DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

STRATEGIC LEARNING AND TEACHING

Colin F. Camerer

Tech H. Ho

SOCIAL SCIENCE WORKING PAPER 1100

August 2000

Strategic Learning and Teaching

Colin F. Camerer

HSS, Caltech

Pasadena, CA 91125

camerer@hss.caltech.edu

Teck H. Ho

The Wharton School

University of Pennsylvania

Philadelphia, PA 19104-6366

hoteck@marketing.wharton.upenn.edu

"A man who carries a cat by the tail learns something he can learn in no other way."
--Mark Twain

Decision makers learn from experience and this learning affects their future decisions. But how? Using game theory, the authors explore how people factor in the payoffs of past decisions in their current choices. The authors introduce the concept of experience-weighted attraction (EWA) to explain how people learn from their experience. Sophisticated players can use this concept to learn which strategies work and "outguess" their rivals. And by understanding how rivals and customers learn, managers can take actions that "teach" rivals and customers what you want them to believe— reassuring partners and intimidating competitors.

Every decision is influenced by learning. The price a firm bids in a B2B auction, a complicated wage-bonus package offered to an executive, and the scale of investment in new business are all influenced by what has worked in the past. Airlines have struggled for years to coordinate industry-wide pricing practices, effectively teaching one another to stabilize fares and keep airplanes full, while skirting laws against explicit collusion. Firms figure out how to balance risky portions of compensation packages (such as options and company stock) with fixed salary portions, to both motivate employees and keep them happy during a stock market downturn. Internet buyers clicking through Websites use trial-and-error-learning to figure out how to get what they want from the web site. Firms also try to teach customers through advertising, service experiences, and so forth.

Understanding how customers learn is invaluable for evaluating marketing strategies. Understanding how rivals learn is crucial for developing competitive strategies. How do managers learn from their own experiences and from one another to make these decisions? How can they effectively teach customers and rivals to affect their decisions?

THEORIES OF LEARNING

Learning is a change in behavior based on experience. How do managers learn? According to economic theories, idealized decision makers run through complicated calculations to derive an optimal decision. But in practice people typically learn by trial and error. Firms try out a group bonus scheme one year, and abandon it if productivity does not improve. Construction companies bid on contracts and learn to build in expected cost overruns after losing money. Firms learn that hiring friends referred by current employees improves morale and productivity, and reduces turnover; so they increase their inside-hiring bonus from $1,000 to $5,000. Financial markets learn, cyclically, whether initial public offerings (IPOs) are good investments or not. And consumer products companies learn whether a successful yogurt brand name can be used to umbrella-brand ice cream, chocolate, or pudding.

In all these cases, not only are people and firms learning how to do better, but what they are learning depends upon what others--consumers, competitors, regulators--are doing.

There are three primary theories of how this kind of interactive learning takes place:

- ***Reinforcement:*** From the 1920s to the 1960s psychologists asserted that the primary way animals and people learn is by reinforcement. If actions are positively reinforced, they are taken again; if they are negatively reinforced, they are not taken again. Reinforcement is undeniably important, especially for learning by animals and children, and for certain kinds of human learning which derive from the "old" part of the brain (e.g., learning what to fear, and developing drug addictions). But later research in psychology showed that humans—and even pigeons-- are sensitive to whether actions they did *not* take would have yielded positive reinforcements. Reinforcement theory omitted this kind of learning from the road not taken.

- ***Belief learning:*** Beginning in the 1950s, learning theorists proposed that managers learn by developing a precise guess (or belief) about what other managers will do, based on observations of the past actions of those managers. Then managers choose a strategy that will yield a good payoff if their belief proves to be right.

- ***Experience-Weighted Attraction (EWA):*** Experience-weighted attraction (EWA) learning theory improves upon these two leading theories by combining their key features. It is based on the concept that strategic choices have different "attractions" for managers (as reflected in a numerical value), based on past experience. The EWA theory is embodied in a mathematical model[1], which is beyond the scope of this chapter, but we will highlight its key features.[2]

The concept of attraction incorporates both the extra reinforcement based on a manager's firsthand experience and the beliefs about the actions of others.[3] The way in which attractions are adjusted to reflect learning is based upon three factors: consideration, change and commitment, as discussed in more detail below.

The EWA, reinforcement and belief theories have been compared in 31 separate sets of data from experiments (conducted by us, as well as six other researchers). In 90 percent of those studies, EWA explained and predicted what people actually chose more accurately than either of the other theories. Experiments also show that players achieve higher payoffs by using the EWA model to forecast the strategic behaviors of others.[4] For example, in one study 43 of 54 players using EWA earned more money than counterparts not using this approach.

Learning Styles: Consideration, Change, and Commitment

Different managers learn from their experiences in different ways. These different learning styles – and the attraction levels assigned to different strategies – are based on three factors: Consideration, change, and commitment.

Consideration Index

People hate missing a plane by five minutes more than they hate missing it by an hour, because they can more easily imagine ways they would have caught the plane; consideration of a missed opportunity is far stronger for a near miss. In the EWA theory, the consideration index is a measure of the relative weight people give to foregone payoffs in past decisions, or how vividly they imagine lost value from missed opportunities. If their imagination is hazy, lost value does not weigh very heavily in learning. If, however, there is powerful emotional regret, as when the passenger is standing at the gate watching the plane taxi onto the runway (or a competitor lands a big account you spent months wooing), consideration will be a much more significant factor in learning.

Consideration of lost opportunities is also embedded in legal principles such as negligence. In an accident, the party who had the "last clear chance" to avert the accident is usually held liable, because it is easiest to imagine the accident being avoided by that party's last-second action.

Economists often refer to the value of a lost opportunity as an "opportunity cost." When the consideration index is higher, managers switch more quickly to the strategies they wish they would have picked, minimizing their opportunity cost or regret.

Smart learners have a high consideration index, which means they are actively considering what they should have done after every move. Learners with a low consideration index can get stuck using a strategy that is not necessarily bad, but is not nearly as good as alternative strategies.[5] In artificial intelligence programs that help machines to learn, clever programmers avoid getting stuck with suboptimal outcomes by having the machine experiment periodically.

Change Index

The second factor that affects the attraction of past strategies is the manager's perception of how rapidly the environment is changing. When change is rapid, players attach little weight to old experiences, because what worked years ago (or months, in internet time) is irrelevant. When the environment is stable, managers can learn over longer periods, weighing previous successes and failures equally. In this case, managers can use a long history of payoffs to choose a strategy.

In learning, information is like machinery. The change index determines the obsolescence policy for disposing of old information, much as companies develop obsolescence policies for replacing old equipment. When technology quickly makes machinery obsolete, managers retire old machinery and upgrade quickly; when obsolescence is slow, they hold onto old machinery longer. When the world is changing rapidly, you should "retire" past experience more rapidly. This advice implies that different businesses in a single firm should be learning at different rates. Think of a restaurant that begins to advertise and accept reservations on its website. Learning about how to operate the website requires very rapid adjustment and learning (which implies forgetting about what worked a year or two ago). But knowing how to train waitstaff and buy fresh produce and seafood benefits from the wisdom of many years of experience.

Commitment Index

The third factor that affects learning is how quickly managers lock in on a strategy that has performed well. If, for example, a given strategy has produced consistently high payoffs, and the environment is relatively stable, then it makes sense to commit to choose that strategy all the time. But if the relative performance of strategies seesaws back and forth, and the environment is changing rapidly, then it is a mistake to commit too strongly to one strategy.

In the EWA theory, commitment is expressed by judging a strategy based on its cumulative performance, simply totalling up all its successes. If attractions are cumulative payoff totals, then a strategy which is good enough to be chosen often will pile up more and more wins, compared to lesser strategies which are not tried much. This process is like the opposite of handicapping a horse race. To make a horse race even, the best horses are forced to carry more weight to slow them down. But suppose you wanted to guarantee that the best horse wins. Then you would want to give less weight to the fastest horses, to give them an advantage based on previous successes. Similarly, suppose you were deciding which scientist in your R&D lab to devote resources to. During an initial trial period, you might judge new scientists by their average output. Then it makes sense to commit more resources to the scientists who did well initially, and judge scientists by their total output.

The EWA theory specifies a precise mathematical way in which past experiences are combined using the consideration, change, and commitment indices. These indices correspond to different learning styles. For example, a manager who is constantly second-guessing and pointing out mistakes has a high consideration index (which is good for learning, but may have other organizational costs such as demoralizing workers and discouraging risk-taking). A manager who always surfs the latest trend has a high change index; one whose mantra is "if it ain't broke don't fix it" has a low change index.

SOPHISTICATED LEARNING

Sophisticated players are aware of the learning process. They use an understanding of how others learn– and the factors of consideration, change and commitment – to anticipate the decisions of their competitors. Even if they don't apply the formal EWA model in their decision making, they use this deeper understanding to outguess their competitors and coordinate with partners.

All decision makers become more sophisticated over time. In our experimental studies we find that a small fraction of players are sophisticated learners, but that fraction grows as subjects play similar games repeatedly. They learn about how to play better and also to anticipate how others are reacting and learning.

How does the presence of sophisticated players affect competitive interactions and outcomes? To explore this issue, we examine a well-known competitive game – the beauty contest.

Beauty Contest Games

Economist John Maynard Keynes once compared investment in the stock market to a kind of beauty contest popular in British newspapers in the 1920s.[6] In these contests, a newspaper printed pictures of 100 people and asked readers to name their favorites. All those who chose the face that proved to be most popular were eligible for a prize. Keynes pointed out that the goal was not to pick what *you* considered to be the prettiest face, but to pick the face that *the most people* would pick. He noted that stock market investment is similar--just substitute “hottest stock next week” for “prettiest face.” While Keynes' metaphor is nearly 80 years old, it is perhaps no more apt than in the modern-day era of tech stocks valued by "new paradigm" principles.

Based on this idea, researchers have developed a type of competitive game called the “beauty contest.” The game has been studied on three continents with several samples of sophisticated adults, including economics Ph.D.'s, CEOs and corporate presidents,[7] and readers of business publications including *The Financial Times.*[8] It illustrates some important properties of strategic thinking in naturally occurring settings.

In the beauty contest game we analyze below, a group of players simultaneously choose numbers between 0 and 100. The player whose number is closest to two-thirds of the average wins a fixed prize. Since all players should choose two-thirds of what they think the average number will be, any kind of introspective reasoning will lead to lower and lower numbers. A typical player might reason as follows: Suppose the average is 50. Then I should choose 33. But if everybody does that first step of reasoning, then the average will be 33; so I should choose two-thirds of that, which is around 22. But if everyone chooses 22, should I choose two-thirds of 22? Where does this thinking stop?

In standard game theory, there is no natural stopping place until one reaches the “Nash equilibrium,” the point at which all players' strategies are profit-maximizing best responses to one another. In this case, the Nash equilibrium is zero. This assumes all the players are acting rationally, according to the dictates of game theory.

In practice, however, everyone is not rational, nor do we expect them to be. Some might be confused and choose around 50, or just choose a lucky number. Others will do one step of reasoning and pick 33. Others will do two steps and come down to 22. The number of levels of reasoning people are likely to use is a question that can only be answered by psychology and observation, rather than by pure mathematics. (The early game theorist Tom Schelling said that many questions in game theory cannot be answered by pure logic. Those questions can only be resolved by putting people in strategic situations and watching what they do, much as you cannot prove a joke is funny until you tell it and people laugh.) Furthermore, a smart player's goal in the game is to be one step smarter than the average person, but no smarter!

In fact, experiments on the beauty contest game have shown a surprising amount of regularity. Figure 9.1 shows some distributions of choices from four interesting groups of subjects. Each interval is a range of five numbers (1-5, 6-10, and so forth). The height of each bar represents the proportion of subjects choosing numbers in that interval. It is easy to see that initial choices are widely dispersed and centered somewhere between equilibrium (in this case, 0) and 50. The average tends to be between 20 and 35, so that choosing a number between 14 and 25 would give you a better-than-average chance of winning. These numbers correspond to one or two steps of reasoning from the presumed average of 50. In fact, about a dozen experimental studies have shown, with very different strategic situations and different subjects, that people typically engage in one to two steps of reasoning. One group of subjects are Caltech undergraduates, who have spectacular analytical skill (and know that their peers do too); they choose numbers about one step closer to the equilibrium than the other groups-- CEOs, portfolio managers, and economics Ph.D.'s.

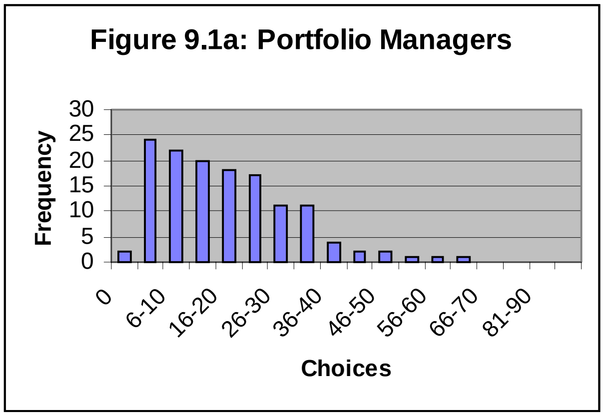

Figure 9.1a: Portfolio Managers

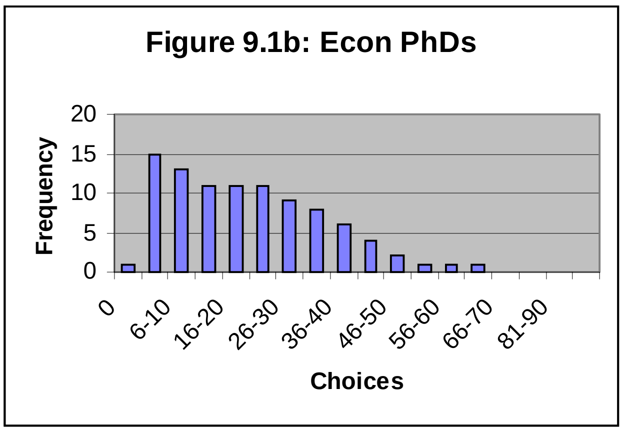

Figure 9.1b: Econ PhDs

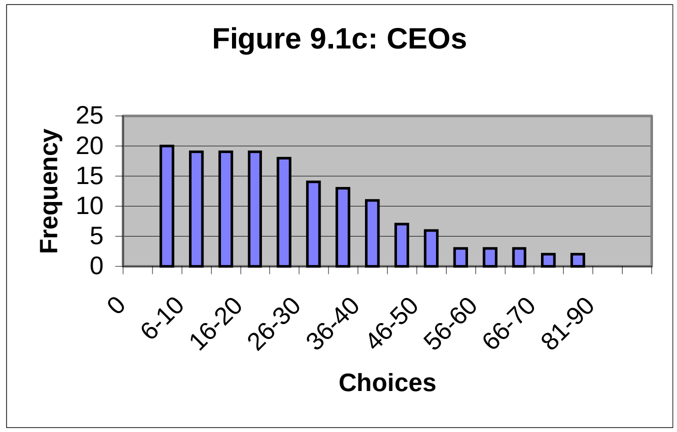

Figure 9.1c: CEOs

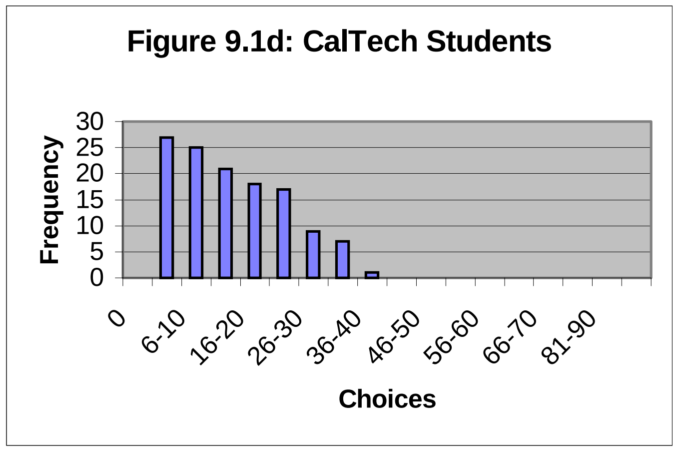

Figure 9.1d: CalTech Students

Players learn as they continue to play. When the game is repeated, numbers gradually converge toward the equilibrium prediction of zero. That prediction is useful as a long-run benchmark, but it is not very helpful for explaining the trajectory of learning or for giving advice on how to play in the first few periods. Fortunes can be made and lost getting to equilibrium.

What happens when sophisticated players enter these games? In any given game, some players are very sophisticated and others are not. The sophisticated players understand the learning process. Suppose the average in one period is 30, so that choosing 21 would have been a winning number. In the next period, most players tend to choose 21, but sophisticated players guess others will choose 21 and choose 70 percent of 21, around 15, leapfrogging ahead. (In the internet economy, for example, the trick is not to be one step ahead of the competition, but always planning for the next step after that one.) The more sophisticated the players, the more rapidly the group will converge toward the equilibrium point of zero. But what is important is that the sophisticated players will lower their number if they believe other players are unsophisticated. Remember, they are trying to anticipate how others will act. For example, if these sophisticated players could actually use the EWA model to predict every player's choice in every period, they would know how to make better choices and would win more often.

Figure 9.2a: Actual choice frequencies in beauty contest game

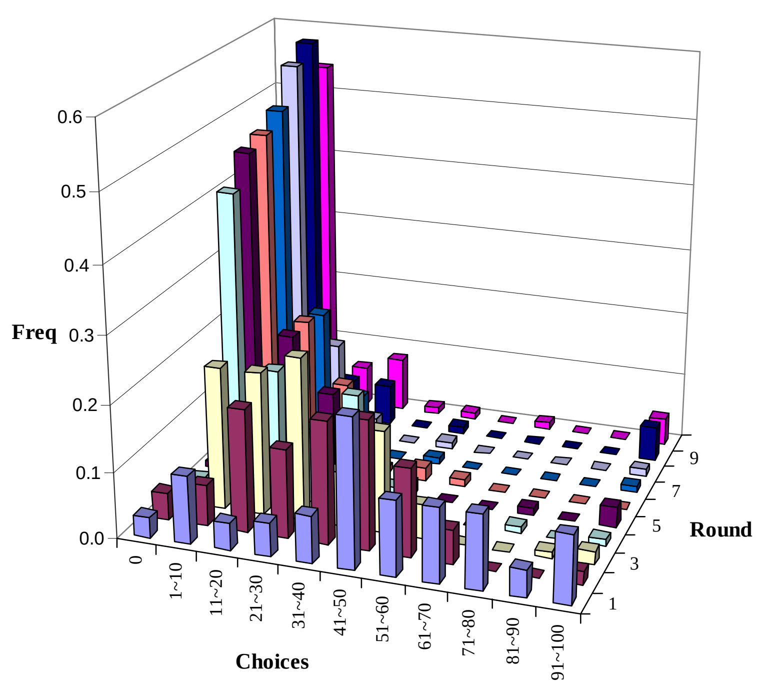

Figure 9.2b: Predictions of EWA theory in beauty contest games

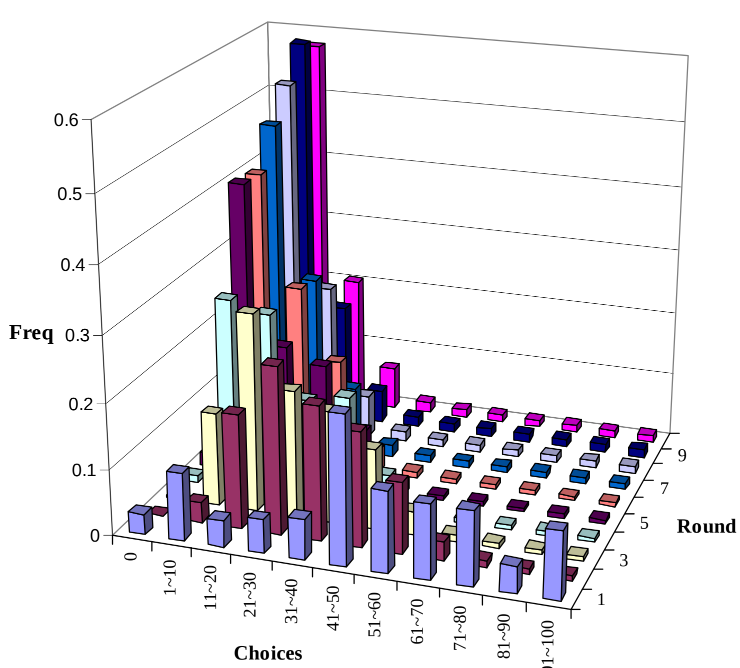

Figure 9.2a shows the distribution of actual choices by subjects who played the game 10 times. The axis on the front left shows intervals of numbers players might pick. The right axis is the period of the game (from 1 to 10). The proportion of players choosing lower numbers grows and grows across the 10 repetitions-- that is, the bars representing low-number choices grow taller and taller-- as they learn. Figure 9.2b shows statistical predictions of the EWA theory (that is, what players are predicted to do). The two graphs look very similar, which means that the theory is doing a good job of predicting how players will actually learn.

STRATEGIC TEACHING

This same process used by managers to *learn* from past experience can also be used to *teach* rivals and customers. If a sophisticated player thinks other players are adaptive, and will respond to their recent experiences, it can then shape those experiences to its own advantage. The sophisticated player has an incentive to choose strategies in the current period that affect how the adaptive players respond in future periods, if those responses benefit the sophisticated player. We refer to this as strategic teaching.

Strategic teaching goes on all the time in business. For example, Wall Street security analysts tend to hammer the share price of a company that reports disappointing earnings--below what was expected. Managers have an incentive to manipulate expectations for earnings as best they can (using accounting methods to keep some potential earnings “in their pocket”). They are effectively trying to teach the Street that earnings will be low so analysts will be pleasantly surprised when the earnings come out higher.

In this process, it is important to understand what a rival is learning from your teaching to avoid teaching too much or too little. Think about a supplier trying to reassure a new client of its goodwill by providing extra goods, lowest-price guarantees, and so forth. Some clients will easily learn that the supplier is trustworthy, so repeated concessions to establish trust are a waste of money (the supplier is over-teaching). Other clients will never quite trust the supplier, so it may be better to sever the relationship. The EWA model offers a way to quantify what another firm or person is learning from the current strategy.

There are many other important potential applications of the idea of optimal teaching. In economics, it is well-established that a spurt of unexpected inflation can lower unemployment (this is called the Phillips curve). As a result, policy makers who benefit from lower unemployment--such as a sitting administration just before a presidential election-- are always tempted to increase inflation. Policy makers who take a long view--such as a Federal Reserve Board chairman with a long-term perspective --would like to build up credibility with the public by keeping price inflation low for years at a time. This certifies his ability to resist the temptation to inflate. In our terms, the policy makers have an incentive to teach the public that inflation will be low in the future. This basic idea has been used to explain the substantial shift from the high double-digit inflation of the 1970s in the U.S., to a more recent post-1985 drop in inflation.[9]

Another example is pricing of goods that fall in price rapidly, such as computers and new high-tech products like cellular phones and personal digital assistants. In these markets, consumers may delay purchase if they believe prices will fall rapidly. Since firms would like to sell goods earlier rather than later, they have an incentive to teach consumers that prices will not fall that rapidly, so they will buy right away.

The Prisoners' Dilemma

A sophisticated player who is teaching other players may even deliberately choose a bad strategy in the current period, if choosing that strategy yields a better return in future periods. In a sense, the current period choice is an "investment" in educating the other players, which pays off for the teacher in the future. To explore this, consider the well-known example of the repeated prisoners' dilemma.

In the prisoners' dilemma, two prisoners are arrested. They have the choice of either cooperating with one another or defecting and ratting on the other. If they cooperate, they each receive an equal payoff of 3, as shown in Table 9.1. If one defects and other doesn't, however, the defector is then rewarded with a payoff of 5 while the betrayed partner is left with zero. If both defect, then they both end up with 1. The dilemma is that both players can do better if they cooperate. But if they cannot be sure what the other player will do, they are better off defecting. At least that guarantees a payoff of 1 and perhaps of 5.

Table 9.1: Prisoners' Dilemma Payoffs

	Cooperate	**Defect**
Cooperate	3,3	0,5
Defect	5,0	1,1

In the repeated prisoners' dilemma, players engage in this same game over several sessions. Players usually cooperate in the initial periods. Then one of the players defects as the end draws near, and both players defect until the last round. How can we explain this?

Strategic teaching explains it by assuming that some players are adaptive, and learn, while others are sophisticated and try to teach the adaptive players to cooperate. If the teachers can convince the adaptive learners to cooperate, then everybody can get payoffs of 3 for several periods.

How does the teaching work? The sophisticated player starts by cooperating, choosing a strategy that maximizes the sum of the current payoff and expected future payoffs (assuming the adaptive players learn from what the sophisticated players do in the beginning). Suppose the adaptive players start out with equal initial attractions on cooperating and defecting, so they are equally likely to choose either one. If the adaptive students choose to defect, they earn 5, and are likely to continue to defect in the future.

Anticipating this, the sophisticated teacher may well give up, and both players defect throughout. More interestingly, suppose the adaptive player starts by choosing to cooperate. Then both players earn a payoff of 3, but also reinforces defection (since they know they would have earned 5 if one had defected). Thus, strategic teaching can explain why there is mutual cooperation for several periods in the repeated prisoners' dilemma. But how does it explain the fact that cooperation usually breaks down a couple of periods from the end? Remember that sophisticated teachers maximize the sum of their current payoff and expected future payoffs (knowing how their current actions "teach" the adaptive players). When the end draws near, there are not many future periods of cooperative payoffs left, and the strategic teachers are tempted to defect to earn the largest payoff of 5. How likely adaptive players are to switch to defect will depend on the amount of cooperation they encountered previously. Thus, it pays for sophisticated teachers to cooperate until close to the end, to build up the chance that an adaptive player will still cooperate in the last period (even after a defection), and to defect a couple periods before the end.

Strategic teaching is sometimes beneficial for both players--as when one firm tries to reassure a business partner, who learns from the firm's past behavior that it can be trusted in the future and that trust enables the two sides to do business profitably. In other cases strategic teaching benefits the “teacher” by enabling the teacher to exploit or manipulate the “student” who is learning. Examples include developing a reputation for aggressively pricing when there is new competition.

CONCLUSIONS

An important component of good strategic advice is a theory of learning. Learning is important because in many naturally occurring strategic situations players do not know their payoffs at all; therefore they can learn more from experience than from deduction and calculation. Understanding how others learn also is important because sophisticated players who know how others learn can stay a step ahead of them. Sophisticated players can teach others what to expect in the future, in a way that often benefits the teacher.

Since its discovery about fifty years ago, game theory has been hailed as *the* mathematical language for analyzing social situations. Unfortunately, an obsession with mathematical details of how rational players would play a game after careful introspection have kept game theorists from tackling the more practically useful question--how do people actually play? And knowing that, how *should* people play? The EWA theory was created because other learning theories left out important pieces of this puzzle. We also had a nagging sense that the two most prominent kinds of learning--learning by reinforcement of successful strategies, and learning by forming beliefs about other players' behavior and responding to them--might have something fundamental in common. Furthermore, we know statistically from analysis of 31 experimental data sets that leaving out any of the three learning-style features (consideration, change and commitment) falls short in explaining how people actually learn.

Implications

What are the implications of this theory for managers? In any competitive interaction, look for opportunities to teach and to learn from the situation. Game theory can help think through the dynamics of multiple rounds of interactions. Don't just look at the impact of poor choices but also consider the impact of regrets over choices that were not taken.

- **Learn how to learn:** As you become more sophisticated in your learning you will better be able to make your own decisions and outguess your competitors. Understand how the three factors of consideration, change and commitment affect learning. For any given decision, identify how these factors will influence your learning about it and decide on what learning style is best. Considering the opportunity costs of all strategies you have not tried is generally wise. Applying a low change index is smart when the world is changing rapidly, but drawing on a large span of experience (i.e., a high change index) is smart when the world is relatively stable. Committing strongly to a particular strategy is wise only if you become convinced from its string of successes that no other strategy is likely to prove better.

- **Understand learning patterns of others:** Knowing how competitors respond to experience enables firms, in principle, to stay one step ahead. Recent evidence from academic studies of financial markets shows that investors tend to underreact to certain kinds of events in the short run, and overreact in the long run. These patterns open up the possibility that a theory of how investors learn could enable investment managers to beat the market.[10]

- **Assess the sophistication of other players:** Determining whether other players are sophisticated or not offers insights into the strategies you should take with them. If players are unsophisticated, how can you use teaching strategies to "educate" them in the direction you want to go?

- **Try using a statistical model:** While we have left out the mathematical details of the EWA theory, applying it statistically may improve a firm's ability to forecast learning of others. Many statistical models are used in business, like the Black-Scholes options pricing formula (and many subtle versions of it), models forecasting consumer purchases, and models forecasting interest rates and other macroeconomic variables. These models are successful because they combine variables in very precise ways, which sharpen managerial intuition. A variant of the EWA theory was applied by one of us (Ho) and Kuan Chong[11] to predict consumer choices of ice cream, diapers,and other consumer products. The theory assumes that people learn what brands they like by trial and error and generalize from experience with one brand to related brands which are similar. The analysis used 130,000 choices by consumers, and was able to predict choices about 10 percent better using only one-fifth as many variables as the previous "best practice" model. (Ten percent is a small improvement, but in consumer forecasting even a 1% increase adds lots of value.) A large supermarket chain rushed the model into their business immediately (before Ho and Chong's paper was even published in a journal!).

Learning has become a central focus of organizations. But less attention has been spent on the process by which people learn – particularly from their own experience and that of others. The EWA learning theory is a powerful tool for strategic decision making for two reasons: First, it enables a company to stay a step ahead of competitors by predicting what they will do. Second, it enables a company to teach competitors how the company will behave in the future, which is useful for both creating trust among partners, and signaling toughness to competitors.

NOTES

[1] Define the attraction for strategy j, after period t's experience is incorporated, as $A^j(t)$, and define the payoff from strategy j in period t to be •(j,t). In the EWA theory, attractions are modified each period according to the formula, $A^j(t)=(\bullet A^j(t-1)+\bullet(j,t).)/(\bullet(1-\bullet)+1)$, if strategy j was chosen in period t, or $A^j(t)=(\bullet A^j(t-1)+\bullet\bullet(j,t))/(\bullet(1-\bullet)+1)$, if strategy j was not chosen. Attractions are then used to determine the chance of playing a particular strategy through an exponential "logit" formula. See Camerer and Ho (1999) for more details. The parameter • is the change index, • is the consideration (or imagination) index, and • is the commitment index.

[2]

[3] Reinforcement learning and belief-based learning are thus closely interrelated: belief learning is a kind of <u>general</u> reinforcement in which even unchosen strategies are strongly reinforced by what they would have earned. This surprising kinship came as a shock to theorists who had thought the two kinds of learning were fundamentally unrelated, like adventurers discovering that two wide rivers in different parts of a country spring from a common source. Having a single unified theory economizes by replacing two different tools with a single tool that can do what the other two could do.

[4] Colin F. Camerer and Teck-Hua Ho, "Experience-Weighted Attraction Learning in Normal-Form Games," *Econometrica* (June 1999); John Van Huyck, Raymond Battalio and Richard Beil, "Strategic Uncertainty, Equilibrium Selection and Coordination Failure in Average Opinion Games," Quarterly Journal of Economics, vol. 106 (1991), pp. 885-909.

[5] This behavior was called `satisficing' by the behavioral economist and Nobel laureate Herbert Simon. Simon's point was that people and organizations often behave this way, while our purpose is to help managers learn faster and understand how others learn

[6] John Maynard Keynes, *The General Theory of Employment*, *Interest and Money*, New York : Harcourt, Brace, 1936.

[7] Colin F. Camerer, "Progress in Behavioral Game Theory," *Journal of Economic Perspectives* (July 1997), pp. 1-48.

[8] R. Nagel, "Experimental Results on Interactive Competitive Guessing," American Economic Review, vol. 85 (1995), pp. 1313-1326.

[9] Sargent, Thomas J. The Conquest of American Inflation. Princeton: Princeton University Press, 1999.

[10] Note that most money managers *underperform* the market (about 75 percent of actively managed mutual funds underperform the S&P 500 after fees are subtracted).

[11] Ho, Teck-Hua and Kuan Chong, "A parsimonious model of SKU choice: Familiarity-based choice and response sensitivity," Wharton School Department of Marketing, 1999, or http://www-marketing.wharton.upenn.edu/ideas/pdf/99-020.pdf