Andreas Duus Pape: Research Statement

My primary vein of research is the application of agent-based modeling techniques to decision theory, a subfield of game theory. This application is a novel computational method, the Case-based Software Agent or CBSA. Agent-based modeling techniques are computational methods, which have, before CBSA, not been applied to decision theory. CBSA has provided a number of empirical insights into human behavior by connecting existing psychological experiments to my field of decision theory. These psychological experiments revolve around 'concept learning' (or, synonymously, 'classification learning') which is the study of how humans and other learners acquire ideas or concepts or theories or models. CBSA is the first bridge between concept learning in psychology and decision theory in economics. There are three core concepts behind CBSA: concept or classification learning; case-based decision theory; and the new kind of empirical, agent-based game theory. I will explain each core concept, relate each to CBSA, and extend it to the other papers in this vein.

CBSA will revolutionize the economic understanding and analysis of the dynamics of social systems. How? As I have shown, CBSA can be calibrated to and tested against human data in any setting which can be represented in a game-theoretic form. As we calibrate and test CBSA against more experimental and field data, we will learn whether and when CBSA provides an adequate explanation of human behavior in these settings, and we will likely extend CBSA in new, useful ways to capture unforeseen nuances of human choice. If and when CBSA provides a robust explanation of human behavior in various settings, it can be put in other agent-based policy models, such as ones describing environmental resource management, markets and government intervention, or tax and voting systems (see secondary vein below). This will provide a kind of grounded, subtle, and precise simulation that economics and policy analysis has never had access to before: a new kind of virtual laboratory for public policy experimentation that is rooted in empirical, psychological facts about human choice patterns.

My secondary vein of research is applied game-theoretic and agent-based modeling of policy questions, including environmental policy. In this vein of work, I develop models collaboratively with applied economists, where the goal is a model that captures essential elements of the problem of interest. This is related to the first vein in several ways. First, of course, many of the techniques are the same. Second, the process of modeling applied questions often leads to insights into how modeling works.

My tertiary vein of research is axiomatic decision theory. Axiomatic decision theories, also called 'representation theorems,' state that when choice behavior satisfies certain axioms (e.g. transitivity or completeness), then a mathematical representation of a particular functional form can be written, which translates the choice behavior into commonly used mathematical objects, such as utility functions and belief distributions. Examples include expected utility (von Neumann & Morgenstern 1944), subjective expected utility (Savage 1954), ambiguity averse maximin expected utility (Ellsberg 1961, Gilboa & Schmeidler 1989), and case-based decision theory (Gilboa & Schmeidler 1995, Gilboa & Schmeidler 1996). In one contribution in this we applied ambiguity aversion to optimal auctions. In another I constructed a new representation theorem.

1 Primary Vein: The Case-Based Software Agent (CBSA)

The cornerstone of my primary vein of research is "Evaluating Case-Based Decision Theory: Predicting Empirical Patterns of Human Classification Learning," co-written with Binghamton University psychologist Kenneth J. Kurtz (I will refer to this paper hereafter as Pape/Kurtz). In this paper, we introduce the computational agent I designed based on a choice theory called Case-based Decision Theory (Gilboa & Schmeidler 1995, Gilboa & Schmeidler 1996). With my agent, the Case-based Software Agent or CBSA, we test whether case-based decision theory can adequately explain observed human choice behavior in psychological 'concept learning' experiments; we find that it does.

1.1 Concept Learning

The nature of concept learning is a core question in cognitive science and psychology. In a broad sense, a 'concept' means a way of organizing or perceiving a choice problem: an idea or theory or model of a problem. One way it is formalized is to consider a 'concept' to be a subset of a given set. An example: suppose there are eight animals: four horses and four dogs. 'Horse' and 'dog' are concepts about these animals. "Horses presumably have something in common (not shared by dogs) such that, after one name has been learned for three horses, the extension of this same name to the fourth horse requires little if any further learning (Shepard, et al. 1961)."

The laboratory study of concept acquisition dates back to at least Shepard, Hovland, and Jenkins in 1961 (hereafter SHJ). In their study, eight objects of three binary dimensions–size (big and small), color (dark or light), and shape (triangle or square)–are presented to test subjects, who are asked to classify the objects as type A or B. The experimenter knows the 'true' category of each object, and tells the test subject whether she is correct. This repeats. There are more or less complicated classifications: for example, "Type I" classifications involve only one dimension: so, for example, all squares are A and all ovals are B, while "Type VI" classifications involve all three dimensions for correct classification. The speed and accuracy of acquiring these concepts–that is, the A and B categorizations corresponding to the "Type" number–is what is measured in these experiments. These are the empirical facts that a theory must explain to be considered by psychologists as an adequate explanation of concept acquisition.

My agent CBSA is one of only two theories that that explain the speed and ordering in this series of experiments.[1,2] This means that, according to the empirical measures of the field of psychology, CBSA appears to be an adequate explanation of human behavior in this domain. My work is an important contribution to the field of Behavioral Economics. Whereas the bulk of the field extends models from psychology and applies them to economic data using economic statistical methods (e.g. Kahneman & Tversky 1979, Tversky & Kahneman 1992, Laibson 1997, Fudenberg & Levine 2006), my work opens a new thread by doing the reverse: I extend a model from economics and apply it to psychological data using psychological statistical methods.

Along with matching human speed and ordering in these experiments, we are also able to establish other facts about Case-based Decision Theory; for example, the so-called 'similarity function' in CBDT which best

[1]The series of experiments include Shepard et al. (1961), Nosofsky, et al. (1994), Nosofsky & Palmeri (1996), Smith, et al. (2004), Minda, et al. (2008), and Kurtz, et al. (2013).

[2]The other model is ALCOVE (Kruschke 1992, Nosofsky & Palmeri 1996). Two other models have been shown to explain some of the experimental results: Love, et al. (2004) and Kurtz et al. (2013).

matches human behavior is of a form consistent with independent evidence in psychology.[3] This work has also provided a stunning set of graphs of human vs. simulated behavior; I refer you to the paper.

1.2 Case-based Decision Theory

Case-based decision theory is a formalization or framework of a choice problem faced by a decision maker, in which the decision maker updates his memory with new experiences, and uses that memory to forecast outcomes of future actions. Decision theory is study of individual choice in the context of game theory (the study of strategy) and economics; other decision theories include Expected Utility Theory, the most commonly used model of choice under uncertainty. In case-based decision theory, the environment of the decision maker is described as having three components: there is a vector of information that the decision maker observes before her choice, called the 'problem' or 'circumstance;' there is a set of actions that the decision maker may choose from in response to the circumstance; and there is a set of outcomes that may result from the choice. In the case of the SHJ series of experiments, the 'circumstance' facing the decision maker is the object she is currently called upon to categorize (and therefore the vector of information has three binary dimensions); the 'actions' available to the decision maker are the categories A and B; and the 'outcomes' are either that the categorization was right or wrong. A triplet consisting of a circumstance, the action taken in response to the circumstance, and the outcome of the action, is called a 'case.' (Hence the name "case-based.") A single case can be thought of as a complete learning experience.

Case-based Decision Theory also supposes a structure for how the decision maker represents the problem. This representation also has three components: First, the decision maker has a memory, which is the set of cases that she has experienced. Second, the decision maker has a 'similarity function,' which is a mathematical function which describes how *similar* two circumstances are. For example, the decision maker presumably thinks that a small, dark triangle is more similar to a large, dark triangle than it is to a large, white square. Third, the decision maker has a utility function, standard in many economic models; it is a mathematical function which assigns a number to each outcome, where higher numbers are better in the mind of the decision maker.

In Pape/Kurtz, we extend case-based decision theory in three ways.

First, I placed this heretofore 'blackboard' model in a computational framework, so the theory could be used to create simulations, a so-called 'Agent-based model.' This is a novel approach in decision theory; prior to this work, decision theory had only three approaches: (1) mathematical structures to represent human choice using axioms; (2) thought experiments, such as the Ellsberg urn experiment (Ellsberg 1961); and (3) more recently, laboratory experiments. Simulation is a valuable method to add to this bouquet because it allows for studying the implications of case-based decision theory in other situations, such as bringing CBDT to bear on experiments that were not designed with CBDT in mind. This is precisely what we do in Pape/Kurtz.

Second, we deepened empirical understanding of the functional form of similarity by incorporating and testing insights from psychology. We established that the best-fitting similarity function in CBSA is one that matches independent research into the concept of similarity in psychology. We found that, consistent with research in psychology (Shepard 1987), similarity functions that are decreasing in vector distance induce

[3]See below for a discussion of the details of Case-based Decision Theory.

the best match to human data. This is some evidence that 'similarity' in case-based decision theory and 'similarity' in psychology refer to the same concepts.

Third, we added two forms of imperfect memory: errors *writing to* memory and errors *reading from* memory. Imperfect memory is necessary to match the speed of human learning; with perfect memory the CBSA learns the solution to the problems about ten times as quickly as humans do. Imperfect memory allowed us to match this speed of learning without incorporating the Luce Choice Rule (Luce 1959, Luce 1977), in which options are chosen with a likelihood proportional to their perceived value. This rule is popular in consumer research and psychology, but it is irrational: according to rationality, if one option is perceived to be better than another, it should always be chosen, not simply be more likely to be chosen. Imperfect memory allows for error–necessary for empirically matching human data–without compromising the rationality of this model.

Before this work, the fields of Agent-based Modeling and Decision Theory did not intersect. I wish to stress this point because it is often lost in the details of my findings. The facts we have learned about human behavior with this tool are, in fact, quite interesting. But if one is to evaluate my research as a whole, thinking that I only have brought results to the table sells this work short. Consider this metaphor: I invented a hammer, and then built a house with it. Looking at the house is only part of the story: one must also look at the hammer, and think about the hammer's role in building future houses. Continuing the metaphor of the hammer, showing the hammer to the field of Decision Theory was not enough. Those trained in agent-based computational economics, the toolmakers in this metaphor, could see how the hammer could be useful. But I wanted decision theory, and therefore game theory, to also recognize the utility of the hammer. And this turned out to be difficult. Originally, I used the hammer to build what amounted to abstract art: it showed that the hammer could be used to build things, but it did not convince Decision Theorists that what it could build could be *useful*. Therefore, the work of developing CBSA had to be attached to new empirical results–these that we learned from concept learning–in order to convince the field of decision theory that this tool was useful. Pape/Kurtz appears in the journal 'Games and Economic Behavior,' the top game theory journal. So I consider its placement validation that decision theory and game theory have accepted that this novel method is useful and the results important.

1.3 Empirical Game Theory

Pape/Kurtz is the first example of what I believe will be the future of empirical game theory: (1) building agents like mine that are consistent with a decision theory, and then (2) simulating their behavior for a variety of parameter values, and (3) finding those parameter values that best explain the human data.[4] As I state in the final paragraph of this paper:

> This computational implementation of Case-based Decision Theory can be calibrated to and tested against human data in any existing experiment which can be represented in a game-theoretic form as we do here. This suggests a model for future studies. As these studies ac-

[4] It's worth noting that, although 'empirical theory' is a contradiction, 'empirical game theory' is not. Both "decision theory" and "game theory" have misleading names; historically, both fields were exclusively theoretical, hence their names. However, they have grown into proper sciences with both a theoretical and empirical component. (Game theory is sometimes defined as 'the science of strategy.') Indeed, some of the best work of both game theory and decision theory are either directly empirical (in that they collect data in laboratory work) or they are theoretical responses to empirical critiques.

cumulate, we will learn whether and when Case-based Decision Theory provides an adequate explanation of human behavior in other decision settings and may also learn which parameters appear to vary by setting and which, if any, remain constant across settings. This could lead to a version of CBDT which can be used to simulate human behavior in a variety of economic models.

The second example of how this method can extend empirical game theory is the paper "Predicting Human Cooperation in the Prisoner's Dilemma Using Case-based Decision Theory," with Todd Guilfoos (hereafter Guilfoos/Pape). While Pape/Kurtz applied CBSA to individual choice data, Guilfoos/Pape applies CBSA to social or group choice. In particular, we apply CBSA and associated methods to aggregate dynamics of cooperation in the repeated Prisoner's Dilemma game, as observed in the experiments performed by Camera & Casari (2009). In these experiments, they observe individuals playing the Prisoner's Dilemma and they vary the amount of information agents have about the other players. In CBSA, this is equivalent to varying the set of variables which appear in the 'circumstance' information vector, so CBSA in some sense automatically comes with a testable empirical hypothesis in this experiment. We not only find a good match with human data; we also find CBDT provides a *better* fit to the dynamics of cooperation than does Camera and Casari's Probit model. This is the first time that CBSA has been directly compared to classic econometric techniques, and is the first time such a result has been found. Guilfoos/Pape demonstrates that the method introduced in Pape/Kurtz is reproducible (i.e. the hammer was used to build a second house). It demonstrates that we can see how much of existing human experimental data CBDT can explain, and learn about parameter values of CBDT at the same time. I intend to write many more papers in this vein and hope that others begin to, as well.

In the future Todd Guilfoos and I intend to run a laboratory experiment about the repeated Prisoner's Dilemma that uses CBSA in the running of the experiment itself: that is, the human subjects will be paired against other humans or against CBSAs, and we will be able to (1) see how effectively and under what circumstances humans can identify whether they are playing against a person or a CBSA (a so-called 'Turing Test') and (2) whether we can statistically distinguish human-human, human-CBSA, and CBSA-CBSA play ex-post.

1.4 Extensions

My work on concept learning has also resulted in some exciting, truly interdisciplinary work, with cognitive scientist Ken Kurtz and artificial-life scientist Hiroki Sayama, named "Complexity Measures and Concept Learning." In this paper we introduce a mathematical metric which explains a large part of existing concept learning results. These results are about the relative difficulty of learning different problems. In essence, the existing literature has been able to explain one ordering of the relative difficulty of problems: the ordering which emerges when adult human learners classify objects with easily distinguishable characteristics such as size, shape, and shading. The literature explains this result with a metric based on the logical complexity of the underlying problem: i.e. the shortest way to express the solution as a set of logical rules. There is another difficulty ordering, which emerges for other known cases, including the cases when adult humans classify objects with characteristics that are not readily distinguished (e.g., brightness, saturation, hue) and for children and monkeys. We are able to explain this ordering with a metric using statistical uncertainty.

This suggests that, when learners do not/cannot form logical rules about the characteristics of objects, they may be using statistical means in their category learning. This is potentially a deep result about how learning occurs.

In the paper "Case-based Learning and the Cobweb Model," with Wei Xiao, we use an empirically-calibrated CBSA to simulate human behavior in a market. We take the decision engine we have calibrated in the above studies and see what happens in a simple macro economy populated with firms who learn as humans do. We are able to show that observing only price can be nearly as good as directly observing the relevant state variable, in terms of convergence, but observing *both* is worse than observing either. Our paper demonstrates CBSA is applicable to domains outside of direct experimental data. This allows for two wholly new types of work: empirical testing of a more traditional kind in economics, where we study observed data from the actual economy in the face of a policy change; and empirical prediction of policies that have not yet been tried. The difference between this and other pure theoretical ventures is that these agents have a firm grounding in *empirical data* about human learning behavior at a basic level.

2 Second Vein: Applied Policy Modeling

I extend agent-based modeling in "An Agent-based Model of Tax Ceilings: Leviathan Extraction and Tax Payment Uncertainty" with Nathan Anderson, Todd Guilfoos, and Jeffrey Schmidt. One criticism of agent-based models is that agents are typically unsophisticated and use simple rules to determine their behavior. Of course, my other work with CBSA helps address this criticism; and in this paper I developed a second method: implementation of rational expectations in an agent-based model. 'Rational expectations' is a paradigm in economics for forming forecasts. The paradigm is: agents in a model have 'rational expectations' if their forecasts of the future are on average correct; that is, if they don't make systematic mistakes. One part of the reasoning behind this paradigm is that the forces of competition should drive out agents who are making systematic errors that could be exploited. It was thought that agent-based models were inconsistent with rational expectations, because agent-based models unfold over time, so the future is unknown when agents construct forecasts. In this paper, we introduce an agent-based model in which rational expectations is achieved. It is a two-period model: in the first period, agents vote on a tax policy when they are uncertain about future economic conditions. In the second period, economic conditions and the effects of the policy are determined. Under the rational expectations framework, agents in the first period ought to be correct about the implications of the policy they are voting on. We are able to achieve this level of sophistication by running the model in reverse chronological order. That is: we run all possible second-period scenarios and determine the effects of each policy in each scenario. Then, we endow the agents with this knowledge of all possible outcomes, and run the first period, in which agents select the policy which they most prefer, given the true set of possible outcomes and their relative likelihoods. This is the first time rational expectations has been achieved in an agent-based model.

This agent-based model is built on the analytical paper "Self-imposed Constraints on Collective Action Under Uncertainty" with Nathan Anderson. In this paper, we investigate under what conditions does, or should, a collective of rational individuals support the imposition of a binding constraint on their own collective action? Our innovation is to allow citizen-taxpayers in a standard political economy model to be risk-averse and uncertain about the future average cost of the collective good, their own future income or wealth, and the future distribution of the tax burden. We show that if citizen-taxpayers face such uncertainty, an agency problem — broadly defined as a flaw in the political process that produces collective decisions incongruent with majority preferences — is neither necessary nor sufficient to justify a binding tax ceiling. This is a striking result, because it contradicts existing theory and can help explain why, empirically, some tax systems are more subject to ceilings than others. This final result leads right into the previous agent-based model with Anderson, Guilfoos, and Schmidt: that agent-based model simulates support for a tax ceiling in Binghamton, NY and Minneapolis, MN, using property tax data from both of those cities.

I designed the agent-based model of the paper "Groundwater Management: The Effect of Water Flows on Welfare Gains," with Binghamton Economics PhD Todd Guilfoos, Neha Khanna (Binghamton University environmental economist), and Karen Salvage (Binghamton University hydrologist). It was published in Ecological Economics 95 (2013): 31-40. In this paper, we model a groundwater aquifer under use for agriculture. There is a "map" which may correspond to an actual physical area or an idealized one. The map consists of a grid of land. Farmers–agents–are located on the map. Each cell on the grid has an amount ("height") of water. This height can change over time through recharge (rain, etc), withdrawal by farmers, and inflow and outflow from/to neighboring cells. The flow between cells is governed by a hydrological rule called "Darcys Law." The specifics of that law are not important here; rather, what is important is that the behavior of water across cells in this aquifer model are governed by principles that are thought to be accurate in the field of hydrology. The farmers each have demand for water that reflects their own cost of withdrawal. The main economic problem here is a negative externality problem: a farmer's use of water increases the cost of water use by her neighbors on the aquifer, and she does not take these costs into account. The central modeling exercise is two-fold: first, design an agent-based computational model of the aquifer itself that captures the true physical properties (flow of water between cells) better than the current models in economics (which assumed that water flows instantly). Second, find a tax policy on water use that solves or, at least, best mitigates the externality problem. While the first is a more standard agent-based modeling problem, the second is fairly novel: it involves putting a choice engine 'on top' of the agent-based model, where the choice engine searches the space of possible tax rates and attempts to find the one that maximizes the welfare of farmers. We show that taking the true physics of water flow into account, versus the models in the literature, significantly increases the value of active management of the aquifer.

With Misuk Seo (Binghamton Economics PhD), I wrote "Reports of Water Quality Violations induce Consumers to buy Bottled Water." The 1996 Safe Drinking Water Act Amendments require that water utilities send quality reports to customers. We test whether receiving WQRs of health violations increases purchases of bottled water. We find that American consumers spend approximately $300 million dollars per year–or about 4% of yearly bottled water expenditures–to avoid water quality violations. With Jessica Harriger (Binghamton Economics PhD) and Neha Khanna, I wrote "Conspicuous Consumption and Inequality," in which we analyze the change in consumer demand following a mean preserving change in consumption inequality when there is conspicuous consumption. We model interdependent preferences including "keeping up with the Joneses" (imitating others) and "running away from the Joneses" (distinguishing oneself from others) with multiple peer groups and peer group effects (envy and snob effects). In both of these papers, my contribution was to help construct and understand the model by which the question could be investigated and answered.

3 Third Vein: Axiomatic Decision Theory

In work started in graduate school with my advisor Emre Ozdenoren and our co-author, Subir Bose, we wrote "Optimal Auctions with Ambiguity." 'Ambiguity' is formal term in decision theory, which refers to one of the two kinds of uncertainty: "risk" refers to uncertainty with some *known* probability distribution, while "ambiguity" refers to uncertainty with an *unknown* probability distribution.[5] In this paper, we study the optimal auction problem, supposing that participants have ambiguous beliefs about each others' perceived value of the objects in the auction. When the bidders face more ambiguity than the seller (i.e. when buyers have less information than sellers about other players' types) we show that (i) given any auction, the seller can always (weakly) increase revenue by switching to an auction providing full insurance to all types of bidders, (ii) if the seller is ambiguity neutral and any prior that is close enough to the seller's prior is included in the bidders' set of priors, then the optimal auction is a full insurance auction, and (iii) in general neither the first nor the second price auction is optimal (even with suitably chosen reserve prices). Since its publication, this paper has, depending on how one counts, garnered nearly sixty citations, including most significantly the following seven citations: two Econometrica articles (in 2006 and 2009), one Journal of Economic Theory article (in 2011), and four Games and Economic Behavior citations (two in 2011, one in 2012, and one in 2013.) These are top three journals in economic theory, and GEB (as mentioned above) is the top game theory journal.

I wrote "Action-Independent Subjective Expected Utility Without States of the World," which appears in Theoretical Economic Letters, 3 (2013): 17-21. This paper develops an axiomatic theory of decision-making under uncertainty that, like Case-based Decision Theory, has no state-space. Some decision theorists feel the state-space is the most problematic part of the expected utility framework: the state-space is the set of all possible conditions of the world that may affect the outcome of the choice faced by the decision maker. Like another decision theory paper by Edi Karni, my paper avoids the state-space by replacing it with bets over outcomes. (Case-based decision theory avoids the state-space by replacing it with an information vector.) While the state-space may be too complicated to correspond to how people 'really' think about a problem, bets on outcomes are easy to conceptualize if one already has access to the set of possible outcomes and the set of possible actions. (Also, of course, bets as a means of soliciting beliefs has a long economic tradition (e.g. Bernoulli 1954).)

4 Conclusion

This research statement describes my three veins of my research. I am a game theorist–specifically, a decision theorist–and an agent-based modeler. I have started a new literature by intersecting decision theory and agent-based modeling, and I think it provides the grounding for a new kind of empirical game theory. I also am a game theorist and agent-based modeler who works well with others, so to speak: I have a long list of projects that are outside of my specific field that I was able to contribute to with my modeling skills, and I have contributed to a large number of collaborative projects with a variety of economists and a number of academics outside of economics: in hydrology, in psychology, and in systems science and bioengineering. Finally, I am a formal decision- and game-theorist who has a high-level understanding of the theoretical and

[5]This new, second sense of uncertainty is attributed to Ellsberg (1961) or Knight (1921).

abstract parts of these fields. These three veins together define me as an economist: what I am, and where I am going.

References

D. Bernoulli (1954). 'Exposition of a new theory on the measurement of risk'. *Econometrica: Journal of the Econometric Society* pp. 23–36.

S. Bose, et al. (2006). 'Optimal auctions with ambiguity'. *Theoretical Economics* **1**(4):411–438.

G. Camera & M. Casari (2009). 'Cooperation among strangers under the shadow of the future'. *The American Economic Review* pp. 979–1005.

D. Ellsberg (1961). 'Risk, ambiguity, and the Savage axioms'. *Quarterly Journal of Economics* **75**:643–669.

D. Fudenberg & D. Levine (2006). 'A Dual-Self Model of Impulse Control'. *The American Economic Review* **96**(5):1449–1476.

I. Gilboa & D. Schmeidler (1989). 'Maxmin expected utility with non-unique prior'. *Journal of Mathematical Economics* **18**:141–153.

I. Gilboa & D. Schmeidler (1995). 'Case-Based Decision Theory'. *The Quarterly Journal of Economics* **110**(3):605–39.

I. Gilboa & D. Schmeidler (1996). 'Case-based Optimization'. *Games and Economic Behavior* **15**:1–26.

T. Guilfoos & A. D. Pape (2013). 'Predicting Human Cooperation in the Prisoner's Dilemma Using Case-based Decision Theory'. *Working Paper* .

D. Kahneman & A. Tversky (1979). 'Prospect theory: An analysis of decision under risk'. *Econometrica: Journal of the Econometric Society* pp. 263–291.

F. H. Knight (1921). *Risk, Uncertainty, and Profit.* Houghton Mifflin, Boston, New York.

J. Kruschke (1992). 'ALCOVE: an Exemplar-Based Connectionist Model of Category Learning.'. *Psychological Review* **99**(1):22.

K. J. Kurtz, et al. (2013). 'Human Learning of Elemental Category Structures: Revising the Classic Result of Shepard, Hovland, and Jenkins (1961)'. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **39**(2):552–572.

D. Laibson (1997). 'Golden Eggs and Hyperbolic Discounting'. *The Quarterly Journal of Economics* **112**(2):443–477.

B. Love, et al. (2004). 'SUSTAIN: A Network Model of Category Learning'. *Psychological Review* **111**(2):309.

R. Luce (1959). *Individual Choice Behavior.* John Wiley.

R. Luce (1977). 'The Choice Axiom After Twenty Years'. *Journal of Mathematical Psychology* **15**(3):215–233.

J. P. Minda, et al. (2008). 'Learning rule-described and non-rule-described categories: A comparison of children and adults'. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **34**(6):1518.

R. Nosofsky, et al. (1994). 'Comparing Models of Rule-Based Classification Learning: a Replication and Extension of Shepard, Hovland, and Jenkins (1961)'. *Memory and Cognition* **22**:352–352.

R. Nosofsky & T. Palmeri (1996). 'Learning to Classify Integral-Dimension Stimuli'. *Psychonomic Bulletin and Review* **3**:222–226.

A. D. Pape & K. J. Kurtz (2013). 'Evaluating Case-based Decision Theory: Predicting Empirical Patterns of Human Classification Learning'. *Games and Economic Behavior* **82**:52–65.

A. D. Pape & W. Xiao (2014). 'Case-based Learning and the Cobweb Model'. *Working Paper* .

L. J. Savage (1954). *The Foundations of Statistics*. Wiley.

R. Shepard (1987). 'Toward a Universal Law of Generalization for Psychological Science'. *Science* **237**(4820):1317.

R. Shepard, et al. (1961). 'Learning and memorization of classifications'. *Psychological Monographs* **75**:1–41.

J. D. Smith, et al. (2004). 'Category learning in rhesus monkeys: A study of the Shepard, Hovland, and Jenkins (1961) tasks.'. *Journal of Experimental Psychology: General* **133**(3):398–414.

A. Tversky & D. Kahneman (1992). 'Advances in prospect theory: Cumulative representation of uncertainty'. *Journal of Risk and uncertainty* **5**(4):297–323.

J. von Neumann & O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.