Cognition and the Theory of Learning by Doing

W. Bentley MacLeod

USC Olin Research Paper No. 00-11

USC
THE LAW
SCHOOL

RESEARCH PAPER SERIES

Sponsored by the John M. Olin Foundation

University of Southern California Law School
Los Angeles, CA 90089-0071

Cognition and the Theory of Learning by Doing

W. Bentley MacLeod*
Department of Economics and The Law School
The University of Southern California
Los Angeles, CA 90089-0253
wmacleod@usc.edu

April 2000
Keywords: decision theory, learning by doing, incentives

Abstract

The standard model of utility maximization assumes that individuals make complete contingent plans before making a decision. This paper shows that when the environment is sufficiently complex the optimal rule entails learning by doing rather than contingent planning. An implication of this result is that spillovers from learning by doing are less than perfect, a hypothesis central to many models of learning by doing, such as Rosen (1972). It is shown that one can derive learning curves that can be fitted to data. The model is also consistent with the observations of behavioral economics, such as Kahneman and Tversky (1979), that in the short run individuals may make biased decisions, though long run behavior is consistent with utility maximization. Finally, the model is applied to Becker (1968)'s deterrence model to show that the optimal amount of monitoring is bounded away from zero.

*I would like to thank Boyan Jovanovic for supplying me with his data, George Akerlof, Colin Camerer, Janet Currie, Dick Day, Rachel Kranton, Steve Postrel and Louis Putterman for helpful discussions on the nature on learning, and Mehdi Farsi for able research assistance. I am also grateful for the comments during presentations of earlier versions of this work at the American Economics Association Meeting, Chicago 1997, a McArthur Foundation Conference, Hoover Institute, Stanford, August 1998, and Conference on Cognition, Emotion and Rational Choice, Center for Comparative Political Economy, UCLA, April 2000. The financial support of National Science Foundation grant SBR-9709333 is gratefully acknowledged.

1 Introduction

Beginning with Knight (1921) and Simon (1955), there has been a great deal of work exploring the implications of human cognition for economic decision making.[1] Despite many advances in our understanding of human cognition, it is still the case that expected utility theory appears to be the best general model of choice behavior (see for example Grether (1980), Harless and Camerer (1994) and Hey and Orme (1994)[2]). The goal of this paper is to develop a simple, stylized model of decision making that begins to bridge the gap between models based upon utility maximization and behavioral models, whose goal is to explain observed deviations from rational choice.[3] The main result demonstrates that when the number of events for which an individual or organization must prepare a response is so large that complete contingent planning is impossible, then the optimal Bayesian decision rule is to acquire behavior via learning-by-doing. The resulting model can be empirically implemented, and is calibrated to fit some empirical learning data. It also provides some new insights into the structure of optimal incentive rules, and can explain why large fines are not used to economize upon monitoring costs, as predicted by Becker (1968).

The hypothesis of complete contingent planning is a central hypothesis in models of rational choice, such as Savage (1972)'s theory of decision making under uncertainty. The important contribution of Savage's theory is to show how that it is possible to construct a theory of rational choice even when there are events for which the objective probabilities are unknown. A central ingredient of this theory is that the decision maker is able to consider every possible future event and can assign a probability to each of these events. As Savage was well aware, both of these assumptions are quite strong.[4] In practice individual planning is necessarily incomplete, moreover, even if planning were complete, the subjective nature of probability assessments weakens the predictive content of the theory.

In section 3 it is shown that when the environment is sufficiently complex (and hence contingent planning is necessarily incomplete), and individuals have no prior information, then learning-by-doing is an optimal Bayesian decision rule. Moreover, this rule has characteristics that correspond to the features of a generic "cognitive" model of decision making.[5] The essential idea is that when the environment is very complex, then the most efficient way for an individual to plan is to wait and respond to events as they occur (an idea that is implicit in Simon (1951)'s theory of the employment relationship). In a stationary environment the events that are observed are more likely to occur again in the future, and hence it is optimal to acquire optimal responses only to those events that occur, rather than attempt to make plans for all hypothetical future events.

Arrow (1962) and Alchian (1963) have shown that learning by doing plays an important role in understanding both the nature of economic growth and the theory of the firm. An important feature of learning

[1] More recently Williamson (1975) has argued that understanding complexity is central to understanding the theory of organization.

[2] These papers find the even stronger result that simple expected value maximization is a good first order model of behavior.

[3] See Camerer (1995) for a recent review of this literature.

[4] See Binmore (1992) for a further discussion of this point in the context of game theory.

[5] See Churchland and Sejnowski (1993) for an excellent review.

by doing is that the acquired skills and techniques are often difficult and costly to transfer to other firms or individuals. Rosen (1972) uses this feature to construct a theory of the firm in which skills acquired through experience form part of the marketable assets of the firm.[6] The main result helps explain not only why learning by doing is optimal, but also why the spillovers are less than perfect. Acquiring the skill from another firm is equivalent to contingent planning, that is one needs to learn the optimal response for many possible contingencies, many of which may not be relevant for the organization or individual acquiring the skill. In the extreme case it may simply be more efficient to acquire the skill via a process of learning-by-doing, rather than attempt to buy the skills from others.

This result distinguishes the paper from Jovanovic and Nyarko (1995), who also use a decision theoretic approach to explore the foundations of learning-by-doing. In their model it is assumed that there is an unknown parameter that determines the optimal decision that can only be learned through costly experimentation. For example, one may not know what is the correct amount of fertilizer to put on a field, and hence it may take several seasons of experimentation to discover the optimal rule. But once the rule has been learned, it can be easily passed on to another person or individual. In their model it is the complexity of the *choice space,* rather than the *event space* that makes it expensive to learn.

In section 4 the model is compared to Jovanovic and Nyarko's results using their data. This data consists of observations on performance as a function of experience, and for the most part corresponds to skill acquisition: coronary anoplasty success, flight control ability and sales performance data. It is found that the model in this paper consistently fits as well or better than the Jovanovic and Nyarko model, except in the case of steel plant productivity. This result is consistent with the hypothesis that productivity in a manufacturing plant depends upon knowledge of a technique rather than the ability to respond to a large number of possible scenarios, and hence we should expect the Jovanovic-Nyarko model to do best in this case.

The basic model assumes that the individual or organization acquires the appropriate behavior for a finite number of events. However in practice, it is reasonable to suppose that stimulus comes from a continuum of possible events. Cognitive scientists explain the human ability to learn appropriate behavior by supposing individuals classify data into finite categories using sophisticated pattern recognition abilities. The precise form that these pattern recognition routines take is a very difficult problem and an active area of research. What is known is that under quite weak conditions a large class of pattern recognition algorithms result in convergence to the optimal rule. In section 5 it is shown that one can incorporate a simple pattern recognition procedure into the model, and, even with a continuum of events, learning-by-doing converges to optimizing behavior.

This result provides a formal bridge between the literature in behavioral economics that focuses upon the many observed deviations from rational choice, and the literature that finds that utility maximization is the most useful general model for explaining observed economic institutions. For example, the classic

[6] See also Lippman and Rumelt (1982) who develop a theory of industry dynamics under the hypothesis of imperfect ability to imitate competitors.

paper of Kahneman and Tversky (1979) documents a number of observed behavioral anomalies that can be explained by cognitive mechanisms relying upon pattern recognition. More recent work by Rubinstein (1988) and Leland (1994, 1998) have shown that some of the deviations from expected utility theory can be explained by supposing individuals use explicit similarity functions in evaluating alternatives. Thus, the use of pattern recognition in the short run can explain observed deviations from rational choice, and when combined with learning-by-doing implies in optimizing behavior in the long run; a result that is consistent with the evidence discussed by Grether (1980), Hey and Orme (1994) and Harless and Camerer (1994).

The penultimate section of the paper illustrates an application of the model to an optimal incentive problem. Standard principal-agent models assume that both the principal and agent are able to anticipate and plan for an agent's response to an incentive system. In the context of deterrence, Becker (1968) has shown that this implies that the cost of monitoring an activity can be made arbitrarily small by matching a decrease in the probability of detection with an increase in the penalty. Not only do we not observe a systematic application of this principle, but there is some evidence that the probability of detection has a small impact on behavior (Grogger (1991)).

The deterrence model is typically applied to illegal behavior, such as theft, speeding, tax evasion, etc. Notice that these activities are complex in the sense of this paper. When deciding on one's speed many situations affect the choice, including where police are likely to be monitoring, traffic and road conditions, and so forth. Increasing the probability of detection, increases the number of times an individual is caught, and hence experiences the fine. The theory of this paper suggests that behavior is more responsive to change with experience, and hence increasing the probability of detection should increase the impact that fines have on behavior. This observation is used to derive optimal level of monitoring for a simple model of deterrence. The final section of the paper contains some concluding comments.

2 An Example

Before proceeding to the general result, let us begin by considering a very simple example that illustrates the role of learning by doing in decision making. Consider two individuals playing the game of tic-tac-toe: they each take turns playing with one placing X's while the other places O's on a 3×3 board, with the first person placing a line of three X's or O's in a row winning. We can define payoffs by supposing the person who loses pays the winner \$10, with no money changing hands in the event of a draw.

Notice that even though the game requires a sequence of moves, each time a player moves her decision corresponds to choosing how to play in response to a board position. The actual choice may depend upon expectations of what will occur in the future, however formally a move is a function from the set of possible board positions to the choice of where to place a O or X, a feature shared by many decision problems.

Suppose that it is X's turn to play and the current board looks like:

$$\begin{array}{c|c|c} \text{O} & \text{X} & \\ \hline & \text{X} & \\ \hline & \text{O} & \end{array} .$$

Under the hypothesis that player X is "rational" she would choose the strategy that would ensure either a draw or a win for X. Economic models typically do not inquiry into the process of how the agent makes such a choice, rather it is simply assumed that an optimal choice is made, though there may be some unmodelled error.

The purpose of this inquiry is to better understand both the source of error, and how behavior adjusts to deal with past errors. For this board position consider the following two possible moves by X:

$$A: \begin{array}{c|c|c} \text{O} & \text{X} & \text{X} \\ \hline & \text{X} & \\ \hline & \text{O} & \end{array} \quad \text{or } B: \begin{array}{c|c|c} \text{O} & \text{X} & \\ \hline & \text{X} & \\ \hline \text{X} & \text{O} & \end{array} .$$

In both cases, in order to avoid a loss, O must block the possible diagonal three in a row by X. However, case A leads to a sure loss by X, while in case B, the outcome is a draw regardless of subsequent play (in both cases I am assuming O makes the obvious blocking move against that X).

Anticipating these outcomes is not difficult, and in fact tic-tac-toe is simple enough to be completely analyzed by hand (Nilsson (1998), page 199). However, despite this, if player X has little experience and is required to move quickly, then a mistake may occur. The fact that children enjoy playing tic-tac-toe suggests that mistakes occur, and sometime a player wins. This is more likely to be the case in situations, such as this one, that require some foresight. Thus the second ingredient of the model is the hypothesis that when individuals are required to decide under time pressure, mistakes will occur even though with more reflection an optimal choice might be made.

This situation is reflected in day to day life where we simply do not have the time to weigh every decision carefully, and hence mistakes are often made. Even though we do not have time to reflect and consider the consequences of every decision, even under time stress we do make *conscious* choices, corresponding to what the noted cognitive scientist, Allen Newell (1990) has called decisions in the *cognitive time frame* (milli-seconds to minutes). These decisions depend upon very fast pattern recognition routines in which the individual uses experiences from previous, similar, events to make a choice.[7] In constrast, economic models typically assume that decisions are taken in the so called *rational time frame* (minutes to hours), where the individual has the time to consider and weigh many of the consequences that may arise from a choice.

Decision making in the cognitive time frame is however ubiquitous. For example Simon and Schaeffer (1992) use a very elegant experiment to show that memorized responses, in addition to thought and planning,

[7] See Churchland and Sejnowski (1993) for an review of how the mind carries out these pattern recongnition routines.

is part of high level chess playing skill.[8] Due to the high costs of exploring the game tree, high level chess players learn to recognize a large number of positions as being either weak or strong, which then use to limit the amount of search. Even in the context of tic-tac-toe pattern recognition proves useful. If the player were rational in the sense of Savage, then she would construct a complete contingent plan at the beginning of the game, where there are in principle 3^8, or 6561 possible board positions. Of course symmetry and other considerations greatly reduce the number of positions to be considered, however it should be remembered that working out such simplifications itself requires sophisticated, time consuming reasoning.

Thus even in the case of tic-tac-toe, as in the game of chess, only a limited number of board positions occur in practice, and hence through repeated play the individual can learn to respond to only those positions that occur repeatedly. In the case of the example above one would very quickly learn to choose plays like B where intuitively one can see that X is blocking the formation of three in a row by both O's on the board. In fact the game of tic-tac-toe is so simple that the game playing programs merely implement the optimal strategy and do not use any reasoning at all. This approach is consistent with the approach of this paper, which i assume that in many situations human behavior is not the result of a conscious weighing of possibilities, but rather the result of implementing a learned response, which depending upon the context may or may not be optimal.

One of the empirical implications of this model of decision making is that skills are costly to change. If an individual has hardwired responses to a number of situations, then if the optimal response to the same situation changes the individuals will initially make mistakes until the new optimal behavior is acquired. The speed at which optimal behavior is acquired depends upon the complexity of the decision problem. We would expect that decisions with levels of complexity similar to tic-tac-toe to approach optimality very quickly, while some skills, such as chess playing, may take years to acquire. Other examples would include pilots who are taught how to respond appropriately and automatically when an emergency occurs, or surgeons who plan in advance how to carry out a surgical proceedure, and have in place contingent plans should certain emergencies occur. With experience these professionals learn how to respond to contingences they did not anticipate that occur during the exercise of their craft.

In summary, our model begins with the observation that with time constraints individuals may make errors, even if the underlying problem is very simple. Secondly, when the number of possible events is very large, then contingent planning may not be possible, and it may be more efficient to learn optimal responses to those events that occur in practice. One reason I began with the example of tic-tac-toe is to illustrate that these problems can arise even when perfect planning is in principle possible. Savage's theory predicts that individual would use complete contingent planning, however in many situations, such as tic-tac-toe, it may be more efficient to acquire the skill through experience. We would certainly expect most adults to play tic-tac-toe well after only of few minutes of play, even though a complete analysis of the game might take a

[8]They compared the ability of novices and experts to remember both legal and illegal board positions. They found that both groups were equally good at remembering a random board position, but the experts could perfectly recall board positions corresponding to actually outcomes of a chess game.

matter of hours.

The final ingredient of the model is based upon the fact that humans are designed to respond quickly and reliably to patterns, and hence individuals can acquire and implement optimal behavior as a reflex response rather than as reasoned consideration. This basic structure is illustrated in figure 1, and developed formally in the following section.

3 A Model of Optimal Decision Making in a Complex Environment

Consider a situation in which an agent at time t is expected to respond with a decision $d_t \in D$ in response to an event $\omega_t \in \Omega$, where Ω is a finite set of events. Without loss of generality the decision set is restricted to a binary choice: $D = \{0,1\}$. In addition, it is assumed that the events are selected by an *i.i.d.* stochastic process, where $\mu(\omega)$ is the probability that ω is chosen in period t. The decision maker does not know this distribution, but updates beliefs over the set of possible prior distributions as events are observed. Suppose that the incentive contract governing the agent's reward yields a utility $U(d|\omega)$ if decision d is chosen when event ω occurs. Throughout, it is assumed that $U(1|\omega) \neq U(0|\omega)$, hence for each event ω there is a unique optimal choice given by $\sigma^*(\omega) = \arg\max_{d \in D} U(d|\omega)$.

It is assumed that the decision is sufficiently complex that there is not sufficient time between the observation of ω and the choice d for individuals to consistently choose the optimal strategy. However, before facing the event the individual may prepare a response either through an explicit contingent contract or by training with different possible events before the time of decision. For example, airline pilots train in flight simulators to prepare responses for various possible aircraft failures. Even if the pilot is capable of deducing the appropriate response for a particular failure, he or she may not have sufficient time to carry out such an analysis. Indeed, the point of pilot training is to ensure that the appropriate decision is taken quickly with little apparent thought.

The recent research in cognitive psychology, as described in Churchland and Sejnowski (1993), shows that the brain is designed to retrieve responses to learned experiences very quickly and reliably. Let us consider a highly stylized learning model based upon this idea that explicitly relates an optimal Bayesian decision making model to a cognitive model of decision making. At time t let $\Omega_t \subset \Omega$ denote the set of events for which the individual has a prepared response. This implies that if $\omega_t \in \Omega_t$ occurs then the individual is able to respond optimally with $\sigma^*(\omega_t) \in \{0,1\}$. If $\omega_t \notin \Omega_t$, then the individual does not have a prepared response, nor does she have sufficient time to determine the appropriate response and hence she randomizes over D. For simplicity suppose that when the optimal response is made the reward is u_g, while if a plan is not in place for an event the expected payoff is $u_b < u_g$. The *behavior* of an individual at time t is given by the function:

$$\sigma_t(\omega_t) = \begin{cases} \sigma^*(\omega_t), & \text{if } \omega_t \in \Omega_t. \\ \left\{\frac{1}{2}, \frac{1}{2}\right\} & \text{if} \omega_t \notin \Omega_t. \end{cases},$$

where $\left\{\frac{1}{2}, \frac{1}{2}\right\}$ denotes the lottery that selects each action in $\{0, 1\}$ with equal probability.

Events are added to the set Ω_t in two ways. First there is *learning by doing,* if ω_t occurs then the agent evaluates her performance *ex post,* and encodes an optimal response to the event ω_t, which is added to the set Ω_t. The second method is through the explicit formation of a contingent plan. The individual can expend effort before the realization of ω_t to add additional events to the set Ω_t. This is assumed to be a costly activity, either because acquiring the behavior requires expensive training, or simply because of the cost associated with adding a large number of contingent plans. The goal is to explicitly model the trade-off between learning by doing and the formation of a contingent plan *ex ante.*

It is assumed that the individual knows that the events in Ω are generated by some unknown stationary distribution $\mu \in \Delta^N$. To simplify the computation of posterior beliefs, let the prior beliefs be given by a Dirichlet distribution, $f(x|\alpha)$, where $\alpha = \{\alpha_1, ..., \alpha_N\}$. The parameter α_i represents the weight associated with the event ω_i.[9] Dirichlet distributions form a set of conjugate priors for the set of measures in Δ^N. If event ω_i is observed in period t, then the updated beliefs of the individual is given by a Dirichlet distribution with parameters $\alpha_j^t = \alpha_j^{t-1}$ if $j \neq i$ and $\alpha_i^t = \alpha_i^{t-1} + 1$, where α_i^{t-1} is the period t belief parameter. The expected value of the probability that event ω_i occurs given a Dirichlet distribution with parameter α is $\frac{\alpha_i}{\sum \alpha_j}$. Initially the individual does not have any information over which event is more likely and therefore let $\alpha^0 = \{b, ...b\}$. The parameter b represents the *intensity* of an individual's beliefs. When b is very large, then the individual believes that the true distribution is close to uniform, and will update beliefs very slowly. In contrast when b is close to zero, then even after one period, the individual believes that the probability of Ω^t is close to one.

The construction of a contingent plan is assumed to be an increasing function of the number of states. Hence planning may be spread over several periods. Specifically between periods t and $t+1$, the individual may decide to make contingent plans for an additional n^t events in $\Omega \backslash \{\Omega^t \cup \{\omega_t\}\}$, (it is assumed that the event that occurs in period t is always added to Ω^t). The cost of this additional contingent planning is $c(n^t)$, where $c(0) = 0$, $c', c'' > 0$. Let the marginal cost of going from $n^t - 1$ to n^t be given by $mc(n^t) = c(n^t) - c(n^t - 1)$. The individual chooses the amount of planning in period t to maximize her discounted expected payoff given her beliefs. The benefit from adding a plan for event ω arises from a net gain of $u_g - u_b$ that accrues only the first time ω_t is observed. This gain is traded off against the probability of that event not occurring in the future.

[9] For $\mu \in \Delta^N$ then:

$$f(\mu|\alpha) = \frac{\Gamma(a_1 + ... + \alpha_N)}{\Gamma(a_1)...\Gamma(\alpha_N)} \mu_1^{\alpha_1 - 1} ... \mu_N^{\alpha_N - 1}.$$

If α_i goes to ∞ while the other parameters remain fixed then this distibution approaches the measure that place probabilty 1 on the $i'th$ event.

Events that are not in $\Omega^t \cup \{\omega_t\}$ have never occurred in the past, and hence the expected probability that such an event occurs in period $t+1$ is $\frac{b}{bN+t}$. Then the probability that the first time an event occurs is in period T, given that it has not occurred before then is:

$$\pi(t,T,b,N) = \left(1 - \frac{b}{bN+t}\right)\left(1 - \frac{b}{bN+t+1}\right)...\left(1 - \frac{b}{bN+T-2}\right)\frac{b}{bN+T-1}.$$

Notice that due to the Bayesian updating, for any event that has not been observed, the probability that such an event will occur decreases with time. Also notice that this probability does not depend upon the events that have occurred since we have assumed that the distribution is an *i.i.d.* process. Using this expression the marginal benefit from adding an event to the set Ω_t that has not been observed in the past is given by:

$$mb(t,b,N) = (u_g - u_b)\sum_{n=0}^{\infty}\pi(t,t+n,b,N)\,\delta^n.$$

The optimal amount of planning in period t, $\tilde{n}^t$, entails adding states to Ω^t, as long as the marginal benefit is greater than the marginal cost and hence solves:

$$mc(\tilde{n}^t + 1) \geq mb(t,b,N) \geq mc(\tilde{n}^t)$$

It is immediate that the amount of planning increases with the net benefit $(u_g - u_b)$. The following proposition follows from this expression.

Proposition 1 *The optimal level of planning, $\tilde{n}^t$, has the following properties:*

1. *$\tilde{n}^t \geq \tilde{n}^{t+1}$, and if $b < \infty$ there is a T such that $\tilde{n}^t = 0$ for $t \geq T$.*

2. *$\tilde{n}^t$ is increasing with b.*

3. *If $N > (u_g - u_b)/mc(1)$, then for b sufficiently close to zero $\tilde{n}^t = 0$ for all t.*

4. *When $(u_g - u_b)/mc(1) > 1$ and $N \geq \left(\frac{(u_g-u_b)}{mc(1)} - \delta\right)/(1-\delta)$ then $\tilde{n}^t = 0$ for all t.*

Proof. Let $c = mc(1) > 0$ be the marginal cost of planning for a single state. Notice that $\lim_{t\to\infty}\pi(t,t+1,b,N) = 0$, and $\pi(t,t+n,b,N) > \pi(t,t+n+1,b,N)$, hence it follows that $\lim_{t\to\infty} mb(t,b,N) = 0$, and hence for large enough c, $mb(t,b,N) < c$, from which statement 1 follows.

For any $T > t$, $\pi(t,T+1,b,N) = \pi(t,T,b,N)\left(\frac{b(N-1)+T-1}{bN+T}\right)$, from which one can show recursively that $\partial\pi(t,T,b,N)/\partial b > 0$ for $T > t$, and hence the marginal benefit is an increasing function of b from which statement 2 follows

As b approaches zero this corresponds to the agent placing almost all probability mass on Ω_t for $t \geq 1$. Thus the only benefit from planning occurs in period 0, before any observations have been made. In that case the marginal benefit is $(u_g - u_b)/N + \varepsilon(b)$, where $\lim_{b\to 0}\varepsilon(b) = 0$, hence if $N > (u_g - u_b)/mc(1)$ the parameter b can be chosen sufficiently small that no planning takes place, which combined with 1 implies 3.

Given that the level of planning is increasing with b, then to obtain a global bound on N, let $b \to \infty$. In that case no learning occurs, and $\pi(t,T,b,N) = (1-1/N)^{T-t-1}/N$, from which one concludes that $mb(t,\infty,N) = (u_g - u_b)/(1-(1-1/N)\delta) > (u_g - u_b)/mc(1)$, when $(u_g - u_b)/mc(1) > 1$. The conclusion of statement 4 is immediate. ■

The first result is that the amount of planning or training decreases with time. This is consistent with our intuition regarding the acquisition of any skill. At the beginning it requires a great deal of thought and practice, but over time individuals are able to perform automatically, though performance quality is always increasing with time. Secondly, increasing b corresponds to decreasing the importance of learning, hence makes planning relatively more beneficial than learning by doing. In other words contingent planning is more important when uncertainty *decreases,* and the environment is more predictable.

The last two results show that when the number of possible events is very large, then it is optimal to learn only by doing. For these conditions to hold one must suppose that the event space is very large, an assumption that in many situations is quite reasonable. For example, suppose that a person make stock trades as a function of the pattern of past price increases and decreases. After k periods there are 2^k such patterns, a number that grows very quickly with k. Suppose that the individual spends a second considering each pattern, then the following table outlines the amount of time needed to explore all patterns possible after k periods:

k	Time to examine all patterns:
10 periods	17 minutes
20 periods	12 days
40 periods	349 centuries

.

This example is typical of many day to day decision problems. An implication of this result, that we exploit in the empirical implementation of the problem, is that the learning rule is *belief independent.* A difficulty with theories of decision making under uncertainty, such as Savage (1972)'s, is not only do they presuppose an unrealistic amount of planning, but they also suppose that an individual makes decisions based upon some unspecified prior distribution, implying a certain amount of indeterminacy. This model illustrates that when uncertainty is sufficiently large, then one can form predictions regarding behavior that depend only upon the observable features of the decision problem.

Simon (1955) and Day (1967) have shown that under the appropriate conditions learning via simple "satisficing" behavior can lead to an optimal choice. Here it is shown that regardless of the level of planning, learning ensures convergence to the optimum:

Proposition 2 *Regardless of the planning rule, behavior converges to the optimum:*

$$\lim_{t\to\infty} \Pr\{\sigma_t(\omega_t) = \sigma^*(\omega_t)\} = 1.$$

Proof. $\Pr\{\sigma_t(\omega_t) = \sigma^*(\omega_t)\} = (1 - \Pr\{\omega_t \notin \Omega_t\}) + \Pr\{\omega_t \notin \Omega_t\}/2$

$> 1 - \Pr\{\omega_t$ has not occured previously given $\mu\} > 1-(1-\varepsilon)^{t-1}$, where $\varepsilon = \min\{\mu(\omega) | \omega \in \text{support } \mu\} > 0$, since Ω is finite. The conclusion of the proposition immediately follows. ■

Finally, notice that this model can explain why skills are expensive to transfer. In this model there is no asymmetric information, rather all that is assumed is the existence of a small cost for encoding behavior for every possible event in Ω. This cost could also be the cost of acquiring the behavior for a particular event from another individual or firm. When the number of possible events are sufficiently large the individual will *choose* not to learn from another individual or firm, but rather the individual chooses to learn by doing. With this interpretation the amount of planning, n^t, corresponds to those responses that are acquired from another person or organization, and hence can be viewed as the level of spillovers from one individual to another. The fact that in some cases n^t is less than the number of possible events implies that the level of spillover is both endogenous and less than 100%.

4 Learning Dynamics

The purpose of this section is to illustrate an empirical implementation of the simplest version of the model. Suppose that it is optimal to use only learning by doing, and that the N events are generated by a uniform, and hence the probability of any given event is $1/N$. As before suppose that for each event, the gain from making a contingent plan is $\Delta = u_g - u_b$. Under these assumptions, a sufficient statistic for the performance of a decision maker at the beginning of date t is given by the number of situations that have been experienced, denoted by $m^t \in \{0, 1, .., N\}$.

Learning can then be modelled as a Markov chain, where the state transition is given by the following rule:

$$m^{t+1} = \begin{cases} m^t, & \text{if } \omega_t \in \Omega_t. \\ m^t + 1 & \text{if not.} \end{cases}$$

The probability of previously unexperienced event occurring is $1 - m_t/N$, hence the probability transition matrix from state i to state j for the Markov process is the following $(N+1) \times (N+1)$ matrix:

$$P_{ij} = \begin{cases} i/N, & \text{if } j = i, i \neq 0, \\ 1 - i/N, & \text{if } j = i + 1, \\ 0, & \text{in all other cases.} \end{cases} \tag{1}$$

If the decision maker is in state m, then the expected payoff is $\hat{u}_m = u_g \cdot m/N + u_b (1 - m/N)$. Let $\hat{U} = \begin{bmatrix} \hat{u}_0 \\ \vdots \\ \hat{u}_N \end{bmatrix}$ be the $N + 1$ dimensional vector giving the expected payoff in each state. If decision maker begins with no experience, then the initial state is $x_0 = \{1, 0, .., 0\}^\top$. The evolution of the expected payoff at time t, is given by the following dynamic process:

$$V_t (u_g, u_b, N) = x_0^\top P^t \hat{U}. \tag{2}$$

The matrix P is upper diagonal, thus the eigenvalues of the matrix are the diagonal elements and are given by $\lambda_i = i/N$, $i = \{0, 1, .., N\}$. And hence we may write formula (2) in terms of powers of the eigenvalues, or:

$$V_t = \sum_{i=1}^{N} \alpha_i e^{t \ln \lambda_i}. \tag{3}$$

This formula describes an individual whose performance increases greatly at the beginning, and approaches U_g in the limit. As Newell and Rosenbloom (1981) observe, this rule approximates a log-linear learning rule for the initial period of learning (see also the work of Alchian (1963) who also identified a log-linear learning rule for production costs). More surprisingly, equation (3) has the same functional form as the learning curve that Newell and Rosenbloom (1981) use for a more precise fit with the data.

Following the lead of Jovanovic and Nyarko (1995) this simple three parameter learning model (parameters: u_g, u_b, N) is estimated using the data sets described in Jovanovic and Nyarko (1995). The purpose of this exercise is mainly description, and not intended to be a careful empirical test of the theory. Thus the parameters are estimated to minimize the squared error between the model and the data:

$$\min_{u_g, u_b, N} \sum_{t=1}^{T} \left(v_t - V_t\left(u_g, u_b, N\right)\right)^2,$$

where v_t is observed performance in period t.

Jovanovic and Nyarko also use a three parameter model, and hence it is reasonable to compare the fit between the two models. We present the results in four of the cases (the ones with a larger number of data points). The results for the other data sets discussed by Jovanovic and Nyarko (1995) are similar. Each data set records the average performance or productivity as a function of the number of trials or experience.

Data Sets:

- Anoplasty: The productivity is the probability of success for a group of doctors having the same level of experience as measured by the number of operations.

- Flight Control Experiment: Number of successful plane landings on a flight simulator as a function of the number of trials. In this, and the subsequent data sets, we follow Jovanovic and Nyarko (1995) and normalize the maximum performance to one.

- Insurance Sales: Relative sales performance averaged over three month periods.

- Steel Plant: Relative monthly productivity for a newly installed steel-finishing production line over the first two years of operation.

The data, the fitted curves, and estimated parameters are illustrated in figures 2-5. It is interesting to note, that only in the anoplasty data set we find that complexity, as measured by N, is very large. The value of N cannot be literally related to the complexity of the task, though high N does imply that performance increases more slowly. In table 1 the R^2 statistics for the two models are reported. In the case of the

Data Set	R^2 Jovanovic-Nyarko Model	R^2 Cognitive Model
Angioplasty	0.31	0.85
Flight Control Experiment	0.93	0.99
Insurance Sales	0.56	0.64
Steel Plant	0.74	0.74

Table 1: R-squared fit for Learning Curves

Jovanovic-Nyarko model, the results for the three parameter model are reported. Their four parameter model not surprisingly fits much better.

Given that neither of these models are literally true, these results do not prove that either model is the "true" model. Rather, it demonstrates that a very parsimonious model based upon learning when an individual faces a large number of possible events does a good job of fitting this data. Even though human cognition is very complex, these results illustrate one way of using the cognitive model to produce an empirically implementable learning model.

5 Pattern Recognition and Learning by Doing

The learning model described in the previous section makes the extreme assumption that all events that have not been experienced yield the same payoff, and once experienced, result in a discrete increase in payoff when responding to the event again in the future. A key ingredient of this model is the use of a finite state space, though it may be more natural to view the environment as a continuum where data is complex and multi-dimensional. One way the brain deals with complex stimuli is to categorizing data, and search for similarities among events. Even if an event is new, if it is very similar to one that has been experienced in the past, then the optimal response is likely to be close to the optimal response derived for a previously experienced similar event.

Beginning with Kahneman and Tversky (1979), there have been a number of contributions showing how the use of similarity judgements can help explain observed decision making. For example Rubinstein (1988), Leland (1998) and Leland (1994) show that the use of similarity judgements can help explain some aspects of observed behavior in the face of risk. Gilboa and Schmeidler (1995) extend this work to provide an axiomatic foundation for the theory of decision based upon similarity judgements using the case based reasoning model from the artificial intelligence literature. These contributions do not study the impact of repeated experience upon behavior guided by pattern recognition, the question that is addressed in this section.

Suppose that the event space is a compact, convex subset of $\Re^d$, for some $d \geq 1$, and suppose that $\{\omega^t\}_{t=1}^{\infty}$ is an *i.i.d.* process represented by a measure μ on $\Re^d$ that is absolutely continuous with Lebesgue measure. Further suppose that the space of decision is given by $Z = \{0,1\}^n$, for some $n \geq 1$, and that an individual's utility function is $U(z|\omega^t)$. The only condition that we place upon preferences is that the optimal choice,

$\sigma^*(\omega^t) = \arg\max_{z\in Z} U(z|\omega^t)$, exists and is Borel measurable.

Let us suppose that an individual acquires behavior only through learning by doing. The problem now is since μ has no mass point, the probability that an individual sees the same event again is zero. Instead, suppose that each period the individual observes ω^t, and then finds the event that occurred in the past that is most similar to ω^t, say $\omega^{t'}$, and then chooses the decision that is associated with the event $\omega^{t'}$. As in the previous section, it is assumed that after a decision is made the individual has sufficient time to contemplate her response, and to associate with the experienced event ω the optimal response $\sigma^*(\omega)$. Hence, the decision z^t at time t is the optimal response to the most similar previous event $\omega^{t'}$, that is $z^t = \sigma^*\left(\omega^{t'}\right)$.

Our concern here is not with the optimal similarity measure (a difficult open question), but rather to illustrate how the use of a very simple similarity measure can nevertheless result in optimal behavior in the long run. Given a history of events and corresponding optimal responses at time t, $H_t = \left\{\left(\omega^1, z^1\right), \left(\omega^2, z^2\right), \ldots, \left(\omega^{t-1}, z^{t-1}\right)\right\}$, and given a new event ω^t, suppose the individual uses the response for the previously experience event that minimizes the Euclidean distance, $\|\omega^\tau - \omega^t\|$, for $\tau = 1, 2, ..., t-1$. In that case the individual's behavior is given by the function:

$$\sigma\left(\omega^t|H^t\right) \in \left\{ \begin{array}{l} z^\tau = \sigma^*(\omega^\tau) \mid \|\omega^\tau - \omega\| \leq \left\|\omega^{\tau'} - \omega\right\|, \\ \forall \tau, \tau' \in \{1, \ldots, t-1\}, (\omega^\tau, z^\tau) \in H^t \end{array} \right\}. \tag{4}$$

Given this rule, we have the following proposition:

Proposition 3 *Suppose that Ω is a compact subset of $\Re^d$, the optimal behavior, $\sigma^*(\omega)$, is Borel measurable, and μ is absolutely continuous with respect to Lesbesgue measure. Then*

$$\lim_{t\to\infty} E\left\{\left\|\sigma^*\left(\omega^t\right) - \sigma\left(\omega^t|H^t\right)\right\|\right\} = 0, \tag{5}$$

where $H^t = \{(\omega_1, \sigma(\omega_1, T)), (\omega_2, \sigma(\omega_2, T)), \ldots, (\omega_{t-1}, \sigma(\omega_{t-1}, T))\}$.

Proof. Let $g_{tn}(\omega) = |\sigma_n^*(\omega^t) - h_n(\omega^t|H^t)|/2$, where n denotes the $n'th$ coordinate. The optimal response is deterministic, and hence there exists a function f such that $E\{\|\sigma^*(\omega^t) - f(\omega^t)\|\} = 0$ (the general theory of pattern recognition allows the optimal choice to be random). The rule used here is called a nearest neighbor classifier, and from Devroye (1981) it follows that

$$\lim_{t\to\infty} \Pr\left\{g_{tn}\left(\omega^t\right) \neq 1\right\} = 0. \tag{6}$$

By the Cauchy Schwarz inequality:

$$E\left\{\left\|\sigma^*\left(\omega^t\right) - \sigma\left(\omega^t|H^t\right)\right\|\right\} \leq \sqrt{E\left\{\sum_{i=1}^N \left(\sigma_i^*(\omega^t) - \sigma_i(\omega^t|H^t)\right)^2\right\}} \tag{7}$$

$$\leq 2\sqrt{\sum_{n=1}^N \Pr\{g_{tn}(\omega^t) \neq 1\}}. \tag{8}$$

from which the result follows. ■

Given that any compact state space can be approximated by a finite partition, this result can be extended to show formally that the discrete model of the previous section can approximate the optimal decision rule for a continuum. Though this result uses a very simple form of pattern recognition, namely the Euclidean distance between two events, it can be extended to include large classes of other similarity measures (see Devroye, Györfi, and Lugosi (1996)). What the result demonstrates is that even if short run behavior is guided by the use of pattern recognition (as is suggested by the framing anomalies identified by Kahneman and Tversky (1979)), long run behavior should still conform to the hypothesis of utility maximization.

This convergence result also distinguishes this paper from the case based decision theory of Gilboa and Schmeidler (1995). They elucidate the logic structure of a cognitive approach by axiomatizing case based decision theory. However, their model does not in general imply the convergence of behavior to utility maximization, and hence does not address the issue of why utility maximization is a good first order approximation to observed behavior.[10] Moreover, while the goal here has been the construction of an empirically implementable model, Gilboa and Schmeidler emphasize that the purpose of their theory is to provide an alternative language for the description of decision making under uncertainty.

6 Incentives and the Probability of Detection

The principal-agent model assumes that both the principal and agent are able to anticipate and plan for the consequences of any incentive system. In particular, for a fixed incentive system, an agent's behavior should be time invariant. However the learning by doing model suggests that in complex environments such behavior may change over time as the individual learns how to better respond to the incentives created by the principal. The purpose of this section is to illustrate an application of the learning by doing to the problem of optimal deterrence.

Becker (1968) has shown that one can arbitrarily lower the cost of ensuring compliance to certain standards of behavior by simultaneously lowering the probability of detection and increasing the fine associated with any observed transgressions of the performance standard. This result explicitly depends upon the hypothesis that individuals understand the system of deterrence, and rationally avoid offending the standard of behavior due to the high expected cost of non-compliance. This result has been quite controversial. First, there is the puzzle of why we do not observe more examples of harsh penalties combined with low probability of detection? Secondly, there is some evidence suggesting that individuals may not carry out a one for one trade-off between increased fines and lower probability of detection. Grogger (1991) has found in a study of California offenders that increasing the probability of detection increases deterrence, but the severity of punishment has a small or insignificant effect.

The learning by doing model provides a solution to both of these questions. In practice, for a wide range of activities, from theft to speeding, we observe regular violation of the accepted standards of behavior. In

[10] See the example on page 624.

the context of learning by doing, the probability of detection controls the number of times an individual offended has to pay a fine or face punishment. The model explicitly assumes that in the absence of experience the individual may not in fact respond optimally to the incentive system, but converges to optimal behavior with experience.

This hypothesis is sensible only if the behavior that is being regulated is complex in the sense discussed above. There is evidence that even simple offenses, such as shoplifting, involve potentially complex skills. Weaver and Carroll (1985) study the strategies used by novice and experienced shoplifters, and find that the experienced shoplifter reacts to a variety of problems that arise, such as the location of security cameras, the location of the goods, the nature of the exit from the store, etc. These are parameters that vary from store to store, and hence skilled shoplifting entails reacting to a large number of possible situations.

Weaver and Carroll also find that novice shoplifters tend to misperceive the risks, and are often motivated by other factors than the physical difficulty of carrying out the act. In any situation where there is probabilistic monitoring of performance, there is a potential for individuals to learn about the effectiveness of the system, and to respond strategically. In this section I consider a reduced form model of these phenomena, consistent with the learning model described above, in which individual behavior is more optimal as the agent acquires more experience with being detected. In other words increasing the probability of detection, increases the number of times that an individual is actually caught at an act, and hence increases the speed of learning.

Before considering the effect of learning by doing on incentives, first consider a simple model illustrating Becker's original result. Suppose that there is a principal who sets the level of detection and fine for agent to maximize social welfare. The agent selects a level of activity $\lambda \in [0,1]$ corresponding to the probability that she obtains a benefit B. The personal cost of this activity is $V(\lambda)$, where $V(0) = 0$, $V' > 0$, $V'' > 0$ and $\lim_{\lambda \to 1} V'(\lambda) = \infty$. In the legal context this activity might correspond to speeding, trespass, theft etc., while in an employment context it might correspond to the level of care taken in an activity, whose intensity may lead to accidents. In the absence of any incentive system the agent selects λ to solve $B = V'\left(\lambda^A\right)$. From the principal's perspective when even the agent receives B, there is a social cost S, and hence the activity entails an expected social cost λS. Therefore the first best activity level satisfies:

$$B - S = V'(\lambda^*).$$

Clearly $\lambda^A > \lambda^*$, and thus the principal wishes the agent to reduce her activity level. This can be achieved with a monitoring system that has a probability of detection $\rho \in [0,1]$ and the level of fine $F \geq 0$ that are imposed whenever the principal observe the agent receiving B. In this case the utility of the agent is given by:

$$U_A = \lambda B - \rho\lambda F - V(\lambda).$$

Let the cost of monitoring be given by $C(\rho)$, where $\rho \in [\bar{\rho}, 1]$, $\bar{\rho} > 0$, and $C' > 0, C'' > 0$, $C(\bar{\rho}) = k > 0$. The fixed k corresponds to setting up a monitoring system with a minimum level of detection $\bar{\rho}$. The net payoff to the principal is $U_P = -\lambda S + \rho\lambda F - C(\rho)$. Let the social welfare before monitoring costs be given by $U(\lambda) = \lambda(B - S) - V(\lambda)$. In the standard principal agent model it is assumed that the agent chooses the optimal λ given the parameters ρ and F. Assuming that the parameters are chosen to maximize social welfare, the principal solves:[11]

$$\begin{aligned} \max_{\rho,F,\lambda} U_A + U_P &= \max_{\rho,F,\lambda} U(\lambda) - C(\rho), \\ \text{subject to } B - \rho F &= -V'(\lambda). \end{aligned}$$

It is straightforward to see that the solution to this problem when k is sufficiently small is:

$$\begin{aligned} 0 &= U'(\lambda^*) = B - S - V'(\lambda^*), \\ \rho^* &= \bar{\rho}, \\ F^* &= S/\bar{\rho}. \end{aligned}$$

In other words the level of care is set to the first best λ^*, while the probability of detection is set at the lowest possible level. In the absence of a bound on ρ, one obtains Becker's result in which the probability of detection is set arbitrarily close to zero, while the fine would be arbitrarily large.

To explore the implications of learning by doing for incentives, consider a simple overlapping generations model with a continuum of agents, normalized to one per generation, that live for two periods. Experience when they are young affects their behavior when old. In practice individuals can either learn to be more cautious because they have been fined, or they may become less cautious because they are not fined even though they may have offended. Thus suppose that there are two groups of individuals.

Each period a fraction γ of the young generation are "socially conscious" (or viewed as the novice shoplifters as in Weaver and Carroll (1985), and hence have no intention to offend) and who begin by choosing the socially optimal rule, $\lambda = \lambda^*$.[12] A fraction $1 - \gamma$ of the agents initially ignore the consequences of their actions upon the probability of a fine, and choose $\lambda = \bar{\lambda}$ to solve $B = V'\left(\bar{\lambda}\right)$, where $\bar{\lambda} > \lambda^*$. The parameter γ is assumed to remain fixed with time, and can be interpreted as representing different background and education levels of the young.

Under the assumption of no discounting, the principal chooses the level of monitoring and the fine to maximize expected welfare subject to the behavior of the agents. With probability λ the socially conscious agents cause an offence that if detected would be fined. If such an event does not occur, then they continue to choose λ when old. Similarly, if they offend and are caught they also choose λ when old. However, if

[11] One could assume that the principal maximizes U_P subject to the agent getting at least $\bar{U}_A$, but given the assumption of transferrable utility, we would get the same answer. The current setup is more convenient, and also the interpretation as a benevolent principal/government setting the optimal deterrence parameters.

[12] It may seem strange to suppose that it is efficient to have a positive probability of shoplifting. However, there are cases where this might occur. For example a toddler may put a candy in the stroller at a store, which the parent does not notice until he or she is at home. If it is a long drive back to the store, many law abiding individuals would decide it is not worthwhile to return to the store.

they offend and are not caught, then they "learn" that they can get away with offending, and choose $\bar{\lambda}$ when old. Recalling that we have defined $U(\lambda) = \lambda(B - S)\rho - V(\lambda)$, then the social return from the socially conscious group is:

$$W_{SC} = U(\lambda^*) + (1-\rho)\lambda^* U(\bar{\lambda}) + (1 - \lambda^* + \rho\lambda^*) U(\lambda^*).$$

In the case of the group that initially ignored the fine in making their choice, suppose that they choose the socially optimal level of the activity if and only if they offend and are caught. Hence the social welfare from this group is:

$$W_O = U(\bar{\lambda}) + \rho\bar{\lambda}U(\lambda^*) + (1 - \rho\bar{\lambda}) U(\bar{\lambda}).$$

Hence the program solved by the principal is given by:

$$\max_{\rho,F,\bar{\lambda},\lambda} \gamma W_{SC} + (1-\gamma) W_O \tag{9}$$

$$\begin{aligned} \text{subject to} \quad : \quad & \\ & B - \rho F = V'(\lambda) \\ & B = V'(\bar{\lambda}) \end{aligned}$$

Notice that the fine F does not appear in the objective function (due to the assumption of transferrable utility), and hence the first incentive constraint is not binding. The second incentive constraint completely determines $\bar{\lambda}$. Hence λ and ρ are chosen to maximize the payoff, while the fine is fixed by the first incentive constraint.[13] The solution to (9) is summarized in the following proposition.

Proposition 4 *The optimal incentive rule with learning is characterized by:*

$$\begin{aligned} U'(\lambda^s) &= \left[\frac{\gamma(1-\rho)}{\gamma(2-(1-\rho)\lambda^*) + (1-\gamma)\rho\bar{\lambda}}\right](U(\lambda^s) - U(\bar{\lambda})), & (10) \\ C'(\rho) &= (\gamma\lambda^s + (1-\gamma)\bar{\lambda})(U(\lambda^s) - U(\bar{\lambda})), & (11) \\ V'(\bar{\lambda}) &= B, & (12) \\ F &= (B - V'(\lambda^s))/\rho. & (13) \end{aligned}$$

This proposition follows from a straightforward computation of the program (9). Notice that at the optimum $U'(\lambda^s) > 0$, and since $U(\lambda)$ is concave it follows that $\lambda^s < \lambda^* < \bar{\lambda}$, and the level of care among those individuals who are responding to incentives is *greater* than in the case of no learning by doing. The reason for this effect is that a higher λ^s increases the chance that a socially conscious person learns the benefits of offending, and hence there is a return to lowering λ^s to lower than the first best under complete information. Secondly, a straightforward comparative statics exercise shows that increasing the fraction of socially conscious individuals, γ, decreases λ^s and increases the level of monitoring ρ.

[13] If one assumes that the government uses funds less efficiently than private agents, then the size of F would enter the objective function. Extending the model to deal with this case is straightforward.

Intuitively, one might expect that increasing the number of individuals who have a propensity to conform would decrease the level of monitoring and the level of expected performance, when in fact the opposite occurs. This follows from the fact that individuals who have a propensity to conform need to be monitored in the event of an offence to reduce the chance that they increase their offence level in the future.

7 Concluding Discussion

There is a large and growing literature documenting the fact that individual behavior is often inconsistent with the utility maximization hypothesis. This paper begins by identifying situations where such behavior is a natural consequence of the complexity of the environment, namely situations where individuals must formulate responses to a large number of possible events, a problem that is central to cognitive models of decision making. The economics literature on bounded rationality has also emphasized the role of complexity in explaining observed behavior, however it has focused for the most part upon complexity of solving a particular problem.

For example, Simon (1978), in his Ely Lecture to the American Economics Association, suggests that bounded rationality arises from the difficulty of making an optimal choice, in other words focuses upon the complexity of the decision space. More recently, Rubinstein (1998)'s review of models of bounded rationality also takes the complexity of the decision as a starting point for understanding boundedly rational choice. In this paper I have explicitly assumed that for a particular situation, the decision problem is relatively simple, and have focused instead upon complexity that arises because an individual cannot formulate responses to the large number of contingent events that she may face. This new perspective has a number of useful economic implications.

The main result of the paper shows that when the event space is sufficiently large, then a Bayesian decision maker may optimally choose to make decisions via learning by doing, rather than through the formulation of contingent plans, as is usually assumed.[14] One implication of this result is that observed behavior is independent of prior beliefs. In Savage (1972)'s model individual behavior can be very sensitive to subjective beliefs over the uncertain states, information that is not likely to be observable, and hence reducing the empirical content of theory. In contrast, this model predicts the use of learning by doing when uncertainty (complexity) is very large, something that is potentially observable, and hence the model may be more amenable to empirical implementation.

In this paper the usefulness of the model is illustrated with three applications: the parameterization of empirical learning curves, explaining why utility maximization, despite many observed anomalies, is a useful first order model, and finally explaining why we do not observe contracts that entail large fines combined with low rates of monitoring. The common thread in each of these examples is that the quality of individual decision making is a function of the complexity of a decision, as measured by the number of contingencies

[14] The observation that cognitive behavior is optimal is not new. For example the psychologists Payne, Bettman, and Johnson (1993) explicitly make this point, though they do not have any formal proof of the observation.

to which an individual must respond, a potentially observable quantity.

References

Alchian, A. (1963, October). Reliability of progress curves in airframe production. *Econometrica 31*(4), 679 – 693.

Arrow, K. (1962, June). The economic implications of learning by doing. *Review of Economic Studies 29*(3), 155–73.

Becker, G. S. (1968, March/April). Crime and punishment: An economic approach. *Journal of Political Economy 76*(2), 169–217.

Binmore, K. (1992). De-bayesing game theory. In B. Skyrms (Ed.), *Studies in Logic and the Foundations of Game Theory: Proceedings of the Ninth International Congress of Logic, Methodology and the Philosophy of Science*, Dordrecht. Kluwer.

Camerer, C. (1995). Individual decision making. In J. H. Kagel and A. H. Roth (Eds.), *The Handbook of Experimental Economics*, Chapter 8, pp. 587–683. Princeton University Press: Princeton University Press.

Churchland, P. S. and T. J. Sejnowski (1993). *The Computational Brain*. Cambridge, MA: MIT Press.

Day, R. H. (1967, May). Profits, learning and the convergence of satistcing to margnialism. *Quarterly Journal of Economics 81*(2), 302–311.

Devroye, L. (1981). On the inequality of cover and hart in nearest neighbor discrimination. *IEEE Transactions on Pattern Analysis and Machine Intelligence 3*, 75–78.

Devroye, L., L. Györfi, and G. Lugosi (1996). *A Probabilistic Theory of Pattern Recognition*. New York, NY: Springer-Verlag.

Gilboa, I. and D. Schmeidler (1995, August). Case-based decision theory. *Quarterly Journal of Economics 110*(3), 605–640.

Grether, D. M. (1980, November). Bayes rule as a description model: The representativeness heuristic. *Quarterly Journal of Economics 95*(3), 537–557.

Grogger, J. (1991). Certainty vs. severity of punishment. *Economic Inquiry 29*, 297–310.

Harless, D. W. and C. F. Camerer (1994, November). The predictive utility of generalized expected utility theories. *Econometrica 62*(6), 1251–1289.

Hey, J. D. and C. Orme (1994, November). Investigating generalizations of expected utility theory using experimental data. *Econometrica 62*(6), 1291–1326.

Jovanovic, B. and Y. Nyarko (1995). A bayesian learning model fitted to a variety of empirical learning curves. *Brookings Papers on Economic Activity* (1), 247–305.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decisions under risk. *Econometrica 47*, 262–91.

Knight, F. H. (1921). *Risk, Uncertainty, and Profit.* New York, NY: Hart, Schaffner, and Marx.

Leland, J. W. (1994). Generalized similarity judgements: An alternative explanation for choice anomalies. *Journal of Risk and Uncertainty 9*, 151–172.

Leland, J. W. (1998, May). Similarity judgements in choice under uncertainty. *Management Science 44*(5), 659–672.

Lippman, S. A. and R. P. Rumelt (1982). Uncertain imitability: an analysis of interfirm differences in efficiency under competition. *Bell Journal of Economics 13*, 418–438.

Newell, A. (1990). *Unified Theories of Cognition.* Cambridge, Mass.: Harvard University Press.

Newell, A. and P. S. Rosenbloom (1981). Machanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive Skills and Their Acquisition*, Hillsdale, New Jersey, pp. 1–55. Lawrence Erlbaum Associates, Publishers.

Nilsson, N. J. (1998). *Artificial Intelligence: A New Synthesis.* San Francisco, CA.: Morgan Kauffman Publishers.

Payne, J. W., J. R. Bettman, and E. J. Johnson (1993). *The Adaptive Decision Maker.* Cambridge, U.K.: Cambridge University Press.

Rosen, S. (1972, August). Learning by experience as a joint production. *Quarterly Journal of Economics 86*, 366 – 82.

Rubinstein, A. (1988). Similarity and Decision-Making under risk (is there a utility theory resoltion to the allais paradox?). *Journal of Economic Theory 46*, 145–153.

Rubinstein, A. (1998). *Modeling Bounded Rationality.* Cambridge, MA: MIT Press.

Savage, L. J. (1972). *The Foundations of Statistics.* New York, N.Y.: Dover Publications.

Simon, H. A. (1951, July). A formal theory of the employment relationship. *Econometrica 19*, 293–305.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics 69*, 99–118.

Simon, H. A. (1978). Rationality as a process and as a product of thought. *American Economic Review 68*(2), 1–16.

Simon, H. A. and J. Schaeffer (1992). The game of chess. In R. J. Auman and S. Hart (Eds.), *Handbook of Game Theory*, Volume 1 of *Handbooks in Economics*, pp. 1–18. Amsterdam, The Netherlands: North-Holland.

Weaver, F. and J. S. Carroll (1985, December). Crime perceptions in a natural setting by expert and novice shoplifters. *Social Psychology Quarterly 48*(4), 349–359.

Williamson, O. E. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: The Free Press.

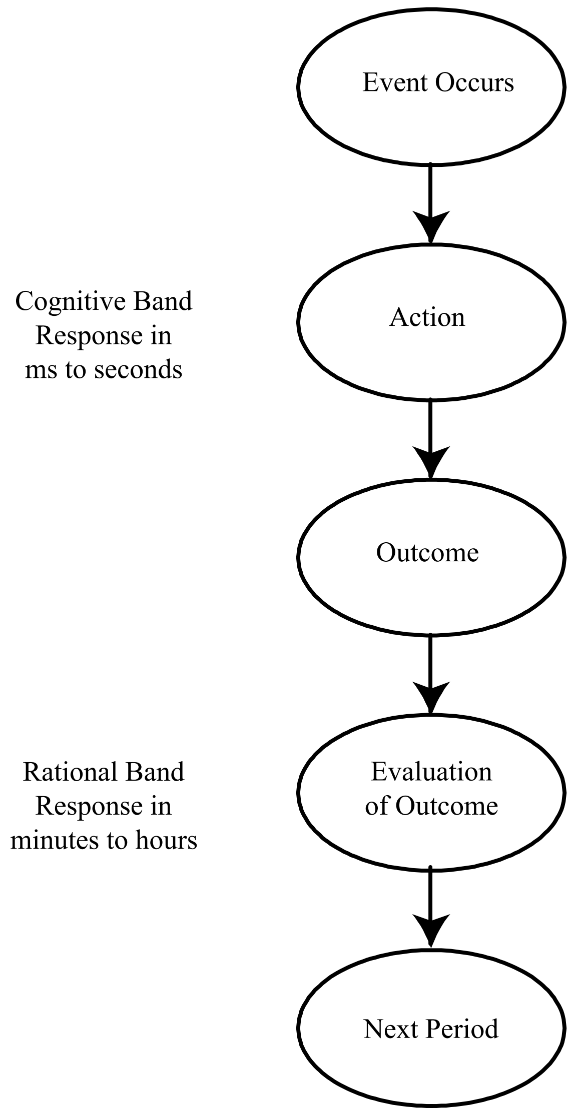
Event Occurs
Action
Outcome
Evaluation of Outcome
Next Period
Cognitive Band Response in ms to seconds
Rational Band Response in minutes to hours

Figure 1:

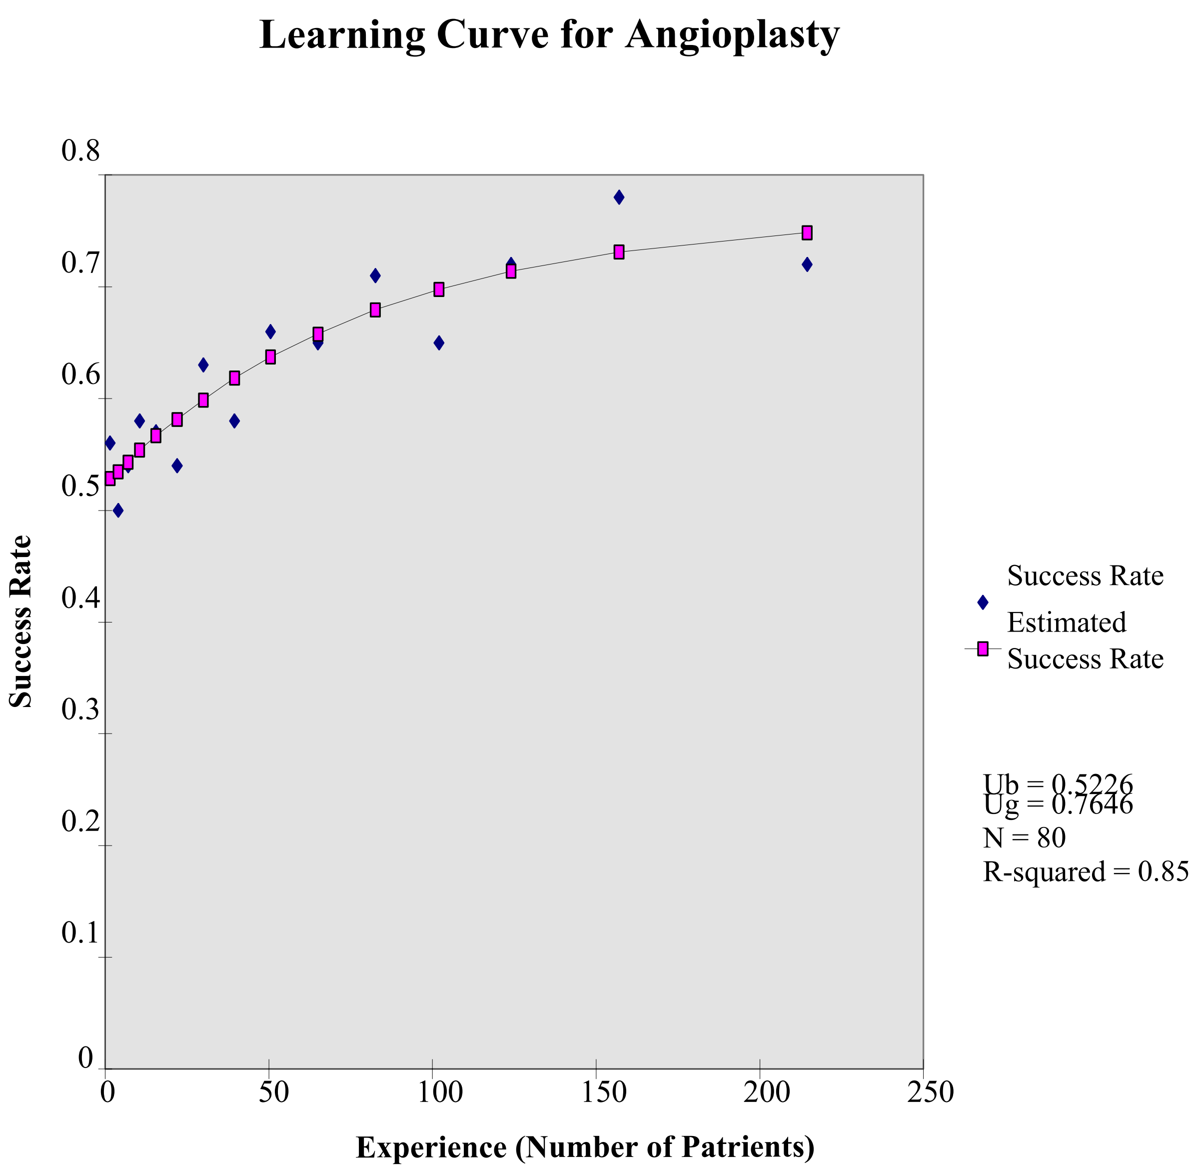
Learning Curve for Angioplasty
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Success Rate
0
50
100
150
200
250
Experience (Number of Patrients)
Success Rate
Estimated Success Rate
Ub = 0.5226
Ug = 0.7646
N = 80
R-squared = 0.85

Figure 2:

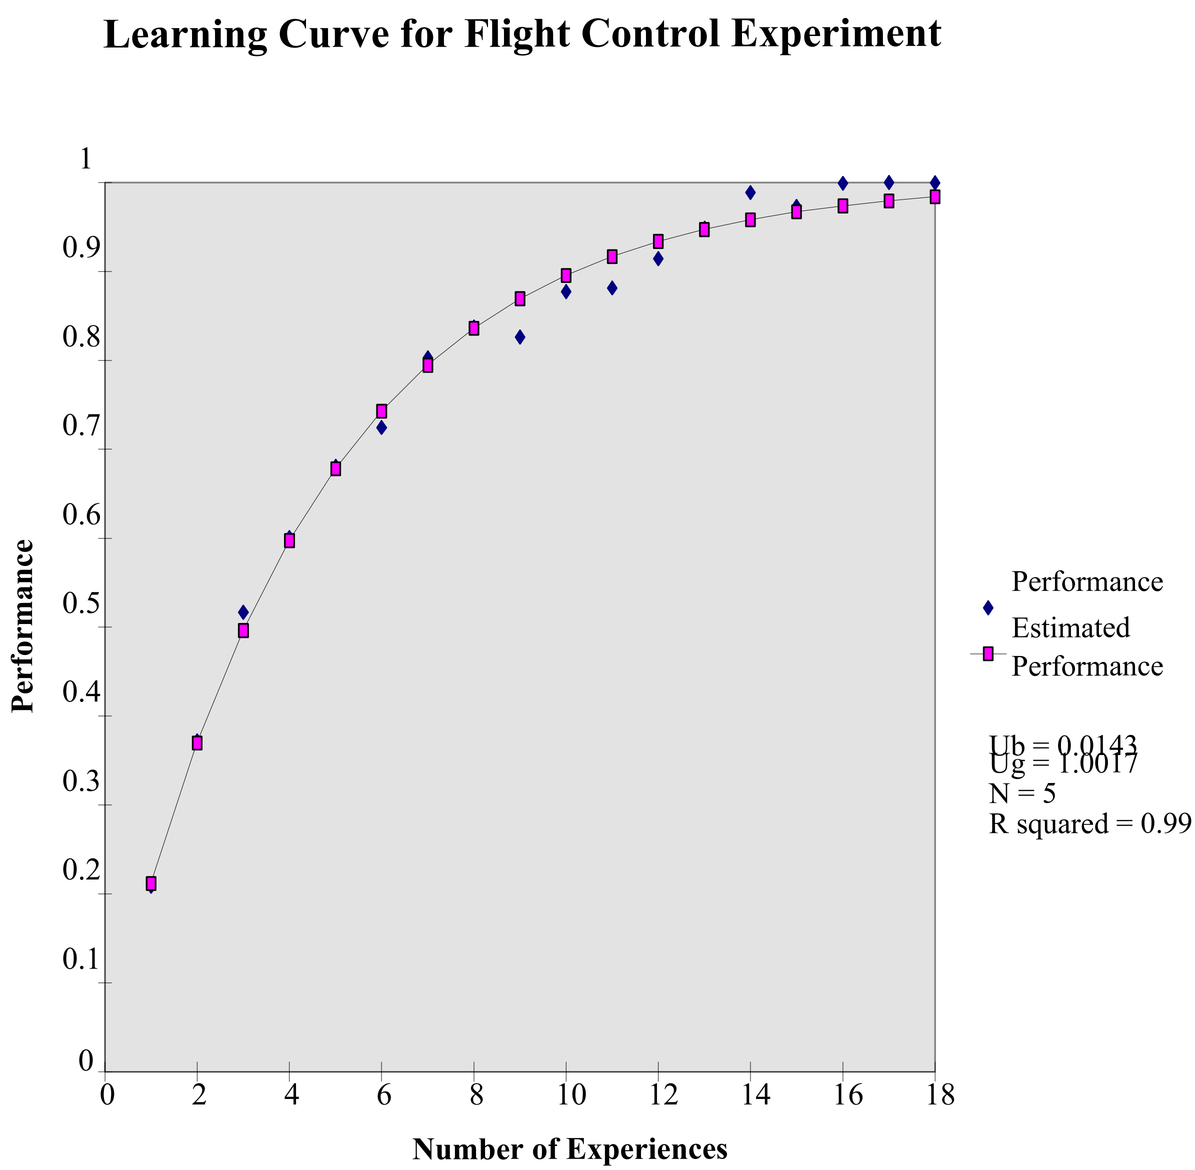
Learning Curve for Flight Control Experiment
Performance
Number of Experiences
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
2
4
6
8
10
12
14
16
18
Performance
Estimated Performance
Ub = 0.0143
Ug = 1.0017
N = 5
R squared = 0.99

Figure 3:

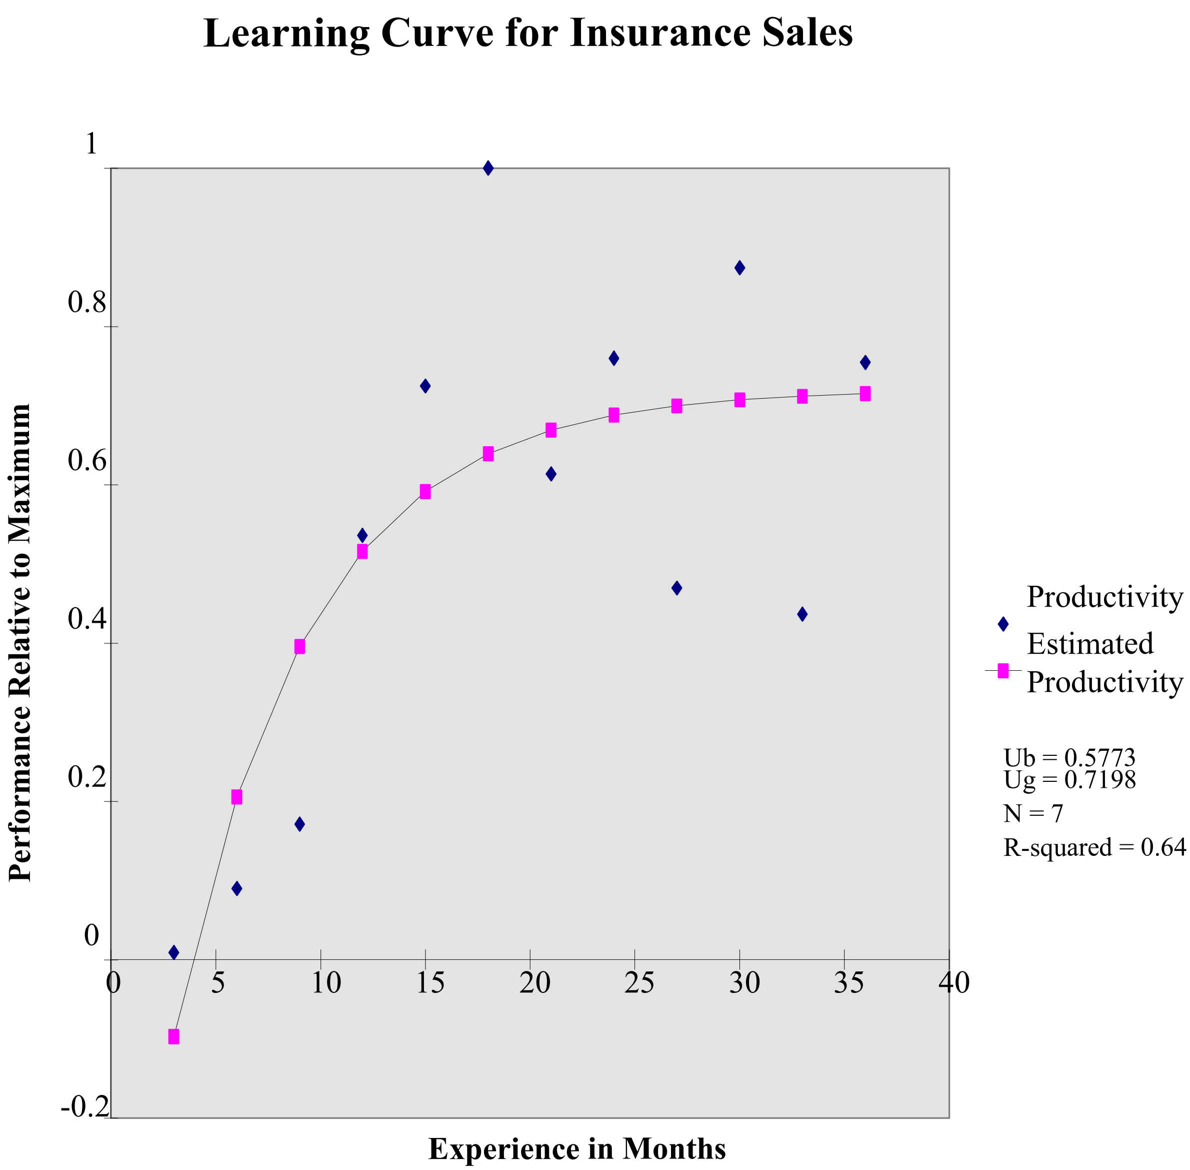
Learning Curve for Insurance Sales
Performance Relative to Maximum
Experience in Months
1
0.8
0.6
0.4
0.2
0
-0.2
0
5
10
15
20
25
30
35
40
Productivity
Estimated Productivity
Ub = 0.5773
Ug = 0.7198
N = 7
R-squared = 0.64

Figure 4:

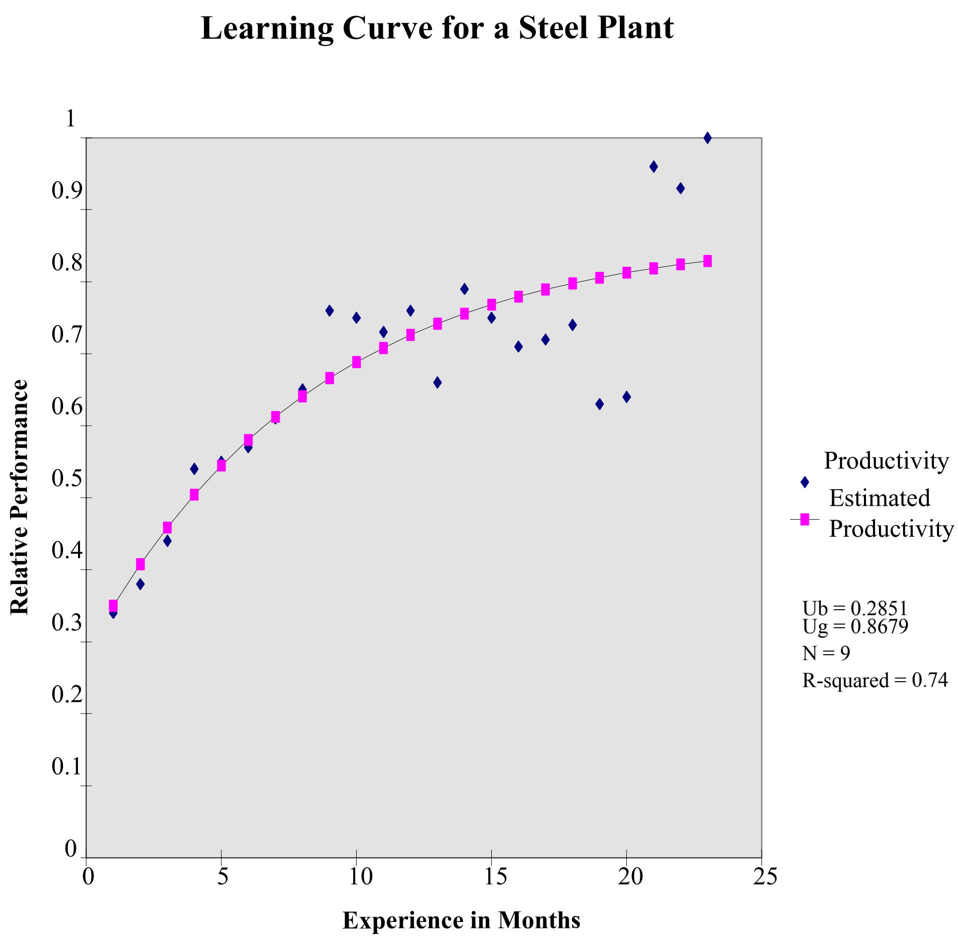
Learning Curve for a Steel Plant
Relative Performance
Experience in Months
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
5
10
15
20
25
Productivity
Estimated Productivity
Ub = 0.2851
Ug = 0.8679
N = 9
R-squared = 0.74

Figure 5:

USC LAW SCHOOL WORKING PAPER SERIES

2000

00-1 W. Bentley MacLeod, *Complexity and Contract* (forthcoming, REVUE D'ECONOMIE INDUSTRIELLE, April 2000).

00-2 Jennifer Arlen, Matthew Spitzer and Eric Talley, Endowment Effects, Other-Regarding Preferences, and Corporate Law.

00-3 Christopher Slobogin, *An End to Insanity: Recasting the Role of Mental Disability in Criminal Cases* (forthcoming in the VIRGINIA LAW REVIEW).

00-4 Christopher Slobogin, *Mental Illness and the Death Penalty*, 1 CALIFORNIA CRIMINAL LAW REVIEW 3 (2000).

00-5 R. Michael Alvarez and Edward J. McCaffery, "Is There a Gender Gap in Fiscal Political Preferences?"

00-6 Ariela Gross, "Between 'Race' and 'Nation': Indian/Black Identity in the Southern Courtroom, 1780-1840."

00-7 Peter H. Huang, *Reasons within Passions: Emotions and Intentions in Property Rights Bargaining*, (forthcoming in 79 OREGON LAW REVIEW, November 2000).

00-8 Marcel Kahan and Ehud Kamar, "Price Discrimination in the Market for Corporate Law."

00-9 Dan Klerman and Nuno Garoupa, "Optimal Law Enforcement with a Rent-Seeking Government."

00-10 Dan Klerman, "The Selection of Thirteenth-Century Criminal Disputes for Litigation."

00-11 W. Bentley MacLeod, "Cognition and the Theory of Learning by Doing."

00-12 W. Bentley MacLeod and Mark Pingle, "An Experiment on the Relative Effects of Ability, Temperament and Luck on Search with Uncertainty."

00-13 Daniel Klerman, "Female Prosecutors in Thirteenth-Century England."

1999

99-1 Howard F. Chang and Hilary Sigman, *Incentives to Settle Under Joint and Several Liability: An Empirical Analysis of Superfund Litigation* (forthcoming in 29 JOURNAL OF LEGAL STUDIES (January 2000)).

99-2 Mary L. Dudziak, The March on Washington, At Home and Abroad.

99-3 Pamela M. Kato, Thomas D. Lyon, John Flavell, Raquel S. Klibanoff, Robin Higashi, and Lynne C. Huffman, Preschoolers' Moral Judgments About Illness and Treatment: Who's Bad?

99-4 Thomas D. Lyon, *Child Witnesses and the Oath: Empirical Evidence* (forthcoming in SO. CAL. LAW REV. (2000)).

99-5 Paul Mahoney and Mark Weinstein, The Appraisal Remedy and Merger Premiums.

99-6 Eric Talley, *Precedential Cascades: An Appraisal*, 73 SO. CAL. L. REV. 87 (1999).

99-7 Eric Talley, *Taking the "I" Out of "Team": Intra-Firm Monitoring and the Content of Fiduciary Duties*, 24 J. CORP. LAW 1001 (1999).

99-8 Antonio Bernardo, Eric Talley, and Ivo Welch, *A Theory of Legal Presumptions*, 16 J. LAW, ECON. & ORG. 1 (2000).

99-9 Antonio Bernardo and Eric Talley, A Note on Presumptions with Sequential Litigation.

99-10 W. Bentley MacLeod and Daniel Parent, Job Characteristics and the Form of Compensation.

99-11 R. Michael Alvarez and Edward J. McCaffery, *Gender and Tax* (forthcoming in GENDER AND POLITICS, Jyl Josephson and Susan Tolleson-Rinehart, eds.).

99-12 Daniel Klerman, *Settlement and the Decline of Private Prosecution in Thirteenth-Century England* (forthcoming in 19 LAW & HISTORY REVIEW (2001)).

99-13 George Lefcoe, Mortgage Prepayment by Defeasance.

99-14 Edward J. McCaffery, *The Burdens of Benefits* (forthcoming in the VILLANOVA LAW REVIEW).

99-15 Edward J. McCaffery, A Life Estate Conception of Property.

99-16 Edward J. McCaffery, *The Missing Links in Tax Reform* (forthcoming in the CHAPMAN LAW REVIEW).

99-17 Alexander M. Capron, *Genetics and Insurance: Accessing and Using Private Information*, 17 SOCIAL PHILOSOPHY AND POLICY FOUNDATION 235 (Summer 2000).

99-18 Alexander Morgan Capron, *Social Science, Bioethics, and the Law: What Contributions to Policy Development?* 128 DAEDALUS (JOURNAL OF THE AMERICAN ACADEMY OF ARTS AND SCIENCES) 295 (Fall 1999).

99-19 Lorne Carmichael and W. Bentley MacLeod, Caring About Sunk Costs: A Behavioral Solution to Hold-up Problems with Small Stakes.

99-20 David B. Cruz, *"The Sexual Freedom Cases"? Contraception, Abortion, Abstinence, and the Constitution*. (A revised version of this paper is forthcoming in 35 HARVARD CIVIL RIGHTS-CIVIL LIBERTIES LAW REVIEW, No. 2, Summer 2000).

99-21 Gregory C. Keating, Fairness and Two Fundamental Questions in the Tort Law of Accidents.

99-22 Kevin Davis, Self-Interest and Altruism in the Deterrence of Transnational Bribery.

99-23 Thomas D. Lyon, *Expert Testimony on the Suggestibility of Children: Does it Fit?* (forthcoming in B.L. Bottoms, M.B. Kovera & B.D. McAuliff (eds.), CHILDREN AND THE LAW: SOCIAL SCIENCE AND POLICY (New York: Cambridge)).

99-24 Thomas D. Lyon, *Questioning Children: The Effects of Suggestive and Repeated Questioning* (forthcoming in J. Conte (ed.)., SUGGESTIBILITY OF CHILDREN AND ADULTS (Thousand Oaks, CA: Sage)).

99-25 Thomas D. Lyon, Karen J. Saywitz, Debra L. Kaplan, and Joyce S. Dorado, *Reducing Maltreated Children's Reluctance to Answer Hypothetical Oath-Taking Competence Questions*. (Revised version of Olin Working Paper No. 98-9). (Forthcoming in LAW & HUMAN BEHAVIOR).

1998

98-1 Cindy Alexander, Jennifer Arlen and Mark Cohen, *Regulating Corporate Criminal Sanctions: Evidence on the Effect of the U.S. Sentencing Guidelines,* 42 JOURNAL OF LAW & ECONOMICS 271 (1999).

98-2 Jennifer Arlen, *Economic Analysis of Tort Damages: A Survey* (forthcoming in the ENCYCLOPEDIA OF LAW AND ECONOMICS).

98-3 Howard F. Chang, *Migration as International Trade: The Economic Gains from the Liberalized Movement of Labor*, 3 UCLA JOURNAL OF INTERNATIONAL LAW AND FOREIGN AFFAIRS 371 (Fall/Winter 1998-99).

98-4 Mary L. Dudziak, Birmingham, *Addis Ababa and the Image of America: Managing the Impact of Foreign Affairs on Civil Rights in the Kennedy Administration* (forthcoming in Brenda Gayle Plummer, ed., AMERICA'S DILEMMA, University of North Carolina Press).

98-5 Ariela J. Gross, *Litigating Whiteness: Trials of Racial Determination in the Nineteenth Century South* (forthcoming in 108 YALE LAW JOURNAL (1998)).

98-6 James R. Hackney, Jr., Law and Neoclassical Economics Theory: Structure, Development, Dissolution and Resolution(?).

98-7 Daniel Klerman, Private Prosecution of Crime in Thirteenth-Century England: The Importance of Settlement and Female Prosecutors.

98-8 Michael S. Knoll, Tax Planning, Effective Tax Rates and the Structure of the Income Tax.

98-9 Thomas D. Lyon and Karen J. Saywitz, *Young Maltreated Children's Competence to Take the Oath*, 3 APPLIED DEVELOPMENTAL SCIENCE 16 (1999).

98-10 Thomas D. Lyon, *The New Wave in Children's Suggestibility Research: A Critique*, 84 CORNELL LAW REVIEW 1004 (1999).

98-11 Eric L. Talley, A Model of Hierarchical Judicial Auditing.

98-12 Charles D. Weisselberg, *Saving Miranda* (forthcoming in 84 CORNELL L. REV. (1998)).

98-13 Scott Altman, A Theory of Child Support.

98-14 Jennifer Arlen, *Comment: The Future of Behavioral Economic Analysis of Law* (forthcoming in VANDERBILT LAW REVIEW, Symposium issue (1998)).

98-15 Susan Athey, Kyle Bagwell, and Chris William Sanchirico, Collusion and Price Rigidity.

98-16 Michael S. Knoll, Hedging in an Economy with Asymmetric Taxes: A Comment on Moshe Ayre Milevsky & Eliezer Z. Prisman, *Hedging and Pricing with Tax Uncertainty: Managing Under an Arkansas Best Doctrine* (forthcoming in the ANNUAL PROCEEDINGS OF THE CHICAGO BOARD OF TRADE).

98-17 Nancy S. Marder, *The Myth of the Nullifying Jury* (forthcoming in 93 NORTHWESTERN LAW REVIEW).

98-18 Alexander S. P. Pfaff and Chris William Sanchirico, Environmental Self-Auditing: Setting the Proper Incentives for Discovering and Correcting Environmental Harm.

98-19 Chris William Sanchirico, Enforcement by Hearing: An Integrated Model of Evidence Production.

98-20 Chris William Sanchirico, Games, Information and Evidence Production: With Application to Legal History and "Decoupling."

98-21 Chris William Sanchirico, Taxes Versus Legal Rules as Instruments for Equity: A More Equitable View.

98-22 Matt Spitzer and Eric Talley, *Judicial Auditing*, 29 J. LEGAL STUD. 649 (2000).

1997

97-1 Scott Altman, *Should Child Custody Rules be Fair?*, 35 U. OF LOUISVILLE J. OF FAMILY LAW 325 (1997).

97-2 Jennifer Arlen, *Corporate Crime and Its Control*, THE NEW PALGRAVE DICTIONARY OF ECONOMICS AND THE LAW, Peter Newman (ed.), Vol. 1, p. 492 (1998)

97-3 Michael S. Knoll, *Products Liability and Legal Leverage: The Perverse Effects of Stiff Penalties*, 45 UCLA LAW REVIEW 99 (1997).

97-4 Paul G. Mahoney, *Contract Remedies: A Survey and Critique* (forthcoming in the ENCYCLOPEDIA OF LAW AND ECONOMICS).

97-5 Paul G. Mahoney, The Stock Pools and the Securities Exchange Act.

97-6 Jennifer Arlen and Reinier Kraakman, *Controlling Corporate Misconduct: An Analysis of Corporate Liability Regimes*, 72 NYU LAW REVIEW 687 (1997).

97-7 Lucian Arye Bebchuk and Howard F. Chang, *The Effect of Offer-of-Settlement Rules on the Terms of Settlement* (forthcoming in 28 JOURNAL OF LEGAL STUDIES 489 (June 1999)).

97-8 Linda S. Beres and Thomas D. Griffith, *Do "Three Strikes" Laws Make Sense? Habitual Offender Statutes and Criminal Incapacitation* (forthcoming in the GEORGETOWN LAW JOURNAL).

97-9 Thomas D. Lyon, *Scientific Support for Expert Testimony on Child Sexual Abuse Accommodation Syndrome* (forthcoming in Jon Conte, ed., THE KNOWNS AND UNKNOWNS OF CHILD SEXUAL ABUSE).

97-10 Judith Resnik, *Changing Practices, Changing Rules: Judicial and Congressional Rulemaking and Civil Juries, Civil Justice, and Civil Judging,* ALABAMA LAW REVIEW (forthcoming 1998).

97-11 Matthew L. Spitzer, *The Economics of Freedom of Expression* (forthcoming in the NEW PALGRAVE DICTIONARY OF ECONOMICS AND THE LAW).

97-12 Christopher D. Stone, *Too Many Fishing Boats, Too Few Fish: Can Trade Laws Trim Subsidies and Restore the Balance in Global Fisheries?* 24 ECOLOGY LAW QUARTERLY 505 (1997).

97-13 Eric Talley, *Interdisciplinary Gap-Filling: A Review of Game Theory and the Law*, J. OF LAW AND SOCIAL INQUIRY (forthcoming 1997).

97-14 H. Lorne Carmichael and W. Bentley MacLeod, Fair Territory: Preferences, Bargaining, and the Endowment Effect.

97-15 Ronald A. Cass, Richard D. Boltuck, Seth T. Kaplan and Michael S. Knoll, Antidumping.

97-16 Mary L. Dudziak, *The Little Rock Crisis and Foreign Affairs: Race, Resistance and the Image of American Democracy*, 70 SOUTHERN CALIFORNIA LAW REVIEW 1641 (1997).

97-17 Michael S. Knoll, Financial Innovation, Tax Arbitrage, and Retrospective Taxation: The Problem with Passive Government Lending.

97-18 Michael S. Knoll, The Principle of Marginal Deterrence in Torts: The Potentially Perverse Effects of Stiffer Tort Penalties and Higher Taxes.

97-19 George Lefcoe, How Buyers and Sellers of Development Land Deal with Regulatory Risk.

97-20 Thomas D. Lyon, *Are Battered Women Bad Mothers? Rethinking the Termination of Abused Women's Parental Rights for Failure to Protect,*in H. Dubowitz, ed., NEGLECTED CHILDREN: RESEARCH, PRACTICE AND POLICY (Thousand Oaks: Sage Publications, 1999, Chapter 12, pp. 237-260).

97-21 Thomas D. Lyon and Jonathan J. Koehler, *Where Researchers Fear to Tread: Interpretive Differences Among Testifying Experts in Child Sexual Abuse Cases*, in S. Ceci and H. Hembrooke, eds., WHAT CAN (AND SHOULD) BE SAID IN COURT: EXPERT WITNESSES IN CHILD ABUSE CASES (American Psychological Association: Washington, DC, 1998, Chapter 13, pp. 249-263).

97-22 W. Bentley MacLeod, Complexity, Contract and the Employment Relationship.

97-23 W. Bentley MacLeod and James M. Malcomson, Motivation and Markets.

97-24 Eric L. Talley, *Turning Servile Opportunities to Gold: A Strategic Analysis of the Corporate Opportunities Doctrine* (forthcoming in 108 YALE LAW JOURNAL (1998)).

1996

96-1 Antonio E. Bernardo and Eric L. Talley, *Investment Policy and Exit-Exchange Offers Within Financially Distressed Firms*, 51 J. FINANCE 871 (1996).

96-2 Howard F. Chang, *Carrots, Sticks, and International Externalities*, 17 INT'L REV. OF LAW & ECON. 309 (1997).

96-3 Howard F. Chang, *Liberalized Immigration as Free Trade: Economic Welfare and the Optimal Immigration Policy*, 145 U. PA. L. REV. 1147 (1997).

96-4 Erwin Chemerinsky & Catherine Fisk, *The Filibuster*, 49 STANFORD L. REV. 181 (1997).

96-5 Susanne Lohmann and Hugo Hopenhayn, *Delegation and the Regulation of Risk* (forthcoming in GAMES AND ECONOMIC BEHAVIOR).

96-6 Michael S. Knoll, *Products Liability and Legal Leverage: The Perverse Effects of Stiff Penalties*, 45 UCLA LAW REVIEW 99 (1997).

96-7 Susanne Lohmann, *Demosclerosis, or Special Interests 'R' Us: An Informational Rationale for Political Gridlock*, in Michelle R. Garfinkel and Stergios Skaperdas, eds., THE POLITICAL ECONOMY OF CONFLICT AND APPROPRIATION (Cambridge: Cambridge University Press, 1996, Chapter 7, pp. 119-130).

96-8 Susanne Lohmann, Electoral Incentives, Informational Asymmetries, and the Policy Bias Toward Special Interests.

96-9 Susanne Lohmann, *Federalism and Central Bank Autonomy: The Politics of German Monetary Policy, 1957-1992* (forthcoming in 51 WORLD POLITICS, April 1998).

96-10 Thomas D. Lyon & Jonathan J. Koehler, *The Relevance Ratio: Evaluating the Probative Value of Expert Testimony in Child Sexual Abuse Cases*, 82 CORNELL L. REV. 43 (1996).

1995

95-1 Ron Garet, *Deposing Finnis*, 4 SO. CAL. INTERDISCIPLINARY L. J. 605 (1995).

95-2 Michael S. Knoll, *An Accretion Corporate Income Tax*, 49 STANFORD L. REV. 1 (1996).

95-3 Michael S. Knoll, *A Primer on Prejudgment Interest*, 75 TEXAS L. REV. 293 (1996).

95-4 Mark Weinstein, *Profit Sharing Contracts in Hollywood: Evolution and Analysis*, 27 J. OF LEGAL STUDIES 23 (1998).

1994

94-1 Scott Altman, *Lurking in the Shadow*, 68 SO. CAL. L. REV. 493 (1995).

94-2 Jennifer Arlen and Deborah Weiss, *A Political Theory of Corporate Taxation*, 105 YALE L.J. 325 (1995).

94-3 Lucian Arye Bebchuk and Howard F. Chang, *An Analysis of Fee-Shifting Based on the Margin of Victory: On Frivolous Suits, Meritorious Suits, and the Role of Rule 11*, 25 J. OF LEGAL STUDIES 371 (1996).

94-4 Howard F. Chang, *An Economic Analysis of Trade Measures to Protect the Global Environment*, 83 GEORGETOWN LAW J. 2131 (1995).

94-5 Howard F. Chang, *Patent Scope, Antitrust Policy, and Cumulative Innovation*, 26 RAND J. OF ECON. 34 (1995).

94-6 Erwin Chemerinsky, *The Fifty-Fifth Cleveland-Marshall Fund Lecture: When the Government Must Make Content-Based Choices*, 42 CLEVE. ST. L. REV. 199 (1994).

94-7 Erwin Chemerinsky, *The Values of Federalism*, FLORIDA L. REV. (forthcoming).

94-8 Jeffrey A. Dubin and Matthew L. Spitzer, *Testing Minority Preferences in Broadcasting*, 68 SO. CAL. L. REV. 841 (1995).

94-9 Ronald R. Garet, *Gnostic Due Process*, 7 YALE J. LAW & HUMANITIES 97 (1995).

94-10 Gregory Keating, *Reasonableness and Rationality in Tort Theory*, 48 STANFORD L.REV. 501 (1996).

94-11 Michael S. Knoll, *A Primer on Prejudgment Interest*, 75 TEXAS L. REV. 293 (1996).

94-12 Michael S. Knoll, Put-Call Parity and the Development of the Modern Mortgage.

94-13 Michael S. Knoll, Socially Responsible Investment and Modern Financial Markets.

94-14 Elyn R. Saks, Interpreting Interpretation: The Limits of Hermeneutic Psychoanalysis (forthcoming in her book on Hermeneutics, Yale University Press, 1997).

94-15 Pablo Spiller and Matthew L. Spitzer, *Where is the Sin in Sincere? Sophisticated Manipulation of Sincere Judicial Voters (with applications to other voting environments)*, 11 J. LAW, ECON. & ORGAN. 32 (1995).

94-16 Christopher D. Stone, *What to Do About Biodiversity: Property Rights, Public Goods and the Earth's Biological Riches*, 68 SO. CAL. L. REV. 577 (1995).

94-17 Charles D. Weisselberg, *The Exclusion and Detention of Aliens: Lessons From the Lives of Ellen Knauff and Ignatz Mezei*, 143 UNIV. PENN. L. REV. 933 (1995).