errors per aircraft, or about one for every 80 planes.

Amazingly, the number of quarterly errors per region went up by 3.12 after the OWCP amendments (a number which is quite significantly different than you would expect by chance). It appears that some controllers were *deliberately* letting planes get close together, creating system errors that could be verified by OWCP claims examiners, in order to get out of their jobs and get two years of full pay and job retraining!

The controllers were not *too* reckless, however. A regression of the number of near-miss NAMCs did not show an increase after the OWCP amendments. Furthermore, *before* the OWCP amendments, about 40% of the system errors were reported by the controllers themselves, and a third of those were also reported as NAMC’s. *After* the amendments, the percentage of system errors the controllers reported themselves rose to 58%, but only a sixth were also reported as NAMCs. The data suggest the controllers were deliberately letting planes get close together, but not close enough to create

dangerous near-misses. (If anything, they were more careful to avoid NAMC's while creating more system errors.)

As with any statistical analysis, it is important to disentangle causality from correlation. It is possible that the OWCP amendments were passed because air traffic control was getting more difficult, and more system errors would have happened anyway. (The analysis tries to control for this by including the amount of traffic, however.) A final piece of evidence suggests, again, that the increase in system errors was a deliberate response to the increased incentive to get medical disability. Remember that controllers needed five years on the job to become eligible for the generous Second Career full-pay retraining benefits. In 1974, before the changes, 20% of the controllers had less than five years experience, but accounted for 40% of the controllers with reported system errors (a ratio of about 2); the corresponding ratio of error-prone controllers with 5-10 years of experience, compared to the number of controllers with that much experience, was about .8. These data suggest, not surprisingly, that inexperienced controllers were creating more errors. But by 1976, the ratio of error-prone controllers among inexperienced controllers—those ineligible for the Second Career benefit—was about 1.2, and the ratio for eligible controllers (those with 5-10 years experience) rose to 1.8: That is, the more experienced controllers were making more than their share of errors.

This amazing story shows how hard it might be to detect a terrible response to a change in incentives. System errors are relatively rare. The pattern detailed above only came to light when two clever economists, armed with a theory that led them to be suspicious that maybe increased incentives would change behavior, put together a lot of data and saw the pattern. Furthermore, this true-life tale shows how carefully incentive systems must be balanced. Changing the OWCP rules may have been a good policy to combat air traffic controller burnout, but if the OWCP had anticipated an increase in dubious claims, and had increased the number of investigators when the doctor-shopping amendments and Second Career benefits were passed, the number of system errors might not have risen so sharply, or at all.

What does behavioral economics have to say about the frequency of people exploiting their hidden action? The evidence described in section ? on human nature suggests that while self-interest is common, reciprocity springing from moral obligation is common too. The effects observed in experiments make it hard to believe that self-interest seeking “with guile” is so prevalent in the business world.

Preventing damage from hidden information and action

The most important force that reins in cheating is reputation from repeated interactions. If one person cheats another, the cheated party usually finds out and quits trading with the cheater. Sometimes the cheated party can sue for fraud or publicize what happened. The online auction site EBay developed a novel system of “testimonials”: Buyers who bought goods from a seller can post online comments about the quality of service from that seller. New buyers can see the comments; if they're bad that hurts the seller's business a lot. [give specific example here]

If you think people are usually guileless in trying to cheat others, you'd be surprised how well the EBay system works. People routinely send checks in the mail and

receive goods from people far away who they've never met, and may never buy from again. At the same time, the system does not work perfectly, and the amount of fraud from online internet transactions has increased dramatically in recent years. EBay eventually introduced an escrow system in which the buyer can deposit money with EBay, which is transferred to the seller only after the buyer receives their goods and reports their satisfaction.

Holdup

A special kind of hidden action is called the "holdup problem". The potential for holdup occurs when one party makes a large investment which is only valuable in a trading relationship with another party. After the investment is made, the investing party is at risk of being "held up" by the other party renegotiating the deal or, more subtly, being stingy in negotiations as quality and price are adjusted over time. To the extent that the investment is irreversible (or "sunk"), and specific to the relationship, a profit-maximizing investing firm has no choice but to accept the unpleasant result of new negotiations.

For example, it is common after getting married for men to gain a little weight, and for women to cut their hair shorter or spend less time on their personal appearance. (These kinds of changes don't appear to be permanent changes in tastes due to marriage, since after divorcing, it is common for men and women to go back to their single ways.) The most charitable way to interpret these common changes is that people work harder at their appearance before they get married, and slack off a little afterwards. Marriage raises the marginal costs of exiting the relationship, and newlyweds take advantage of this fact by investing less than they did earlier.

A common economic example is suppliers building a plant near a large customer, to economize on shipping costs. A famous example in economic history is Fisher Body, who were asked to build a plant near General Motors (GM) factories, to build automobile bodies for their large customer GM that would be cheaper to ship to their factories. As the story goes, Fisher Body balked at making this investment, which led GM to vertically integrate by simply buying Fisher Body, to eliminate the post-investment haggling.

Another example comes from journalism. Magazines which are produced every week or month usually do not own the printing presses that physically produce their magazine. [check this] . But daily newspapers usually *do* own their own printing presses. Why are these ownership structures different?

The difference is probably not just due to the fact that the daily papers use the printing press more often. Magazine publishers could own their own presses as well, and simply rent them out when they are not being used to print the magazine.

The holdup problem suggests a possible explanation. Once a daily newspaper's content is delivered to the printing press (or "put to bed"), the newspaper is extremely vulnerable. If the paper's publication is delayed, or a printing mistake is made, it is a small disaster since the paper is expected to come out daily, and on time. If a magazine is published late, or a printing error is made, the damage is lower for a magazine because the news they report is less timely. So if a newspaper owns the printing press facility, they can get more control (e.g. higher priority, and the ability to fix a mistake rapidly). A magazine, in contrast, doesn't benefit as much from added control over last-minute slips. Assuming there is some cost to integrating—typically, that the employees are less

motivated because they are paid a fixed wage, rather than having their earnings depend on market competition— the benefits to a newspaper to owning their presses can be higher than the costs, but the benefits to a magazine are not worth the added costs.

It is an open question how common holdup is in business relationships, and why it occurs. There are not many empirical studies. TBA here.

Limiting moral hazard TBA

Monitoring (e.g. "spying" on retail employees,

Baker et al paper on tracking truckers with GPS). {TBA}

Reputational incentives

Word of mouth

EBay "testimonials" from satisfied buyers

Lawsuits (e.g., class actions)

Escrow

"Bonding" or collateral (e.g., maid services are bonded)

long-lived institutions (e.g., family names, Crazy Karl's Karpet Barn, political parties)

Stock market value after airplane crashes, product recalls etc.

But...markets may under/overreact (e.g., all Arthur Anderson employees get tainted)


Bad workers can get good recommendations to "outplace" (academic hiring-- phone vs letter)

"Capital T" endgame/lame duck problem (e.g., President Clinton final-hour pardons)

SIDEBAR: Honor versus Monitoring¹⁹



Versus



¹⁹ This material was drafted by Galen Loram.

Virtually all corporations older than a month or two have some sort of corporate culture. What about academic institutions? Most do – they range from the emphasis on sports at the University of North Carolina, to the pledge not to drink, have sex or even dance that students at Wheaton College in Illinois make, to the cherished liberalism and activism at the University of California in Berkeley.

The California Institute of Technology, a small university in Southern California with a focus on math and science, has at the core of its culture an **honor code**. This code simply states: “No member of the Caltech community shall take unfair advantage of any other member of the Caltech community.”

According to Rutgers University Professor Donald McCabe with the Center for Academic Integrity, the recognized expert on the topic, 75% of college students admit to cheating at some point during their college years. But at Caltech students, faculty and staff estimate that only 2-4% cheat. What is responsible for the 20-fold discrepancy between Caltech and other universities?

Many schools have some sort of honor code but most are taken lightly. Students often do not know about the existence of a code, or its details or “teeth” (penalties for violating it). The Caltech honor code seems to work extraordinarily well. So what differentiates an effective honor code from one that hardly makes it into the guide for incoming freshman?

One reason is that the honor code is visible in practice at many places, and makes life easier. Students and professors leave their doors open, and often unlocked. Virtually all exams are take-home exams (sometimes “closed book”), typically with a maximum time limit that students can spend taking the exam stated. It is also assumed that people will read, understand and follow detailed policies on exam rules and homework collaboration– or clarify misunderstandings.

A recent study by the Center for Academic Integrity found that 82% of its members – colleges nationwide – felt that if a school did not already have an effective honor code, it was either ‘very difficult’ or ‘not possible’ to create one. This suggests a strong founder effect: when the university was founded tens or hundreds of years ago, what the founders decided to implement in the way of an honor code has a marked impact on the community years later. Cultures do not arise overnight, and one which requires a vast amount of trust on the part of all concerned and rests on peer enforcement is fragile at best; thus while in the process of attempting to create such a culture a few ‘bad apples’ could ruin the barrel.

A key feature of the Caltech honor code is peer enforcement in two forms:

First, if a student sees an honor system violation it is his or her responsibility to report it. *Failing to report an honor system violation is, in and of itself, an honor system violation.* This meta-norm works best if there is a culture in which new members of the community are indoctrinated into the system, under the assumption that the code generally works so that violations are rare and will be noticed. (If people thought the code was not followed, then violations would congest the system and violators would not worry as much about being caught and brought to justice.) Perception is incredibly important: if students, faculty or staff do not believe that the code is followed – whether or not it actually is - they will not feel obligated to follow it themselves. Thus a strong sense of community—upholding the code-- is necessary to maintain the code.

Second, adjudication of honor code violations is conducted entirely by students, with minimal faculty oversight. This saves faculty time, and also means the students know they are being judged by peers. For the students who participate in the Board of Control (BoC) process, it provides a way to contribute to campus life and learn to cope with emotional, nuanced issues of morality which helps prepare them for tackling these questions later in their scientific, corporate or public life.

While the benefits of the honor code come in a variety of forms, one of the most important is reduced monitoring costs. Unproctored exams liberate professors to focus on their research instead of wasting time fruitfully looking for cheaters. Conversely, students benefit from the chance to take exams in a situation of their own choosing – blasting Nine Inch Nails on their stereo, sitting next to a tranquil stream, or sitting at a desk in the classroom where they learned the material. Another benefit comes in the form of gift exchange. Students feel that they are trusted and thus feel obliged to return the favor in the form of not cheating. Graduate TAs are freed from deeply probing for similarities, instead only looking at something if it catches their eye.

The honor code has another huge advantage at Caltech: Most students at Caltech study science and engineering, and many will pursue lives working in universities or corporations. These students will work in environments where they have enormous personal responsibility and are relatively unsupervised (e.g., managing a large R&D lab). The pressure to cut corners, take unfair credit, downplay results which discredit a pet theory, or even fabricate data can be severe. The honor code enables scientists-to-be to “practice” and exercise their ability to do the right thing in a relatively benign environment, as a warmup to the outside scientific world where the moral issues are often subtler and the stakes higher.

Surprisingly, the vast majority of the honor system violations are blatant and occur when someone is under extreme stress and “cracks” (perhaps a close family member is diagnosed with terminal cancer or a similarly grave tragedy occurs).

The other main source of violations is ignorance. For example, a student from a foreign country may not understand the idea of plagiarism and might submit a paper that heavily relied on a couple of sources and failed to cite them. The increased atmosphere of trust also allows for other benefits, people can accidentally leave their bag around campus with hundreds of dollars inside and come back hours later to find it still sitting there unperturbed. The unspoken threat of social ostracism of a cheater tends to be enough to keep those who would still consider cheating in line.

There are drawbacks too, of course. Doubtless, without the close scrutiny a higher percentage of those who do cheat manage to ‘slip through the cracks’ and not get caught. For those 15-20 cases of suspected cheating each year that do get reported the chair and secretary determine whether or not the issue at hand is an honor code issue, and whether enough evidence is present to warrant the attention of the Board of Control.

Examples of something that would not warrant the attention of the Board would be someone being rude to someone else or a case of suspected copying when the two exams had nothing in common. The BoC then views the evidence, interviews any potential defendants and witnesses and then makes three decisions – whether or not an honor system violation has occurred, how to nullify any unfair advantage gained, and how to protect the community from such a thing happening again. These are recommendations to the deans of students, and are accepted about 95% of the time.

Because there are so few cases each case is treated meticulously, in a time consuming process. An average case of cheating or “overcollaboration” that goes to the BoC takes about 60 person-hours to resolve. The fact that cases are taken so seriously also upholds faith in the system—a student who loses a BoC case can at least take consolation in the fact that the system worked hard at adjudicating the case, even if the outcome is unpleasant for him or her.

END SIDEBAR

5. Summary TBA

Homework

Homework #2: Due 130 Wed 15 October.

ch 9 (3, “What is the relation..”,6, “Two basketball players...”,10, “Suppose you...tennis balls”)

Additional question 1. The President of Red Lobster (645 middle-end seafood restaurants) was recently fired when a promotion backfired. The chain offered “all you can eat” specials on Alaskan King Crab (and side dishes) and charged \$20. Describe briefly the mistake the chain probably made in offering the promotion. (see http://www.stpetersburgtimes.com/2003/09/26/State/All_you_can_eat_wa_s_t.shtml for details).

2. When you rent a car, it is usually very expensive to buy special insurance to cover any damage to the car (e.g. on a recent trip to New Mexico they quoted a rate of \$17.95 per day!?).

(a) Explain why, *in theory*, adverse selection and moral hazard (or “hidden information and hidden action”) could justify this large charge.

(b) In fact, there is little evidence of adverse selection or moral hazard in regular car insurance. (That is, drivers who buy insurance do not have more accidents than uninsured drivers, controlling for observable variables like the driver’s gender, age, and accident record.) Give a psychological explanation of why the two “hiddens” (information and action) do **not** lead to more accidents, contrary to the prediction in part (a) above?

Wall Street Journal (1996). Building blocked, A1. December 12 (byline J. Sapsford).

Business Week. (1999). That Russian company you bought? Maybe you didn’t. p. 70, December 13 (byline M. Coker).

Los Angeles Times (2004). Militiamen ‘reclaim’ oil for Nigerians. July 19, p. C3 (byline Michael Peel).

Hoff, Karla and Joseph Stiglitz. (2004). After the big bang? Obstacles to the emergence of the rule of law in post-Communist societies. American Economic Review, June, 94, 753-763.

Staten, Michael E. and John Umbeck. 1982. *American Economic Review*, 72(5), 1023-1037.