The Boundaries of Instance-Based Learning Theory for Explaining Decisions from Experience

Cleotilde Gonzalez

Dynamic Decision Making Laboratory

Social and Decision Sciences Department

Carnegie Mellon University

coty@cmu.edu

208 Porter Hall, Pittsburgh, PA 15213

Phone: 412-268-6242

Fax: 412-268-6938

Acknowledgements: This is a summary of research from many years, and thus I owe recognition to a large number of organizations that supported my work, including the Army Research Laboratory, the Office of Naval Research, the National Science Foundation, and the Army Research Office, among others. I also owe recognition to doctoral students, post-doctoral fellows and other research staff in the dynamic decision making laboratory.

from rational behavior have been identified, replicated, and extended upon using laboratory experiments, to the point where this type of research has dominated the field of behavioral decision making for the past six decades.

However, despite the many years of effort, we have only limited answers to the question of *how* people make decisions; rather, most research has aimed at demonstrating how people don't make decisions. The large collection of cognitive biases cannot be all explained by one comprehensive theory and most importantly, we do not know how the biases develop and how do they emerge in the first place. As a result, we know little of how to prevent them. Most empirical studies up to date focus on the observable processes such as choice selection, and ignore cognitive processes that lead to choice, such as recognizing alternatives, deciding when to search for information, evaluating and integrating possible outcomes, and learning from good and bad decisions, among other processes.

A recent development in decision sciences has great potential to expand our understanding and provide insights into the decision making process. A shift of attention to how decisions are made from experience (i.e., *decisions from experience*), rather than from explicit description of options, opens a window towards a better understanding of cognitive processes that including: information search, recognition and similarity processes, integration and accumulation of information, feedback, and learning. Researchers use experimental paradigms that involve repeated decisions rather than one-shot decisions, the estimation of possible outcomes and probabilities based on the observed outcomes rather than from a written description, and learning from feedback. All of which are natural processes for making decisions in many real-world situations in which alternatives, outcomes, and probabilities are unknown. The experimental paradigm often involves two alternatives, represented by two unlabeled buttons, each representing a probability distribution of outcomes that is unknown to participants. Clicking a button yields an outcome as a result of a random draw from

Although there are multiple paradigms for the study of decisions from experience (Hertwig & Erev, 2009; Gonzalez & Dutt, 2011), a common paradigm is the "sampling" paradigm (see Figure 1), in which people are able to explore the outcomes of the options without real consequences before they decide to make a final choice.

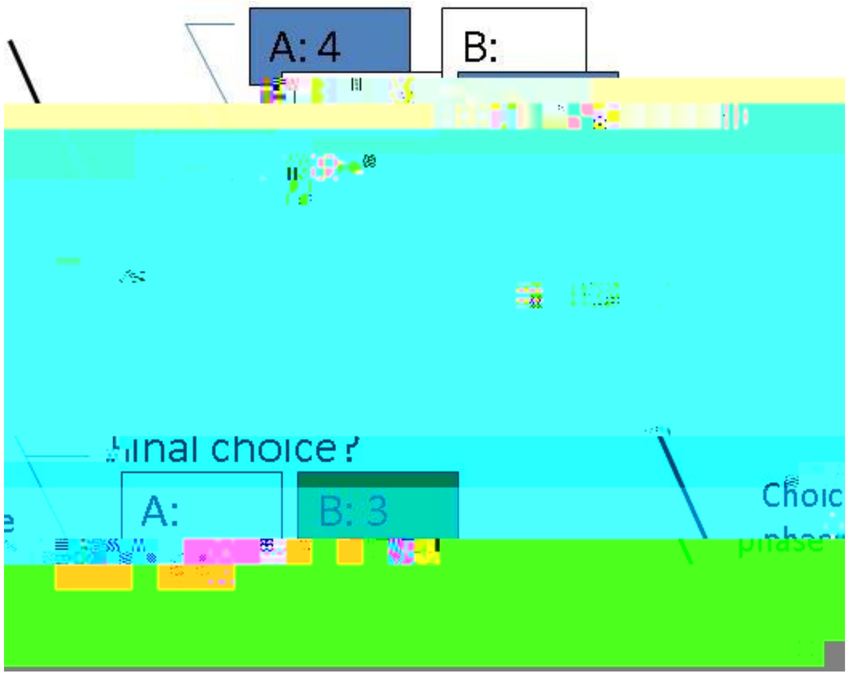

Figure 1. The sampling paradigm of decisions from experience.

A key observation that contributed to the initial success of the theoretical development of decisions from experience was the "description-experience gap" (Hertwig, Barron, Weber, & Erev, 2004): that the choice that an individual makes depends on how information about the problem is acquired (from description or experience); particularly in problems involving outcomes with low probabilities

(probabilities less than .2, "rare events"). A robust finding across a range of paradigms for decisions from experience is that people behave as if rare events have *less* impact than they deserve according to their objective probabilities. More importantly, this finding contradicts the prediction from prospect theory that people behave as if rare events have *more* impact than they deserve. However, this theory only applies to "simple prospects with monetary and stated probabilities" (Kahneman & Tversky, 1979 pp. 274). Thus, although prospect theory seems to provide good explanations for decisions from description, findings from decisions from experience may contradict those predictions from prospect theory in many cases (Hertwig, 2012).

Although prospect theory (Kahneman & Tversky, 1979) has been a prominent model to explain human-choice behavior in descriptive choices, a comprehensive model that can explain decisions from experience has not yet been found. In fact, a challenge in understanding the cognitive processes involved in making decisions from experience is the proliferation of highly task-specific cognitive models that often

To test theories of human behavior, we use *computational models*: representations of some or all aspects of a theory as it applies to a particular task or context. Thus, the value of models is that they can solve concrete problems and provide explicit mathematical and computational representations of a theory, which can then be used to make predictions of behavior.

IBLT constructs and processes were implemented into a computational model (called Cog-IBLT) that helped make the theory more explicit, transparent, and precise (Gonzalez et al., 2003). Cog-IBLT demonstrated the overall mechanisms and learning process proposed by the theory in a dynamic and complex resource allocation task (the "water purification plant", reported in Gonzalez et al., 2003). Cog-IBLT was constructed within the ACT-R cognitive architecture (Anderson & Lebiere, 1998), using the cognitive mechanisms existent in ACT-R. Specifically, Cog-IBLT used the ACT-R's experimentally-derived mathematical representations of: *Activation* (a value that determines the usefulness of an instance from memory and experience and the relevance of the instance to the current context); *Partial Matching*

Instances in a model of the decision from experience paradigms (e.g., that shown in Figure 1) have a much simpler representation compared to instances in Cog-IBLT or in other IBL models. The instance structure is simple because the task structure is also simple. Each instance consists of a label that identifies a decision option in the task and the outcome obtained. For example, (Left, $4) is an instance where the decision was to click the button on the left side and the outcome obtained was $4. The details of this IBL model and its relevance were fully explained in Gonzalez and Dutt (2011), but the main aspects of this model are summarized here.

The IBL model of decisions from experience ("IBL model" hereafter) assumes that choices from experience are based on either a repetition of past choices (i.e., "inertia") or on the aggregation of past experiences instances of payoffs in memory that have been observed as a result of past choices (i.e., "blending"). At trial , the model starts with a random choice between the two options. Then, in each trial , the model first applies a probabilistic rule (based upon a free parameter called *pInertia*) to determine whether to repeat its choice from the previous trial or not. If this probabilistic rule fails, then inertia does not determine the choice and the model chooses the option with the highest *blended* value. An option's blended value is a weighted average of all observed payoffs on that option in previous trials. These observed payoffs are stored as instances in memory and are weighted such that payoffs observed more frequently and recently receive a higher weight compared to the less frequent and distant payoffs. This weight is a function of the recency and frequency of the

Formally, the model works as follows:

In choose randomly between the two choice options (1)

For each trial ,

If the draw of a random value in the uniform dis Pw$%¿I&u4&u0 1D$gQYU$Q4%Í¼ where refers to the total number of payoffs observed for option up to the last trial, and is a noise value defined as (Lebiere, 1998). The variable is a free noise parameter expected to capture the imprecision of recalling instances from memory from one trial to the next.

The activation of each instance in memory depends upon the activation mechanism originally proposed in the ACT-R architecture (Anderson & Lebiere, 1998). The IBL model uses a simplified version of that activation mechanism. In each trial *t*, activation of an instance is

(4)

where is a free decay parameter, and refers to previous trials when payoff contained in the instance was observed (if a payoff occurs for the first time in a trial, a new instance containing this payoff is created in memory). The summation will include a number of terms that coincides with the number of times that a payoff has been observed after it was created (the time of creation of instance itself is the first timestamp) variability in behavior from one trial to the next. The variable is the same noise parameter defined in equation 3 above. A high implies a high noise in activation.

The most recent developments of the IBL model of decisions from experience are important given the simplicity of this model and the broad predictions that it can make (e.g., Gonzalez & Dutt, 2011; Gonzalez et al., 2011; Lejarraga et al., 2012). Next section describes some examples of what the model is able to explain and what the model in its current form does not explain. All examples below rely on two parameters: the decay, , and the noise with values 5.0 and 1.5 respectively. However, the models reported below vary in the inclusion or not of the *pInertia* parameter (see Dutt & Gonzalez, 2012 for a discussion on the value of this parameter), and also on the specific values of the parameters. As explained next, we have used a fit and generalization procedure, in which the parameters values are fit to particular data sets and then used these parameters to predict the behavior in a new data set.

What the IBL model explains and what it does not explain

Existent demonstrations from IBL models suggest the generality of the theory, and not only the descriptive power of the theory but the explanatory one. That is, the theory not only describes the kind of constructs and processes existent in dynamic decision making, but it helps explain why decision making in dynamic tasks occur in the way described. But with generality and robustness also comes the lack of specificity: What are the effects and phenomena that the IBL model can explain and predict? Here we first summarize this tradeoff between generality and specificity, then we present the concrete phenomena that the model in its current form is capable and not capable of explaining.

What the IBL model explains.

Two comprehensive and important demonstrations of the IBL model robustness are the fitting and predictions obtained against a large and publicly available data set, the TPT (Erev et al., 2010). TPT was a competition in which different models were submitted to predict choices made by experimental participants. Competing models were evaluated following the generalization criterion method (Busemeyer & Wang, 2000): they were fitted to choices made by participants in 60 problems (the estimation set) and later tested using the parameters that best fitted the estimation data set to predict a new set of choices in 60 problems (the test set). This process of fitting and generalization procedure is useful as generalization is regarded as pure prediction of behavior.

TPT involved 2 types of experimental paradigms of decisions from experience, Sampling and Repeated choice; and all the problems in the TPT involved a choice between two options:

Safe: M with certainty

Risky: H with probability Ph; L otherwise (with probability 1-Ph)

A safe option offered a medium (M) payoff with certainty, and a risky option that offered a high (H) payoff with some probability (pH) and a low (L) payoff with the complementary probability. M, H, pH, and L were generated randomly, and a selection algorithm assured that the 60 problems in each set differed in domain (positive, negative, and mixed payoffs) and probability (high, medium, and low pH).

An example of the IBL model

can be seen, the IBL model accurately predicted learning in most of the problems (see detailed tests in Lejarraga et al., 2012).

Figure 2. Learning curves from human and IBL model data in the test set of the TPT. Each panel represents one of the 60 problems, each problem ran for 100 trials (both for the IBL model and human data),and the panels show the proportion of risky choices averaged in blocks of 25 trials. The SD in each graph denotes the squared distance between the observed R-rate and the IBL predictions across 100 trials. The IBL model was run in exactly the same experimental paradigm as humans were. The model included the same simulated participants as the human data set.

The 60 problems represent a large diversity of behavioral effects, and in creating this diversity of problems, the organizers of the TPT (Erev et al., 2010) aimed at extending the traditional view of using counter-examples of particular behavioral effects by demonstrating the robustness of general learning effects. This demonstration and additional ones in Lejarraga et al. (2010) and in Gonzalez and Dutt (2011) indicate the IBL model to capture these general learning effects too.

However, reliance on quantitative model comparison and numerical model predictions may lead this work to need of a "help line" (Erev et al., 2010) to guide potential users on what phenomena that this model can explain and the predictions that it can and cannot currently make. Although the TPT problems represent a large diversity of behavioral effects, these are difficult to isolate. This is because the problems were created with an algorithm that randomly selected outcomes and probabilities in such a way that 1/3 of the problems involve rare High outcomes (Ph<0.1) and about 1/3 involve rare Low outcomes (Ph>0.9); also 1/3 of the problems are in the gain domain (all outcomes are positive) and 1/3 are in the loss domain (all outcomes are negative). Thus, effects such as those found in other studies (e.g., Erev & Barron, 2005) may be difficult to isolate in the TPT .

We aim to address the question of robustness and specificity for the IBL model in the following sections, where I summarize results from the model in data sets where different type of phenomena were clearly identified: payoff variability effect, underweighting of rare events, loss rate effect, individual differences (Erev & Barron, 2005), probability matching, and adaptation to nonstationary environments (Lejarraga et al., 2012).

The payoff variability, underweighting of rare events, and loss rate effects.

Erev and Barron (2005) demonstrated

problems differ on the variance of the two payoff distributions. We developed a computer program for data collection and we ran an experiment where each of 60 participants, undergraduate and graduate students at Carnegie Mellon University, worked on one of the three problems. We followed almost identical instructions as in the original experiments: individuals did not receive any information about the payoff structure. They were told that their task was to select one of the alternatives by clicking on one of two unmarked and masked buttons on the screen and were not informed of the trial number. They were provided with the payoff value of the button they clicked on. Payoffs were drawn from the distribution associated with the selected button. There are two differences between our methods and Erev and Barron (2005): (1) we did not use a performance-based incentive structure. Participants were paid a flat fee for performing the repeated choice task, and (2) we ran 400, rather than 200, trials for all problems to better explore learning effects. The average proportions of maximization (i.e., Pmax, the rate choices with the highest expected value) in our data set are very similar to those reported in Erev and Barron (2005). The average Pmax for the second 100-problem block (i.e., Pmax2) was 0.82, 0.61, and 0.50 for Problems 1, 2, and 3 respectively (compared to .90, .71, and .57 in Erev and Barron (2005)). The slight but generally lower Pmax2 values in our replication may be due to the difference in the performance-based incentive.

Figure 3 shows the proportion of maximization (Pmax) choices from humans (dark lines) and those from the IBL model (dotted lines) in each of the three problems.

Figure 3. The payoff variability effect in Human (dark lines) and IBL Model (dotted lilnes). data.

the risky option more often because it results in the maximum Blended value. In Problem 3, the risky alternative provides some higher payoffs (e.g., 21), half of the time which raises its expected value and leads to its selection more often. But the value of the risky alternative appears to quickly even out or decrease over time as a series of poor payoffs (e.g., 1) may lower its expected value and make the certain alternative (i.e., 10) more attractive, which in turn would increase the activation of this option by its more frequent selections.

Additional demonstrations of IBL predictions of the Payoff Variability, underweighting of rare events, and loss rate effects.

We ran the IBL model in the 40 problems reported in Erev and Barron (2005), which belong to the three effects described above. We ran the IBL model in each problem over the course of 400 trials for 100 simulated participants. The set of simulations resulted in the predicted learning curves summarized as the average Pmax in four blocks of 100 trials each. Figure 4 shows the learning curves for humans and for the IBL model. The Pmax per block (100 trials in each block) is shown for each of the 40 problems from Erev and Barron (2005)[1]. The figure shows that the IBL model can account for problems that demonstrate the *payoff variability* effect (Problems 1 to 22), the *underweighting of rare events* (Problems 23 to 25), and the *loss rate* effect (Problems 26 to 40).

[1] The human data reported in this section were obtained from Ido Erev and Greg Barron.

Figure 5. Distribution of proportion of maximization in the second block (Pmax2) over Humans and those produced by the IBL model with the simulated participants in 32 of the 40 problems reported in Erev and Barron (2005) and corresponding to the behavior in Figure 4, for the second block. Each panel represents a problem and the distributions of participants' proportion of maximizations. The y-axis shows the proportion of participants (Humans, dark bars, and simulated by IBL model, white bars).

Probability

chips, while incorrect predictions of low frequency lights cost 1 chip. In the 1:1 condition, incorrect predictions cost 1 chip for both lights. When the two lights occurred with the same frequency (in the 50-50 condition), the light assigned a higher gain was also assigned a lower cost.

The IBL model and the predictions as compared to the results in Myers et al. (1961) were reported in Lejarraga et al. (2012) and reproduced here in Figure 6. The figure shows the mean number of choices for one of the options across participants in each of the 27 problems of Myers et al., (1961). The figure shows accurate predictions of the IBL model (white bars) compared to human data (dark bars) in all the 27 problems.

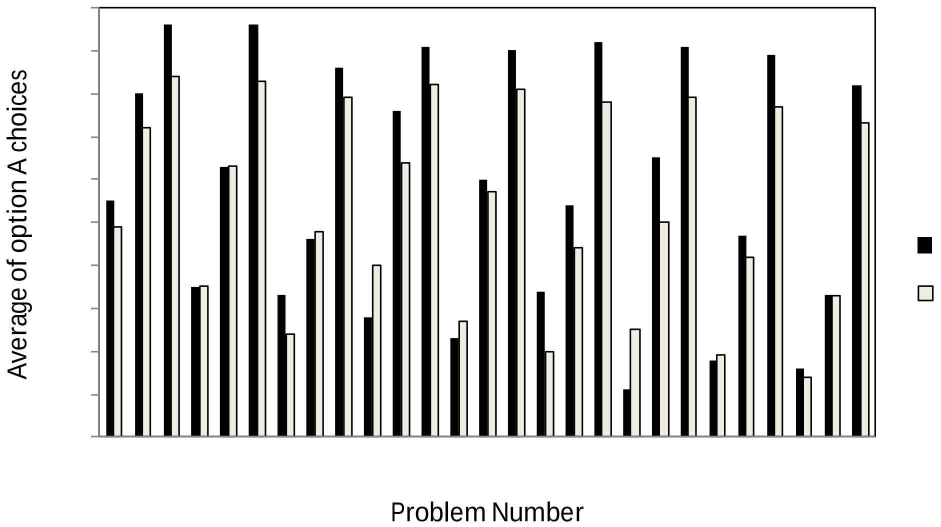

Figure 6. Average choices of option A in 27 problems of Myers et al., (1961) probability learning experiment. The predictions of the IBL model for each problem (white bars) are close to human data (dark bars). For details on the numerical comparison and explanations of the data set see Lejarraga et al. (2012).

Adaptation to nonstationary environments.

Rakow and Miler (2009) explored repeated choice in situations where the outcome probabilities for one of the two options changed over trials. In their

Experiment 1, 40 participants made 100 repeated choices between two risky options in four problems. In all of these problems, each of two options involved a positive and a negative outcome, so participants could win or lose money with each decision. The novelty of the problems studied by Rakow and Miler (2009) is that for one of the options, the probability of the positive outcome remained constant across trials (i.e., the stationary option, S), while this probability changed across trials in the other option (i.e., nonstationary option, NS). Changes in the probabilities for the NS option were gradual: the probability changed .01 per trial and over 40 trials. For example, problem 1 involved a choice between S that offered 10 with a .7 probability or -20 otherwise, and NS that initially offered 10 with a .9 probability or -20 otherwise. From trials 21 to 60, the probability of 10 in NS reduced by .01 in each trial, such that the probability of 10 in m[bat12 7( )]TJBT1 0 0 1 (pr)-16(o)]TJ16(w)-27(i)-5(r)-16(yd)-10( ))-27(i)-5(r)-1

in fhe probability oir

Although the IBL model provides robust predictions across a wide diversity of problems and explains a good number of well-known effects in decisions from experience, the model is not expected to predict behavior accurately in a number of situations. Below there are examples of situations in which the model does not provide accurate predictions. We know there might be many other effects that the model cannot predict and we hope to address the model's miss-predictions in future research.

Pure Risk Aversion

In the demonstrations of the payoff variability effect, Erev and Barron (2005) interpreted the difference between problems 1 and 3 (see Figure 2) as reflecting risk aversion (the high alternative is less attractive when the payoff variability increases), and the difference between problems 1 and 2 as reflecting risk seeking preferences (the low alternative is less attractive when its payoff variability increases). In these problems, however, risk is confounded with expected value, and thus it cannot be interpreted cleanly as a pure risk aversion effect. To explore the pure risk aversion effect, we collected data on a fourth problem not reported in Erev and Barron (2005), in which alternatives are of equal value but they only differ in the variability of the payoff:

Problem 4.	Certain	11	points with certainty
	Risky	21	points with probability 0.5
		1 Otherwise	

Using the same methods as in the first 3 problems, we collected data from 20 participants in problem 4. Results shown in Figure 8 indicate that humans starting at an indifference point (solid line), reduce the proportion of risky choices over time.

The IBL model in contrast (dotted line), starts with a larger preference towards the certain alternative (11) than the risky alternative (21,.5; 1,.5) and moves towards indifference over time. Although the effect is relatively small, the trends are in opposition to the humans', and they would be expected to continue in the same direction with even more practice.

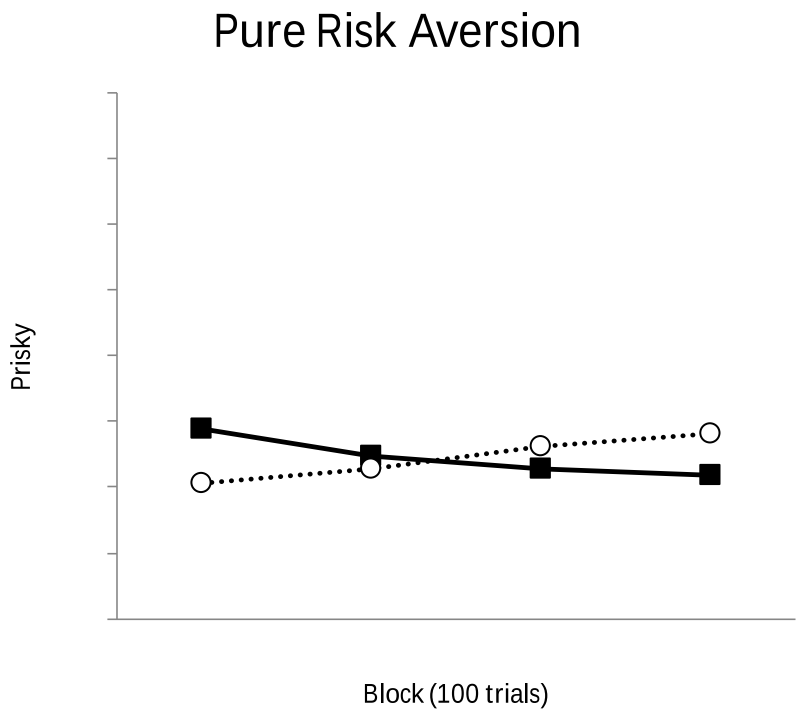

Figure 8. Average human proportions of risky choices (solid line) and the predictions of the IBL model (dotted line), in Problem 4 during 400 trials, averaged in 4 blocks of 100 trials each.

The key insight is that initial experiences of the "1" outcome in the risky option produce a higher blended value for the certain alternative (11) than the risky alternative in the IBL model. The periods in which the risky alternative is selected and the lowest outcome (i.e., 1) is obtained must be longer than the periods of selecting the certain alternative in the first block. Over time, the model "balances out"

experience in the sampling paradigm of the TPT indicate no difference in risky behavior between gains and losses (= .308, p=.580). The IBL model, however, predicts a difference between gains and losses, which although small, it is significant (= 12.462, p<.001). These effects are illustrated in Figure 9. Interestingly, human behavior as well as the IBL model prediction are in disagreement with the predictions from prospect theory: Humans do not show higher risk-seeking tendency in problems involving losses than gains and the IBL model, shows a higher tendency for risky cognitive actions. However, we have started to investigate how the IBL model may account for situations involving two or more individuals involving non-cognitive aspects (e.g., emotions, power, trust). We propose that IBL models may also help in understanding how conflictual social interactions are influenced by the prior experiences of the individuals involved and by the information available to them during the course of interaction (Gonzalez & Martin, 2011). Some initial steps have been taken to use IBL models in multi-person games. For example, Gonzalez and Lebiere (2005) reported a cognitive model for an iterated prisoner's dimella (IPD), initially reported by Lebiere, Wallach, and West (2000), that assumes instances are

payoff. More recently, the IBL model was used in more complex multi-person task, the market entry game (Gonzalez et al., 2011). This model, which obtained the runner-up prize in a modeling competition, shares basic features with IBL models of individual choice (e.g., Lejarraga et al., 2012), and importantly no explicit modifications were included in the model to account for the effects of the market entry task.

Many models of individual decisions from experience are incapable of representing human behavior in social contexts. For example, Erev and Roth (2001) noted that simple reinforcement learning models predicted the effect of experience in two-person games like the Iterated Prisoner's Dilemma (IPD) only in situations where players could not punish or reciprocate. A simple model predicts a decrease in cooperation over time, even though most behavioral experiments demonstrate an increase in mutual cooperation due to the possibility of reciprocation (Rapoport & Chammah, 1965; Rapoport & Mowshowitz, 1966). To account for the effects of reciprocation, Erev and Roth (2001) made two explicit modifications to the basic reinforcement learning model: if a player adopts a reciprocation strategy, he will cooperate in the next trial only if the other player has cooperated in the current trial; the probability that a player continues to do so will depend on the number of times the reciprocation strategy was played. Although these tweaks to the model may accurately represent the kind of cognitive reasoning that people actually use in the IPD, they are unlikely to generalize to other situations with different action sets or outcomes. The IBL model appears to account for these reciprocity effects without the need for explicit and situation-specific rules (Gonzalez, Dutt, Martin, & Ben-Asher, 2012; in preparation; Juvina et al., 2011). However, much work is needed for understanding how the IBL model can be extended to account for the effect of non-cognitive variables (e.g., emotions, social considerations such as power, fairness, envy, etc.) on decision making.

Conclusions

Research on decisions from experience has demonstrated great potential to expand our understanding of the processes involved in making decisions. Experimental and cognitive modeling approaches to study of experience-based choice help open a window to understanding processes beyond the observable choice. With simple experimental paradigms, researchers have improved our understanding of the processes that lead to a choice, such as the recognition of alternatives, the formation of preferences, the evaluation of outcomes, the integration of experiences and the projection of costs and benefits. With cognitive models, researchers have helped to explain how these processes develop, and to predict behavior in some novel circumstances.

A problem, which I have aimed to address in the past years, is the lack of a comprehensive model for experience-based choice behavior and the proliferation of task-specific models of decisions from experience. Several on-going efforts have addressed this issue in many different ways through comprehensive model comparison and demonstrations (Gonzalez & Dutt, 2011; Lejarraga et al., 2012), and through model prediction competitions (Erev, Ert, & Roth, 2010; Erev et al., 2010). These efforts are converging over how decisions from experience are explained: via cognitive memory processes, including recency and frequency of events. Our explanations come from models based on IBLT that have shown robust and accurate predictions in multiple tasks.

This chapter summarizes the history of IBLT and IBL models. Furthermore, it highlights and attempts to start addressing an important problem in this research program: the robustness and specificity tradeoff. Although the IBL models have shown robustness and generality, they also need to clearly and more specifically guide the potential users of these models to explain concrete phenomena in decision sciences. We summarized some phenomena that the IBL model explains: payoff variability effect, underweighting of rare events, loss rate effect, individual differences, probability matching, and adaptation to nonstationary environments. We also summarized some phenomena that the model in its current form is unable to capture: the pure risk aversion effect, more risk seeking in losses compared to in gains domains, and emotions, social, and non-cognitive effects. Future research will address these and many other challenges that the IBL model faces.

References

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Edwards, W. (1962). Dynamic decision theory and probabilistic information processing. *Human Factors, 4*, 59-73.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review, 112*(4), 912-931.

Erev, I., Ert, E., & Roth, A. E. (2010). A choice prediction competition for market entry games: An introduction. *Games, 1*, 117-136.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S., Hau, R., et al. (2010). A choice prediction competition for choices from experience and from description. *Journal of Behavioral Decision Making, 23*, 15-47.

Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making, 21*, 575-597.

Erev, I., & Roth, A. E. (2001). Simple reinforcement learning models and reciprocation in the Prisoner's Dilemma game. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 215-231). Cambridge, MA: MIT Press.

Ert, E., & Erev, I. (2011). On the descriptive value of loss aversion in decisions under risk. Unpublished manuscript.

Gibson, F. P., Fichman, M., & Plaut, D. C. (1997). Learning in dynamic decision tasks: Computational model and empirical evidence. *Organizational Behavior and Human Decision Processes, 71*(1), 1-35.

Gilboa, I., & Schmeidler, D. (1995). Case-based decision theory. *The Quarterly Journal of Economics, 110*(3), 605-639.

Gilboa, I., & Schmeidler, D. (2000). Case-based knowledge and induction. *IEEE Transactions on Systems, Man, and Cybernetics-PART A: Systems and Humans, 30*(2), 85-95.

Gonzalez, C. (2004). Learning to make decisions in dynamic environments: Effects of time constraints and cognitive abilities. *Human Factors, 46*(3), 449-460.

Gonzalez, C. (2005). Decision support for real-time dynamic decision making tasks. *Organizational Behavior & Human Decision Processes, 96*(2), 142-154.

Gonzalez, C., Best, B. J., Healy, A. F., Bourne, L. E., Jr, & Kole, J. A. (2010). A cognitive modeling account of simultaneous learning and fatigue effects. *Journal of Cognitive Systems Research, 12*(1), 19-32.

Gonzalez, C., & Dutt, V. (2010). Instance-based learning models of training. In *Proceedings of the Human Factors and Ergonomics Society 54th Annual Meeting* (pp. 2319-2323). San Francisco, CA: Human Factors and Ergonomics Society.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating decisions from experience in sampling and repeated choice paradigms. *Psychological Review, 118*(4), 523-551.

Gonzalez, C., Dutt, V., & Lebiere, C. (in press). Validating instance-based learning mechanisms outside of ACT-R. *Journal of Computational Science*.

Gonzalez, C., Dutt, V., & Lejarraga, T. (2011). A loser can be a winner: Comparison of two instance-based learning models in a market entry competition. *Games, 2*(1), 136-162.

Gonzalez, C., Dutt, V., Martin, J., & Ben-Asher N. (2012). Decisions from Experience in Conflict Situations: Cognitive Model of the Effects of Interdependence Information. *Behavioral Decision Research in Management*

Conference. BDRM 2012. June 27-29, 2012. Leeds School of Business, Boulder, CO

Gonzalez, C., & Lebiere, C. (2005). Instance-based cognitive models of decision making. In D. Zizzo & A. Courakis (Eds.), *Transfer of knowledge in economic decision-making* (pp. 148-165). New York: Macmillan (Palgrave Macmillan).

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science, 27*(4), 591-635.

Gonzalez, C., & Martin, J. M. (2011). Scaling up Instance-Based Learning Theory to account for social interactions. *Negotiation and Conflict Management Research, 4*(2), 110-128.

Haruvy, E., & Erev, I. (2001). On the application and interpretation of learning models. In R. Zwick & A. Rapoport (Eds.), *Advances in experimental business research* (pp. 285-300decisM49(I)3318 428.1 Tm[()-16(pp. 281 0 0 1 816(e)-5(v,)o.)10( )-49(i)-

Schunn, C. Lebiere & P. Munro (Eds.), *Proceedings of the Sixth International Conference on Cognitive Modeling* (Vol. 420, pp. 178-183). Pittsburgh, PA: Lawrence Erlbaum Associates Publishers.

Myers, J. L., Reilly, R. E., & Taub, H. A. (1961). Differential cost, gain, and relative frequency of reward in a sequential choice situation. *Journal of experimental Psychology, 62*(4), 357-360.

Pew, R. W., & Mavor, A. S. (1998). *Modeling human and organizational behavior*. Washington: National City Press.

Rakow, T., & Miler, K. (2009). Doomed to repeat the successes of the past: History is best forgotten for repeated choices with nonstationary payoffs. *Memory & Cognition, 37*(7), 985-1000.

Rapoport, A., & Chammah, A. M. (1965). *Prisoner's dilemma: A study in conflict and cooperation*. Ann Arbor: University of Michigan Press.

Rapoport, A., & Mowshowitz, A. (1966). Experimental studies of stochastic models for the Prisoner's dilemma. *System Research and Behavioral Science, 11*(6), 444-458.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124-1131.

Zsambok, C. E., & Klein, G. A. (Eds.). (1997). *Naturalistic decision making*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.