

The Evolution of Our Preferences: Evidence from Capuchin-Monkey Trading Behavior*

M. Keith Chen[†]

Yale School of Management and Cowles Foundation

Venkat Lakshminarayanan

Laurie Santos

Yale College

Yale Department of Psychology

July 19, 2005

Abstract

Behavioral economics has demonstrated systematic decision-making biases in both lab and field data. But are these biases learned or innate? We investigate this question using experiments on a novel set of subjects – capuchin monkeys. By introducing a fiat currency and trade to a capuchin colony, we are able to recover their preferences over a wide range of goods and risky choices. We show that standard price theory does a remarkably good job of describing capuchin purchasing behavior; capuchin monkeys react rationally to both price and wealth shocks. However, when capuchins are faced with more complex choices including risky gambles, they display many of the hallmark biases of human behavior, including reference-dependent choices and loss-aversion. Given that capuchins demonstrate little to no social learning and lack experience with abstract gambles, these results suggest that certain biases such as loss-aversion are an innate function of how our brains code experiences, rather than learned behavior or the result of misapplied heuristics.

*We would like to acknowledge a very generous Whitebox Advisors' Research Grant, without which this research would not have been possible. Alisia Eckert provided invaluable research assistance throughout this project. Many people gave generously of their time on this and earlier drafts; we would especially like to thank Nick Barberis, Dan Benjamin, Judy Chevalier, Barry Nalebuff, Emily Oster, Jesse Shapiro, Peter Schott, and seminar participants at Brown, UCSD, Caltech, The University of Chicago, and Yale for their insightful comments and suggestions.

[†]Corresponding Author. Comments are welcome at keith.chen@yale.edu or at 135 Prospect St. Box 208200 New Haven, CT 06520. The most recent version of this paper and several accompanying film clips are available at: <http://www.som.yale.edu/Faculty/keith.chen/>

"Nobody ever saw a dog make a fair and deliberate exchange of one bone for another with another dog. Nobody ever saw one animal by its gestures and natural cries signify to another, this is mine, that yours; I am willing to give this for that."

Adam Smith, *Wealth of Nations*

1 Introduction

Over the past few decades, behavioral economists have identified that human decision makers exhibit a number of systematic biases both in the lab and in the field. Two of these biases, reference-dependence and loss-aversion,¹ have received a substantial amount of empirical attention, both from economics and neighboring disciplines such as psychology and sociology. Evidence that agents treat losses differently than comparable gains is found in the behavior of: individual investors (reluctant to realize losses; Odean 1998), house sellers (unwilling to sell below buying price; Genesove & Mayer 2001), and consumers (asymmetrically more sensitive to price increases than decreases in numerous markets; Hardie et al. 1993). Despite the mounting evidence of the importance of this behavior, relatively little attention has been paid to the question of where these biases come from in the first place. Are such systematic biases the result of social or cultural learning and specific environmental experiences? Or could they be more intrinsic to our biology, perhaps resulting from universal mechanisms that arise regardless of context or experience? Put another way, could there be an innate component to our behavioral biases? The answer to these questions of origin may effect how we think about both the potential scope of behavioral biases, and the degree to which we believe market incentives will act to reduce their effects.

Traditionally, economists have remained agnostic as to the origins of human preferences, and usually assume their stability over both time and circumstance. For example, Becker (1976) writes, "generally (among economists)... preferences are assumed not to change substantially over time, nor to be very different between wealthy and poor persons, or even between persons in different societies and cultures."² Indeed, coupled with maximizing behavior and market equilibrium, Becker asserts that the assumption of stable preferences "forms the heart of the economic approach." If much of the fundamental structure of our

¹These biases, along with a probability-weighting function, make up "prospect theory," first introduced in Kahneman and Tversky (1979). For an excellent summary of the recent empirical work on prospect theory, see Camerer (2000).

²It must be stressed that Becker was referring not to preferences over market goods, but to more primitive "underlying objects of choice" such as "health, prestige, sensual pleasure, benevolence or envy." (Becker 1976).

preferences were innate, this would both bolster this assumption and help explain price theory's positive successes.

The origins of our preferences and biases may also be important from a policy perspective; in particular it may be important to distinguish between those behaviors that are *universal* and similar behaviors that are both *universal and innate*. For instance, a learned, non-innate heuristic may arise in many (if not all) cultures if it easily and quickly solves a common task or set of tasks. However, if this universal heuristic is (at least in part) learned, we would not expect it to persist in settings in which it was highly sub-optimal, or in which market forces strongly discipline behavior. This would limit the potential scope and scale for welfare losses, and may suggest that policy interventions which increase feedback or learning may eliminate what losses do exist. In contrast, if a particular bias has a strong biologically-innate component, it may persist even in the face of large individual costs, ample feedback, and market disciplining. This would greatly change both the potential for successful policy and the types of remedies available.

Given the similarity in spirit to questions addressed by biologists and psychologists working in comparative behavior and cognition, we adopt an interdisciplinary approach to this question that applies traditional economic choice methodology far from its usual subjects. When examining questions of origin, it is widely accepted in both cognitive science and evolutionary psychology that a mechanism is most likely evolutionarily ancient if it explains analogous behavior in both humans and primates. That is, since primates and humans are closely related, it is unlikely that a common trait evolved in parallel between our two species, and much more likely common traits evolved once during our common evolutionary heritage. The tufted capuchin monkey (*Cebus apella*) has been widely studied in biology and psychology for just this purpose. They make excellent subjects since they are relatively quick and adept problem solvers, skilled tool users, and a close evolutionary neighbor to humans.³

Exploiting this evolutionary proximity to examine origins of human behavior, we first demonstrate that the standard tools of utility maximization and price theory describe capuchin purchasing behavior in situations akin to those faced by human consumers. By introducing a fiat currency to a capuchin colony we are able to conduct a number of revealed-preference experiments analogous to canonical human choice experiments. Our capuchins have learned through repeated exposure that human experimenters are willing

³Within the set of primates though, capuchins are actually very distantly related to humans. Capuchins diverged from our common ancestral line in what biologists call the new-world primate radiation. This is when all the primates who inhabit the new world split off from the old-world primates, the line humans emerged from. While the exact date of this split is not known, molecular-clock estimates suggest capuchins split off as a genus around 23 million years ago. Estimates of our latest common ancestor date around 40 million years (Schneider et al., 2001).

to trade small metal disks found in their environment for food rewards, and that different humans trade for different rewards at different prices.

Our first set of experiments shows that when endowed with a budget of disks and facing several possible trading options, capuchin purchases look quite sensible. In response to both price and wealth shocks, capuchins display all the hallmarks of rationality; they adjust their purchasing behavior in ways consistent with the Generalized Axiom of Revealed Preferences (GARP).⁴ In this way capuchin choice closely mirrors our own, and admits all the standard tools of utility analysis and price theory. We used this first set of experiments to validate our experimental method, and to provide a context in which to interpret the capuchins' latter departure from rational choice.

Our second set of experiments demonstrates that when faced with decisions involving simple gain-loss frames, capuchins demonstrate both reference-dependence and loss-aversion. Specifically, in our main experiments capuchins express a strong preference for gambles in which good outcomes are framed as bonuses rather than payoff-identical gambles in which bad outcomes are emphasized as losses, and weigh those losses more heavily than comparable gains. Our experiments also allow us to reject most competing models of naïve or unsophisticated choice. In particular, several of our results require a capuchin to choose a trading partner who initially offers a smaller food reward; an ability that several studies of several monkey species and even great apes have failed to find.⁵

We then characterize our results in two ways. First, we discuss what results we can obtain nonparametrically, for example establishing that capuchin display reference-dependence and loss-aversion without specifying the magnitude of either bias. Then, we calibrate a simple utility function that incorporates both loss-aversion and reference-dependence, and nests them in a model of utility maximization. This allows us to estimate parameters such as the capuchins' degree of loss-aversion, and compare these parameters to analogous human results. While parametric results are always to be taken with caution, our analysis shows that

⁴Early papers by Samuelson (1938), Houthakker (1950), and Afriat (1967) established the revealed preference approach to evaluating whether any given set of choices is consistent with rational behavior. Varian (1982) generalized this approach, and showed that the Generalized Axiom of Revealed Preference (GARP), is a necessary and sufficient condition for any set of choices to arise from the maximization of a continuous, concave, weakly monotonic, and locally non-satiated utility function.

⁵Both new-world monkeys and great apes fail to perform optimally on reverse-contingency tasks (games in which an experimenter presents a large reward whenever the agent reaches for a smaller treat, and presents a small reward whenever the agent reaches for the larger treat). However, when the game is modified such that the large reward is presented whenever the agent reaches for a *picture* of the small reward (and the small reward is presented whenever the agent reaches for a picture of the large reward) both apes and new-world monkeys succeed. A contribution of our work to the psychology literature is that we report the surprising result that capuchins have no trouble solving a reverse-contingency task when treats are obtained by exchanging fiat currency for the rewards, rather than simply reaching for them. That is, token-mediated exchange allows primates to overcome the impulse to simply reach for the greater reward – just as picture-mediated choice does in chimpanzees and tamarins.

the relative strength of losses to gains in capuchin decisions (the coefficient of loss-aversion) is indistinguishable from similar human estimates, suggesting common mechanisms underlie these behavioral preferences in our two species.

Arguing against loss-aversion as a learned trait among our capuchins is the novelty of the situation; abstract gambles were first introduced to these capuchins by our experiments, and subjects encountered them alone, away from others. In addition, an abundance of research has shown that capuchins lack all but the most rudimentary forms of social learning. Indeed, the commonly held belief that monkeys are skilled imitators (hence the adage, "monkey see monkey do") has not been born out in either lab or field data.⁶ Knowledge transfers such as imitative learning have been found to be rare in all non-human primates, including apes (see Tomasello and Call, 1997).⁷ In the absence of social learning then, our results suggest an early-evolutionary origin for loss-averse behavior in humans. That is, our results suggest that loss-aversion is innate and evolutionarily prefigured, a function of decision making systems which evolved before the common ancestors of capuchins and humans diverged.

2 Previous Empirical Work and Experiments

Recently several papers have employed methodologies which attempt to shed light on the underpinnings of behavioral biases. A growing literature in the field of neuroeconomics has attempted to use imaging technology to map brain activity as subjects make economic decisions, and correlates these measures of brain activity to subjects' decisions. For example McClure, Laibson, Lowenstein, & Cohen (2004) show the spatial distribution of brain activity is correlated with decisions involving intertemporal choice. While this approach is extremely useful in shedding light on the mechanisms of decision making, the ability of this approach to address questions of universality and stability is limited by the scope of activities that can be scanned (subjects must be securely restrained inside a large magnet) and by the difficulty of translating neural correlates of behavior into causal statements.⁸

⁶For instance, though capuchins are highly-skilled tool manipulators, observing a fellow capuchin successfully use a tool to solve a problem does not help a capuchin learn to use that tool any faster (Visalberghi 1993.)

⁷A notably outlier in this literature, a recent study by Brosnan and deWaal (2004) shows that when forced to directly observe a conspecific displaying a strong choice preference, capuchins subsequently alter their choices in line with the observed behavior. We address this possible concern by not allowing capuchins to observe each other's choices during our experiments. For an excellent summary of the literature on capuchin social learning, see Frigaszy, Visalberghi & Fedigan (2004).

⁸McClure et al (2004) sheds light on the neural mechanisms which facilitate intertemporal decision making; it uses brain scan data to support the already well establish model of Laibson β - δ discounting (Laibson, 1997) without addressing how stable or universal this mechanism might be. In neuroscience and biology more broadly, this problem is conceptualized as one of separately identifying the proximate (mechanical) cause of an behavior and the ultimate cause (evolutionary forces or incentive structures.) In short, the neuroeconomic approach addresses a different set of questions than the approach we employ in this paper,

More similar in goals to our approach, Henrich et al. (2001) perform behavioral experiments in fifteen small scale societies, all of which are relatively isolated and have had relatively limited market contact. Essentially, their approach exploits the extreme cultural variation between these societies and finds large differences in how they play an ultimatum game. We also hope to shed light on the origins of human economic behavior and the role of environmental experience, but exploit a very different source of variation than Henrich and colleagues. Our experiments can be seen as exploring which aspects of our behavior are not confined to the heavily-socialized human species, but extend to primates that lack any previous market experience. Specifically, if loss-aversion emerged in our evolutionary past we would expect that closely related species would exhibit analogous behavior – and may better understand the origins of our biases by understanding their expression in our close evolutionary neighbors.

2.1 Economic Experiments with Children and Animals

Harbaugh, Krause and Berry also conduct experiments with similar goals to our own, exploiting age instead of cultural or species variation. Harbaugh et al. (2001a) conduct numerous simple budgeting experiments on children between the ages of seven and eleven, and find that violations of GARP are relatively rare. Harbaugh et al. (2001b) in contrast, finds evidence of the endowment effect in children as young as five, and finds no evidence that the effect diminishes with age up through college.⁹ This suggests that the endowment effect is not reduced by market exposure, though leaves open the possibility that children learn this behavior sometime before age five.

While the use of animal subjects is widespread in psychology, their use as subjects in economics is relatively scarce. A notable exception is the work of Kagel, Battalio, Green, and colleagues (Battalio et al., 1981; Battalio et al., 1985; Kagel et al., 1975; Kagel et al., 1981; Kagel et al., 1990; Kagel et al., 1995). These researchers systematically explored a variety of economic decisions (e.g., consumer demand, labor supply, risk aversion, and intertemporal choice) in two classic exemplars of associative learning: rats and pigeons. Having been trained that different levers each delivered a unique reward at an experimentally-variable rate, subjects signaled preferences via their lever choices. Kagel and colleagues then employed a simple revealed preference method in which they examined their subjects' choices when presented with a “budget” of limited lever presses.

since both learned and innate behaviors have neural correlates that may be common across human subjects.

⁹Closely related to loss-aversion, the endowment effect is the observation that consumers often seem to value goods more after possessing them than they do when they do not have them. This is often characterized by a set of people randomly endowed with an object exhibiting a higher willingness to accept (price for selling the good) than the control group's willingness to pay. For a good overview of this bias and its connection to loss aversion, see Kahneman, Knetsch and Thaler (1991).

Most applicable to our work, Kagel et al. (1975), explores how rats and pigeons respond to a compensated price shift. They find that subjects' choices during such a shift largely respected GARP; in fact, utility maximization does a much better job of explaining their data than any other available choice theory (including the canonical non-human psychological choice model, the matching law).¹⁰ In later experiments involving gambles, Kagel and colleagues observed that on balance rats and pigeons obeyed expected-utility theory, but do display some systematic biases. However, unlike results on human (and our capuchin) subjects, Kagel and colleagues find that prospect theory does not explain the deviations from expected-utility theory that are present in rats and pigeons.¹¹

We depart from the important work of Kagel and coauthors in two key ways. First, since rats and pigeons are very distantly related to humans, experiments on them are of limited use in answering questions about high-level human decision making, since most of the relevant neural architecture emerged after our common evolution. The contribution of Kagel and coauthors is more closely akin to that of Becker (1962), demonstrating the robustness of price-theory to large variation in the sophistication of agents.

Second, since rats and pigeons lack the cognitive sophistication of humans, researchers working with these species can only carry out relatively simple choice experiments (i.e., choice between trained levers). These tasks seem unlikely to lead to the classic biases observed in humans, such as framing or reference-point effects. Capuchin monkeys on the other hand, are socially sophisticated organisms whose native ecology requires successful management of scarce resources and risky tradeoffs. This sophistication and their evolutionary proximity to humans make capuchins far better-suited subjects with which to study the mechanisms that enable economic decision-making; yet since our subjects have all been raised in captivity we can limit the possibility that behavior analogous to human behavior evolved socially, in response to similar environments.

3 Subjects: The Tufted Capuchin

The tufted (or brown) capuchin is a New-World monkey native to tropical climates within Venezuela, Columbia, Ecuador, Peru, and Bolivia. A cohabiting capuchin breeding-group is usually characterized by a male-dominance hierarchy. A single alpha-male and several sub-alpha males and females normally live together, with the alpha-male holding sexual monopoly over the females within the group. Capuchins are often referred to as "extractive foragers"; they prefer easy to eat fruit but when pressed are capable of pounding apart

¹⁰For a good summary of the psychological literature on the matching law and its relationship to more modern theories of choice see Herrnstein and Prelec (1991).

¹¹Instead, they find evidence of non-standard probability weighting that is best represented by some mix of fanning out and fanning in (see Kagel, MacDonald & Battalio 1990).

hard nuts, stripping tree bark, raiding bee hives and even killing small vertebrates. For an excellent survey of the species covering all aspects of their native ecology see Fragaszy, Visalberghi & Fedigan (2004).

Besides this ability to explore the origins (and potentially the neural basis) of our preferences and biases, running economic experiments with capuchins carries with it several tradeoffs compared to conventional human subjects. Experiments with human subjects must inevitably assume some independence between the effect being studied in the laboratory and such things as subjects' selection into the subject pool, as well as pre- and post-experimental conditions outside the laboratory. With non-human subjects, we can control selection and directly manipulate various features of their daily environment and social interactions. However, because of the difficulties involved in housing and maintaining a rewarding environment for capuchin subjects, it is prohibitively costly to achieve sample sizes to which economists are accustomed in testing humans. As such, we have chosen a sample size typical for comparative cognition studies with primates (e.g., Brosnan and de Waal, 2003; Brosnan and de Waal, 2004).

3.1 The Experimental Subjects

Our test subjects were all born in captivity and live in a single social group. Five adult capuchins, two male [FL, NN] and three female [MD, HG, JM] ranging from seven to eight years old, participated in this experiment. All were genetically unrelated with the exception of JM [mother of MD]. Individuals were isolated from the rest of the group during each trial in order to minimize the effects of social interaction on experimental performance.

All subjects had previously participated in experiments concerning visual cognition, social cognition, and tool use. For the purposes of this experiment, the most relevant were those studies requiring the ability to trade a token for a food reward, most notably the recent study by Brosnan & de Waal (2003).¹² The scope for contamination was minimal though, as our experiments required the capuchins to learn a considerably more difficult skill than in B&dW. In our studies, capuchins must decide how to intensively allocate a budget of twelve tokens between a pair of options. In contrast the studies of B&dW only provided subjects with one token and allowed only one trading option; subjects decided only whether to trade or not (extensively). Additionally, ours were the first experiments in which these subjects encountered choices that were risky or involved the offering of one reward, but the possibility of receiving another. In other words, all of the variation in choices from

¹²Brosnan & de Waal (2003) use a token trading methodology to study capuchin other-regarding preferences when capuchins observe the results of other conspecifics' trades. They find the fascinating result that capuchins seem less willing to accept a food reward in trade for an otherwise worthless token, if they had recently observed another capuchin receive a better food reward in trade.

which we identify behavioral biases was novel to our subjects.

4 Methods: Setting and Apparatus

In all the following experiments subjects were allowed to trade tokens with one of two experimenters. Each experiment is composed of several sessions, each session constitutes twelve trials, and each trial is an opportunity to trade a token for one of two possible food rewards. Every capuchin was endowed with a budget of tokens at the beginning of each session and was allowed to allocate this budget however they saw fit; however trades had to be conducted one at a time. Identical inch-wide aluminum discs were used as tokens in all exchanges.

Trading was conducted in a cubical testing chamber (28 inches wide) that was adjacent to the main cage, and into which subjects entered voluntarily. Two panels on opposite sides of the chamber allowed participants to interact with experimenters through rectangular openings, large enough for the capuchins to reach out of [and experimenters to reach into] the testing chamber (see figure one). In each trial, two potential trades were offered on opposite sides of the cube, and the subject made its choice between these two options by choosing which experimenter to exchange a token with. All sessions were videotaped in addition to an RA recording each actor's string of choices. The experimental trading protocol is pictured below.



Figure 1: *A capuchin decides which experimenter to trade with. The subject enters the testing chamber (A), takes a token from a tray (B), places it in the hand of an experimenter*

(C), and receives a food reward from a tray in his other hand (D). The film clip from which these are drawn is available from the corresponding author on request.

4.1 General Methods

Before each session of twelve trials, two experimenters (E_1 and E_2 , wearing different colors)¹³ arranged an endowment of tokens on a tray, which was in view but out of reach of the subject. To begin each session an experimenter pushed the tray within reach of the subjects through the front of the testing chamber. Then to begin each trial the two experimenters simultaneously positioned themselves in front of opposite side panels (panel A in figure 1). Each experimenter held a dish with a food reward (in clear view of the subjects) approximately six inches above the opening closest to the interior of the cage, and extended an empty hand into the other panel opening. If the capuchin took a token from the tray (panel B) and placed it in an experimenter’s hand (panel C), then the experimenter would lower his food dish and present the capuchin with the food reward (panel D).

In later conditions the experimenters presented capuchins with risky choices; before lowering the food dish the experimenter would sometimes alter the amount of food, either taking away or adding to the amount of food in the dish. Between each trial the experimenters swapped positions (replenishing the food in their dish if necessary), and resumed their initial stance, with the food reward held several inches above the opening closest to the main cage and an empty hand extended into the cube through the opening nearest the tokens. The session ended after the subject has exchanged all twelve of their budgeted tokens for food rewards. Non-standard trades (including those in which tokens were thrown from the enclosure or those in which multiple tokens were pressed into an experimenter’s hand), were not rewarded and the subject was allowed to make that choice again. So as to minimize subject confusion, each experimenter represented a consistent choice for the capuchin throughout each experiment. No capuchin was allowed to participate in more than two experimental sessions on the same day, but could participate every day if they wished to.

Each subject participated in experiments one, two, and three in sequence, moving from one experiment to the next when their choices in the previous had stabilized. In each of our experiments this criterion was set as five consecutive sessions in which a capuchin allocated their tokens in near-constant proportion.¹⁴ We took these final choices to express each actor’s preferred split between the choices each experiment affords. In our data analysis

¹³For expositional simplicity each experimenter is denoted by E_x where x is how many pieces of food the experimenter would initially display before possibly adding or taking away pieces.

¹⁴This meant their token allocations moved no more than one out of each session’s twelve trials, for five consecutive sessions.

we use only the last five sessions for each actor; each capuchin took between six to twelve sessions to stabilize. Once a subject was finished with an experiment, we transitioned them to the next by running several days worth of "forced trials." These trials were identical to the subject's next experiment, except that only one of his future choices (randomly selected each trial) was available at any given time. In this way the capuchin both became aware that the trading environment had changed, and was "forced" to become equally familiar with both of his new options.

5 Preliminary Experiments: Capuchins Obey Price Theory

Our preliminary experiments closely mirror those of Kagel and coauthors, and allow us to directly test that capuchin choice looks broadly rational and admits standard price theory. In order to do so we first found a set of two goods for each subject between which they were roughly indifferent, then elicited their choice over a simple budget set between these two goods. We then subjected each capuchin to a compensated price shift and examined how they respond.

5.1 Methods: Identifying Preferences

Before beginning the pricing experiment, each participant was tested to identify two food rewards between which the subject was roughly indifferent. That is, starting with apples as the first good, we looked for another good such that when allocating a budget of twelve disks, the capuchin would reliably consume at least some of each good. Each experimenter was assigned a different good to display and exchange for a single token. When a subject reliably consumed a positive quantity of both apples and the other food over at least ten sessions, it was determined that the subject was roughly indifferent between the two goods offered (these foods ended up being either grapes or jello cubes). Until this combination was found, the non-apple experimenter changed the good they offered until this interior budgeting condition was satisfied.

5.1.1 Baseline and Compensated Price Shift

Once an appropriate good was found, the next steps of our price shift experiment are very straightforward. To establish a baseline measurement, each one of three subjects were repeatedly asked to allocate a budget of twelve disks between food one (apples) and food two (either grapes or jello). This was done exactly as described in the general methods above, with each experimenter trading one token for one piece of their respective food reward.

Each capuchin was run on this baseline condition until their choices stabilized; that is

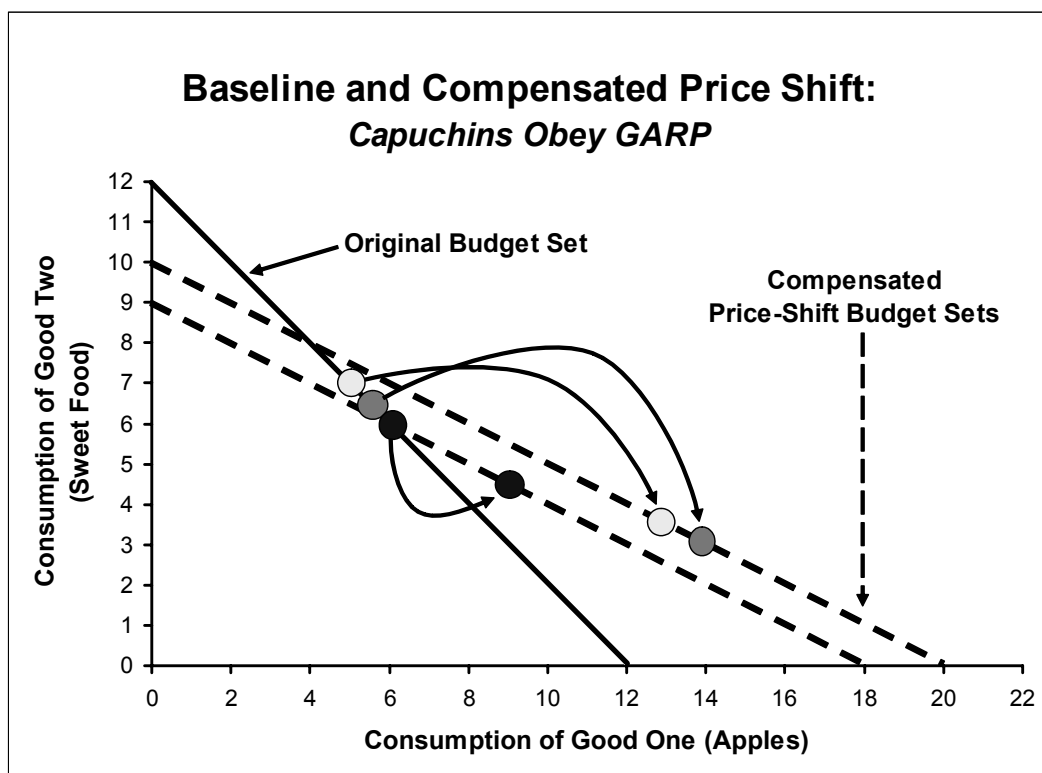
until their choices didn't change by more than one token for a span of five sessions. Once an actor had stabilized, their behavior over the next week was averaged into a representative consumption bundle (see the solid budget line in graph one). Using this bundle, a new budget of disks was assigned to each actor for use in a compensated price shift.

This compensated price shift took the form of the experimenter who trades for apple changing the amount they were willing to trade for a token. Instead of trading a token for one piece of apple, the experimenter would now always display and trade two pieces of apple for each token. This represented a fall in the price of apples by a half, and in order to compensate for this each subject's budget was reduced. This reduction was from twelve to either nine or ten tokens, depending on which most closely shifted back the new budget set such that the bundle the subject originally consumed was close to still lying on the new budget line (see the dotted budget lines in graph one). Each subject's preferences were again allowed to stabilize, then another week's worth of sessions were elicited under this new price régime.

6 Results: Preliminary Price Theory Experiments

The results of the preliminary pricing experiment are summarized in table and graph one. In graph one, the solid line represents the initial budget set each actor was presented in the baseline condition, while the dotted line represents each actor's compensated budget after the price of apples falls in half. In order to satisfy GARP, an actor must consume (weakly) more apples after the shift than before. All subjects choices are aggregated over at least ten sessions, and every actor's choices easily satisfies GARP at the 1% level.

Graph One: *Budget Sets and Subject Choices*



Each point represents the purchasing behavior of a subject after their choices had stabilized, averaged over a week.

Table One: *Price Theory and Compensated Price Shifts*

<i>Baseline Experiment: Subjects given a budget of 12 disks, both goods have price 1.</i>			
	Actor 1 (FL)	Actor 2 (NN)	Actor 3 (AG)
Food used for good one:	Apples	Apples	Apples
Food used for good two:	Jello	Grapes	Jello
Percent of budget spent on good one:	47%	42%	51%
Number of trials:	132	72	144
<i>Compensated Price Shift:</i>			
<i>Good one's price falls from 1 to $\frac{1}{2}$, good two's price stays 1, and the budget shrinks.</i>			
New budget of disks:	10	10	9
Percent of budget spent on good one:	69%	64%	50%
Good one consumed before \rightarrow after shift:	5.6 \rightarrow 13.8	5.0 \rightarrow 12.8	6.1 \rightarrow 9.0
Number of trials:	140	100	216
Choice % responds to price shift: ¹⁵	0.001	0.004	0.966
Change in choice satisfies GARP: ¹⁶	0.001	0.001	0.001

Tests of significance reported as p-values of a two-sided test.

Note though that with only one compensated price shift, satisfying GARP is a weak test. Indeed as Becker (1962) points out, many forms of random behavior can satisfy GARP in response to compensated price changes. Accordingly we apply a more rigorous test as to whether each actor's behavior could arise solely from random behavior. In table one we examine the percent of their budget each capuchin spends on good one both before and after the price shift, and test whether this responds significantly to the change. Two out of three subjects showed a significant response to the compensated price change at the 1% level in this more stringent test.

7 Main Experiments: Are Capuchins Reference-Dependant?

Once our original three subjects completed these initial experiments, an additional two subjects were recruited for our main set of experiments. In this set of experiments the same budgeting procedure was used to elicit choices, with each session composed of twelve opportunities to trade a token for one of two possible food rewards. Unlike the initial experiment though, only apples were used as food reward, and the experimenters no longer

¹⁵This is a two-sided p-test that the *fraction of trials* the subject chooses apples responds to the price shift; NN and FL change significantly while AG does not.

¹⁶Our test of GARP is a two-sided p-test that the number of pieces of apple that the actor consumes weakly increases after the compensated price-shift.

automatically presented the capuchin with the apples displayed in their tray when given a token. Now, experimenters sometime altered the food in the presentation tray before making that tray available to the subject. In this way, we were able to independently vary what the capuchin was initially shown and what the capuchin would receive in exchange for a token, with the latter sometimes consisting of a gamble.

7.1 Methods: Experiment One, Stochastic Dominance

In experiment one, a capuchin could trade their tokens with one of two experimenters. Experimenter E_2 represented a random payoff of one or two apple pieces each with equal probability, and experimenter E_1 represented a sure payoff of one piece. E_1 and E_2 also differed in how many pieces they initially showed the capuchin; E_2 displayed two squares of apple, while E_1 displayed only one square of apple. This experiment tests whether capuchin choice respects first-order stochastic dominance; that is if they prefer gambles that weakly dominate another option.

Specifically, after being given a token experimenter E_1 always lowered his dish to present the subject with one apple piece – exactly as many as he had displayed. In contrast E_2 started every trial displaying two apple pieces in her tray, but would only deliver both pieces half the time she was traded with. The other half of the time E_2 would remove one of her two apple pieces and deliver only the remaining piece to the subject. A random-number generator determined beforehand whether any given trade would result in a payoff of two or one; when a apple piece was removed it was placed into an opaque receptacle underneath the testing table that was both out-of-sight and out-of-reach of the subjects.

7.2 Methods: Experiment Two, Reference-Dependence

In experiment two, subjects chose between experimenters who both delivered identical gambles – differing only in whether they added to or subtracted from their initial displayed offering of one or two apple pieces. This was designed to test whether capuchins would respond to a simple framing manipulation, presenting some payoffs as gains and some as losses, while holding constant the underlying payoffs.

Specifically, E_1 and E_2 would stand on opposite sides of the testing chamber displaying one and two apple pieces, respectively. Upon being presented with a token, E_2 would present the subject with either the two apple pieces he had displayed, or would visibly remove one piece and deliver only the remaining apple piece to the subject. When an apple piece was removed, it was placed into an opaque receptacle underneath the testing table that was out-of-sight and out-of-reach of the subject.

When the subject traded with E_1 however, she would either present the single piece she

displayed, or add one apple piece and deliver two pieces. When this bonus piece was added it was drawn from an identical receptacle.

Essentially then, both experimenters represented a fifty-fifty lottery of one or two apple pieces. They differed only in whether they initially had displayed one or two apples, framing for the marginal apple piece as either a gain or a loss. A random-number generator determined beforehand whether any given trade would result in a payoff of two or one; we call these the bonus vs. penalties conditions.

7.3 Methods: Experiment Three, Loss-Aversion

In experiment three, subjects chose between experimenters who both delivered a payoff of one apple piece – differing only in whether they initially displayed one or two pieces. This experiment was designed to test for the presence of reference-effects in riskless situations, and when combined with experiment two, allows us to measure loss-aversion.

Specifically, E_1 and E_2 would stand on opposite sides of the testing chamber, displaying one and two apple pieces, respectively. Upon being presented a token, E_2 always removed one apple piece and delivered the remaining piece to the subject. The removed square was placed into the opaque receptacle underneath the testing table. In contrast if the subject traded with E_1 , she always presented the single square she displayed.

Essentially then trading with either E_1 or E_2 delivered identically payoffs. However, on all trades E_1 gave exactly the quantity of apple he displayed, while E_2 displayed a quantity of apple that was always reduced from two to one before it was made available to the subject.

8 Non-Parametric Results: Main Experiments

The results of all our experiments are reported below in tables two through four, broken down by subject. We will first discuss what can be learned from our results without imposing any significant parametric assumptions; in the next section we fit a simplified version of the standard prospect-theoretic utility function to our subjects which allows more precise analysis.

The results of experiment one are summarized below. Table two shows how each subject behaved over five sessions (60 trials), after an initial set of sessions in which their choices stabilized as they learned about the experimental choices.

Table Two: *Experiment One, Gambles and Stochastic Dominance*

E_1 shows 1 and gives 1, E_2 shows 2, then gives 1 or 2 with prob. $1/2$.					
Subject (name):	1 (FL)	2 (HG)	3 (JM)	4 (MD)	5 (NN)
Percent of trials E_1 chosen: [†]	10%	18%	12%	22%	5%
Sessions till stable:	7	6	7	6	11

[†] All subject's choices different from 50% at the 1% level in a two-sided p -test.

The capuchins express a clear preference (87% of trades) for E_2 , the experimenter who displays two apple pieces and delivers either one or two pieces with equal probability. This is of course not surprising; the second option stochastically dominates the first and gives on average a half-piece more of apple.

Given this result though, the results of experiment two are quite surprising. All subjects left experiment one conditioned to favor E_2 , the experimenter who displays two pieces of food. Despite this, in experiment two the capuchins quickly reverse this preference and trade much more with the experimenter who displays only one piece of food. Table three summarizes these results.

Table Three: *Experiment Two, Reference-Dependence in Gambles*

E_1 shows 1 and E_2 shows 2, then both give 1 or 2 with prob. $1/2$.					
Subject (name):	1 (FL)	2 (HG)	3 (JM)	4 (MD)	5 (NN)
Percent of trials E_1 chosen: [†]	68%	70%	70%	70%	78%
Sessions till stable:	11	9	9	9	13

[†] All subject's choices different from 50% at the 1% level in a two-sided p -test.

Contrary to both their conditioning and the intuition that naive subjects would favor greater initial displays of food (experimenter E_2), capuchins express a preference for E_1 , the experimenter who frames the gamble as a 50% chance of a bonus rather than a 50% chance of a loss. Pooled, subjects traded with E_1 in 71% of trials in their last five sessions (again, measured after each subject's choices stabilized). For all five subjects this change was significantly different not just from experiment one but from random (50-50) behavior.

Note that any theory of choice which does not take into account reference-dependence fails to predict this pattern of behavior. Indeed, since our experimenters switch sides of the testing chamber between each trial, in order to express a preference between E_1 and E_2 a capuchin has to actively follow their preferred experimenter from side to side, expending both time and attention.

Experiment three shows this effect is not confined to risky choices, and when combined with experiment two suggest that capuchins are not just reference-dependant, but loss-averse. The results are summarized in table four below.

Table Four: *Experiment Three, Riskless Reference-Dependence*

<i>E</i> ₁ shows 1 and <i>E</i> ₂ shows 2, then both give 1.					
Subject (name):	1 (FL)	2 (HG)	3 (JM)	4 (MD)	5 (NN)
Percent of trials <i>E</i> ₁ chosen:†	73%	75%	80%	82%	87%
Choice % greater than in exp. 2, p-value:	0.27	0.27	0.10	0.07	0.11
Pooled, choice % greater than in exp. 2:	<i>p</i> < 0.023, <i>two-sided p-test</i> .				
Sessions till stable:	9	10	10	8	10

† *All subject's choices different from 50% at the 1% level in a two-sided p-test.*

Subjects strongly preferred experimenter *E*₁ over experimenter *E*₂ (who initially displayed one and two pieces of apple, respectively), despite the fact that both always provide the same, sure payoff of one apple piece. For all subjects this preference (% of trials trading with *E*₁) was stronger than in experiment two, suggesting this was not due to conditioning from the previous experiment. Since the only difference between the two experimenters was that *E*₂ showed more than he eventually gave, these results suggest our capuchins are reference-dependant even in riskless choice settings.

Also, note that in experiment two subjects chose between *E*₂ who gave a *half-chance of a loss* and *E*₁ who gave a *half-chance of a gain*, both of these forces pushing the subject to chose *E*₁. In experiment three however, *E*₁ always gives exactly what they showed while *E*₂ delivers a *sure loss*. An interpretation of the fact that subjects show a stronger preference in experiment three than two (79% to 71%), is that a sure loss has a stronger effect than the combined effects of a half-loss and a half-gain, or that losses affect subject's choices more than gains. In other words:

$$|loss| > \frac{1}{2} |loss| + \frac{1}{2} |gain| \Leftrightarrow |loss| > |gain|$$

Thus, the pooled p-test of experiment three being stronger than experiment two (*p* < 0.023 in a two-sided p-test) can be taken to confirm the presence of loss-aversion in capuchin choice.

We now turn to a simple utility function and choice model which will allow us to test this more formally and to provide a point estimate of the ratio of the effects of losses to that of gains, a number often called the coefficient of loss-aversion. We then compare our results to estimates derived from human behavior.

9 Parametric Results: A Simple Choice Model

In order to more formally organize these results we now introduce a simple choice model which nests both loss-aversion and reference-dependence in a more general model of utility. This allows us to estimate parameters such as the capuchins' degree of loss-aversion, and compare these parameters to analogous human estimates from both lab and field data. Our model is a linear simplification of the utility function found in Koszegi & Rabin (2004) in which the reference point when considering a gamble is what the experimenter initially displays to the subject.

Notation 1 *We'll write a gamble as a three-tuple (r, a, b) , where the experimenter displays r then gives the subject payoffs a or b with equal probability. Note that if $a = b$ the experimenter always gives the subject the same payoff and the offered trade involves a sure bet.*

So for example, if the experimenter displayed one piece of apple and gives the capuchin either one or two pieces (each with equal probability), we would write that experimenter's offered trade as $X = (1, 1, 2)$. Note that since all probabilities are 50-50, the expected value (EV) of any gamble (r, a, b) is just $\frac{1}{2}(a + b)$.

Formally, we assume that a subject S's utility for a simple gamble $G = (r, a, b)$, is given by:

$$U_S(G) = \alpha \cdot (\text{losses from ref. point}) + \beta \cdot (\text{gains from ref. point}) + \gamma \cdot (\text{exp. value}) + \varepsilon_S,$$

or

$$U_S(G) = \alpha \cdot (r - a) + \beta \cdot (b - r) + \gamma \cdot EV(G) + \varepsilon_S$$

That is, we assume that a capuchin's utility is a simple linear function over losses and gains (with respect to what was originally displayed) and the expected value of the gamble. This utility representation imbeds both expected-value maximization and complete reference-dependence, allowing us to estimate both the degree loss-aversion, and its relative strength. This is, we can test not only the relative strength of losses and gains, but also how important these reference-effects are with respect to changes in expected food rewards. Completing the identifying model, we assume that agents choose the gamble with the highest expected utility given a random shock ε for each gamble where ε is drawn each trial (we allow for ε to be correlated within subject.)

Table five below summarizes how each gamble falls with respect to expected gains, losses, and food rewards.

Table Five: *Expected Gains, Losses, and Values for each Experimental Choice*

Experiment:	One		Two		Three	
Experimenter:	E₁	E₂	E₁	E₂	E₁	E₂
Gamble offered:	(1, 1, 1)	(2, 1, 2)	(1, 1, 2)	(2, 1, 2)	(1, 1, 1)	(2, 1, 1)
Gains:	0	0	$\frac{1}{2}$	0	0	0
Losses:	0	$\frac{1}{2}$	0	$\frac{1}{2}$	0	1
Expected Value:	1	$1\frac{1}{2}$	$1\frac{1}{2}$	$1\frac{1}{2}$	1	1
% of trials chosen:	13%	87%	71%	29%	79%	21%

Table constructed pooling all subjects' last five sessions after choices stabilize (60 trials).

Using these values we can fit our simple choice model to the data using standard regression techniques. The objects of interest are the weights subjects place on losses, gains, and expected food rewards: our utility parameters α , β , and γ . We are also interested in the ratio of the slope of the utility function over losses and gains, $\frac{-\alpha}{\beta}$. This measure of the "kink" in a subject's utility function is commonly used to quantify loss-aversion in human behavior, and is commonly called the coefficient of loss-aversion.

9.1 Regression Analysis: Parameter Values

In order to provide point estimates for our parameters α , β , and γ , we fit our choice model to the observed behavior of our capuchins over all three experiments, pooling across all five subjects but allowing ε to be correlated within subject. We also use only the last five sessions for each subject in each experiment so as to study only the latter, stable choices of our subjects.

To fit our model we assume that agents maximize utility subject to the random ε utility shocks each option draws on each trial. This can be accomplished with standard regression techniques under two distributional assumptions on the ε term. If ε is normally distributed this reduces to a simple Probit regression; if ε is distributed type-one extreme value this reduces to a Logit regression. Specifically, if the subject is choosing between gambles $G_1 = (r_1, a_1, b_1)$ and $G_2 = (r_2, a_2, b_2)$, then we estimate the following equation:

$$\Pr(G_1 | \{G_1, G_2\}) = \Phi(\alpha \cdot \Delta_{LOSSES} + \beta \cdot \Delta_{GAINS} + \gamma \cdot \Delta_{EV}), \text{ or}$$

$$\Pr(G_1 | \{G_1, G_2\}) = \Phi(\alpha \cdot [(r_1 - a_1) - (r_2 - a_2)] + \beta \cdot [(b_1 - r_1) - (b_2 - r_2)] + \gamma \cdot [EV(G_1) - EV(G_2)])$$

where Φ is the CDF of the appropriate distribution. Note that we specify this equation without a constant term; this restriction has the natural interpretation that if two gambles

are identical in every way, the capuchin chooses each option with equal probability. The results of both regressions are reported in table six below.

Table Six: *Estimating the Choice Model*

	Probit		Logit	
	Coef.	Std.Err.	Coef.	Std.Err.
$\Delta_{LOSSES} (\alpha)$:	-0.818	0.084	-1.345	0.146
$\Delta_{GAINS} (\beta)$:	0.308	0.058	0.478	0.096
$\Delta_{EV} (\gamma)$:	3.040	0.313	5.089	0.146
# of observations:	900		900	

Regressions clustered by subject, all coefficients significant at the 1% level.

In both specifications the coefficients on losses and gains are of similar magnitudes, are of the sign predicted by reference-dependence, and are significant at the 1% level. In both regressions the coefficient on *EV* is also roughly five times the magnitude of those on gains and losses. This is perhaps not that surprising, given that our treatment only looks at the long-run behavior of our subjects after facing the same choice many times and our reference-point treatment is a relatively mild framing intervention.

In much of the work on human loss-aversion, the ratio of the coefficients on losses and gains is commonly used as a measure of loss-aversion, representing how "kinked" the utility function is at the reference point. That is, the ratio $\frac{-\alpha}{\beta}$ (commonly called the coefficient of loss-aversion) measures how much more sensitive subjects are to losses than gains.

We estimate this ratio under both specifications; table seven reports both sets of results.

Table Seven: *Estimating the Coefficient of Loss-Aversion*

	Probit		Logit	
	Coef.	Std.Err.	Coef.	Std.Err.
coefficient of loss-aversion, $-\alpha/\beta$:	2.654	0.551	2.813	0.642
test if $-\alpha/\beta = 1$:	$p < 0.003$		$p < 0.005$	
test if $-\alpha/\beta = 2.5$:	$p < 0.780$		$p < 0.625$	
# of observations:	900		900	

All test statistics computed using the delta method and clustering by subject.

Across both specifications the ratio of losses to gains is roughly 2.7 to 1, meaning our capuchin subjects seem weigh losses more than twice as heavily as gains. While these estimates are parametric and should be interpreted with caution, in magnitude our estimates are strikingly similar to measures of human loss-aversion in both lab and field data.

Tversky & Kahneman (1991) summarize a large body of survey evidence on minimally acceptable gambles and find a ratio of roughly 2.5 to 1. Average ratios of willingness to pay to willingness to accept found in most endowment-effect experiments (see for example Kahneman, Knetsch, & Thaler 1990) yield a ratio of around 2.7 to 1. Bernartzi and Thaler (1995) calibrates a ratio of approximately 2.3 from the aggregate risk preferences of stock investors.

10 Conclusion

Our results indicate that loss-averse behavior is not confined to humans but is present in our closest evolutionary neighbors and is most likely the result of an evolutionarily-ancient and common behavioral mechanism. As field evidence increasingly suggests that reference-dependant and loss-averse behavior is widely present in the economy, the origins of these biases may influence how we incorporate them into an adequate model of individual decision making. If these biases are an innate feature of how our brains represent choices, we may be more inclined to believe they will be broadly present in both common and novel settings, stable across time and cultures, and resistant to elimination by either market exposure or education.

Our paper also suggests the utility of methodological exchanges between economics and the closely related behavioral sciences. Bringing the analytic framework of revealed preference to bear on questions at the intersection of economics, biology and psychology carries the possibility of insights useful to each.

References

- [1] Afriat, Sidney N.; "The Construction of Utility Functions from Expenditure Data,"
International Economic Review, February 1967, 8(1), pp. 67-77.
- [2] Battalio, Raymond C; Green, Leonard; Kagel, John H.; "Income-Leisure Tradeoffs of
Animal Workers"
The American Economic Review, Vol. 71, No. 4. (Sep., 1981), pp. 621-632.
- [3] Battalio, Raymond C.; Kagel, John H.; MacDonald, Don N.; "Animals' Choices over
Uncertain Outcomes: Some Initial Experimental Results"
The American Economic Review, Vol. 75, No. 4. (Sep., 1985), pp. 597-613.
- [4] Battalio, Raymond C.; Kagel, John H.; Rachlin, Howard; Green, Leonard;
"Commodity-Choice Behavior with Pigeons as Subjects"
The Journal of Political Economy, Vol. 89, No. 1. (Feb., 1981), pp. 67-91.
- [5] Becker, Gary S.; "The Economic Approach to Human Behavior"
Chicago: University of Chicago Press, 1976
- [6] Becker, Gary S.; "Irrational Behavior and Economic Theory"
The Journal of Political Economy, Vol. 70, No. 1. (Feb., 1962), pp. 1-13.
- [7] Benartzi, Shlomo; Thaler, Richard H.; "Myopic Loss Aversion and the Equity Premium
Puzzle"
The Quarterly Journal of Economics, Vol. 110, No. 1 (Feb., 1995), pp. 73-92.
- [8] Brosnan, Sarah F.; de Waal, Frans B. M.; "Monkeys Reject Unequal Pay"
Nature 425, 297 - 299 (18 September 2003)
- [9] Brosnan, Sarah F.; de Waal, Frans B. M.; "Socially Learned Preferences for Differen-
tially Rewarded Tokens in the Brown Capuchin Monkey (*Cebus apella*)"
Journal of Comparative Psychology 2004, Vol. 118, No. 2, 133-139
- [10] Camerer, C. F. (2000). Prospect Theory in the Wild: Evidence from the Field
In Kahneman & Tversky (Ed.), Choices, Values, and Frames (pp. 288-300). Cambridge:
Cambridge University Press.
- [11] Fragaszy, Visalberghi & Fedigan; The Complete Capuchin
Cambridge: Cambridge University Press 2004.
- [12] Genesove, David; Mayer, Christopher; "Loss Aversion And Seller Behavior: Evidence
From The Housing Market"

- The Quarterly Journal of Economics, MIT Press, vol. 116(4), (Nov., 2001), pp. 1233-1260.
- [13] Harbaugh, William T.; Krause, Kate; Berry, Timothy R.; "GARP for Kids: On the Development of Rational Choice Behavior"
The American Economic Review, Vol. 91, No. 5. (Dec., 2001), pp. 1539-1545.
- [14] Harbaugh, William T.; Krause, Kate; Vesterlund, Lise; "Are Adults Better Behaved than Children? Age, Experience, and the Endowment Effect"
Economics Letters 70 (2001) pp. 175–181.
- [15] Hardie, Bruce G. S.; Johnson, Eric J.; Fader, Peter S.; "Modeling Loss Aversion and Reference Dependence Effects on Brand Choice"
Marketing Science, Vol. 12, No. 4. (Autumn, 1993), pp. 378-394.
- [16] Henrich, Joseph; Boyd, Robert; Bowles, Samuel; Camerer, Colin; Fehr, Ernst; Gintis, Herbert; McElreath, Richard; "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies"
The American Economic Review, Papers and Proceedings Vol. 91, No. 2, (May, 2001), pp. 73-78.
- [17] Herrnstein, Richard J.; Prelec, Drazen; "Melioration: A Theory of Distributed Choice"
The Journal of Economic Perspectives, Vol. 5, No. 3. (Summer, 1991), pp. 137-156.
- [18] Houthakker, Hendrik S.; "Revealed Preference And The Utility Function,"
Economica, May 1950, 17(66), pp. 159-74.
- [19] Kahneman, Daniel; Knetsch, Jack L.; Thaler, Richard H.; "Experimental Tests of the Endowment Effect and the Coase Theorem"
The Journal of Political Economy, Vol. 98, No. 6. (Dec., 1990), pp. 1325-1348.
- [20] Kahneman, Daniel; Knetsch, Jack L.; Thaler, Richard H.; "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias"
The Journal of Economic Perspectives, Vol. 5, No. 1.(Winter, 1991), pp. 193-206.
- [21] Kahneman, Daniel; Tversky, Amos; "Prospect Theory: An Analysis of Decision under Risk"
Econometrica, Vol. 47, No. 2. (Mar., 1979), pp. 263-292.
- [22] Kagel, John H.; Battalio, Raymond C.; Rachlin, Howard; "Demand Curves for Animal Consumers"
The Quarterly Journal of Economics, Vol. 96, No. 1. (Feb., 1981), pp. 1-16.

- [23] Kagel, John H.; MacDonald, Don N.; Battalio, Raymond C.; Rachlin, Howard; Green, Leonard; Basmann, R.L.; and Klemm, W.R.; "Experimental Studies of Consumer Demand Behavior Using Laboratory Animals"
Economic Inquiry, Oxford University Press, 1975, vol. 13(1), pages 22-38.
- [24] Kagel, John H.; Battalio, Raymond C.; Green, Leonard; "Economic Choice Theory: An Experimental Analysis of Animal Behavior"
Cambridge: Cambridge University Press 1995.
- [25] Kagel, John H.; MacDonald, Don N.; Battalio, Raymond C.; "Tests of 'Fanning Out' of Indifference Curves: Results from Animal and Human Experiments"
The American Economic Review, Vol. 80, No. 4. (Sep., 1990), pp. 912-921.
- [26] Koszegi, Botond; Rabin, Matthew; "A Model of Reference-Dependent Preferences"
Working Paper; 2004, Department of Economics, University of California Berkeley
- [27] Laibson, David; "Golden Eggs and Hyperbolic Discounting"
The Quarterly Journal of Economics, Vol. 112, No. 2, In Memory of Amos Tversky (1937-1996). (May, 1997), pp. 443-477.
- [28] McClure, Samuel M.; Laibson, David I.; Loewenstein, George; Cohen, Jonathan D.; "Separate Neural Systems Value Immediate and Delayed Monetary Rewards"
Science 306, (15 October 2004)
- [29] Odean, Terrance; "Are Investors Reluctant to Realize Their Losses?"
The Journal of Finance, Vol. 53, No. 5. (Oct., 1998), pp. 1775-1798.
- [30] Rayo, Luis; Becker, Gary; "Evolutionary Efficiency and Happiness"
University of Chicago Working Paper, 2005
- [31] Samuelson, Paul. A. "A Note On The Pure Theory Of Consumer's Behavior,"
Economica, February 1938, 5(17), pp. 61-71.
- [32] Schneider, Horácio; Canavez, Flávio C.; Sampaio, Iracilda; Moreira, Miguel Ângelo M.; Tagliaro, Claudia H.; Seuánez, Héctor N.; "Can molecular data place each neotropical monkey in its own branch?"
Chromosoma, Volume 109, Issue 8, Feb 2001, Pages 515 - 523
- [33] Tomasello, M.; Call, J.; Primate Cognition.
Oxford University Press, New York, 1997

- [34] Tversky, Amos; Kahneman, Daniel; "Loss Aversion in Riskless Choice: A Reference-Dependent Model"
The Quarterly Journal of Economics, Vol. 106, No. 4. (Nov., 1991), pp. 1039-1061.
- [35] Varian, Hal R. "The Nonparametric Approach To Demand Analysis,"
Econometrica, July 1982, 50(4), pp. 945-73.
- [36] Visalberghi, Elisabetta; "Capuchin monkeys: A window into tool use in apes and humans"
In: Gibson, Kathleen Rita / Ingold, Tim (eds): Tools, language and cognition in human evolution. Cambridge : Cambridge Univ. Pr. (1993)
- [37] Wolfe, John Bascom; Effectiveness of token-rewards for chimpanzees
Baltimore: John Hopkins Press 1936.