How Basic Are Behavioral Biases? Evidence from Capuchin Monkey Trading Behavior
Author(s): M. Keith Chen, Venkat Lakshminarayanan, and Laurie R. Santos
Source: *Journal of Political Economy*, Vol. 114, No. 3 (June 2006), pp. 517-537
Published by: The University of Chicago Press
Stable URL: http://www.jstor.org/stable/10.1086/503550
Accessed: 02/07/2015 08:58

How Basic Are Behavioral Biases? Evidence from Capuchin Monkey Trading Behavior

M. Keith Chen

Yale University and Cowles Foundation

Venkat Lakshminarayanan and Laurie R. Santos

Yale University

Behavioral economics has demonstrated systematic decision-making biases in both lab and field data. Do these biases extend across contexts, cultures, or even species? We investigate this question by introducing fiat currency and trade to a colony of capuchin monkeys and recovering their preferences over a range of goods and gambles. We show that capuchins react rationally to both price and wealth shocks but display several hallmark biases when faced with gambles, including reference dependence and loss aversion. Given our capuchins' inexperience with money and trade, these results suggest that loss aversion extends beyond humans and may be innate rather than learned.

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

We would like to acknowledge financial support from Whitebox Advisors and the Russell Sage Foundation, 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, Steve Levitt, Barry Nalebuff, Emily Oster, Jesse Shapiro, Peter Schott, and seminar participants at Berkeley, Brown, University of California at San Diego, Caltech, the University of Chicago, the Stanford Institute for Theoretical Economics, and Yale for their insightful comments and suggestions. These data and additional materials are available at http://www.som.yale.edu/Faculty/keith.chen/.

[*Journal of Political Economy*, 2006, vol. 114, no. 3]

> another, this is mine, that yours; I am willing to give this for that. (Adam Smith, *Wealth of Nations*)

I. 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, have received a substantial amount of empirical attention, from both economics and neighboring disciplines such as psychology and sociology.[1] Evidence that agents treat losses differently than comparable gains has been found in experimental markets as the endowment effect[2] (Kahneman, Knetsch, and Thaler 1990), in the trades of individual investors who are reluctant to realize losses (Odean 1998), and in the behavior of house sellers who are unwilling to sell below buying price (Genesove and Mayer 2001).

Despite the mounting evidence of the importance of this behavior, less direct attention has been paid to understanding how basic or widespread these biases are. Are biases such as loss aversion the result of social or cultural learning and specific environmental experiences? Or could they be more universal, perhaps resulting from mechanisms that arise regardless of context or experience? The root cause of a behavioral bias may affect how we think about both its potential scope and the degree to which we believe that market incentives will act to reduce its 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, 5) writes that "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."[3] 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 preferences were so deeply rooted as to extend to closely related species, this would bolster the assumption of preference stability.

[1] These 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).

[2] The endowment effect is the observation that the minimum amount subjects are willing to accept to give up a randomly endowed good is often more than twice what they would have been willing to pay for that good had they not been given it.

[3] It 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" (1976, 5).

Indeed, early experimental work found support for the stability of preferences and the applicability of economic choice theory far from its usual subjects; several studies have demonstrated that rat and pigeon behavior seems to obey the laws of demand (Kagel et al. 1975). Unfortunately, while rats and pigeons are easy subjects to work with, their limited cognitive abilities make it difficult to investigate more subtle aspects of economic choice, including many important and systematic human biases.

In this study we test for both adherence to the law of demand and the presence of reference-dependent and loss-averse choice in a sophisticated and closely related primate, the tufted capuchin monkey (*Cebus apella*). To do this, we introduce a fiat currency to a colony of capuchin monkeys, teaching them that small coin-like disks can be traded with human experimenters for food rewards[4] and are fungible across a variety of possible trades. Using this new ability, we are able to conduct a number of revealed-preference experiments analogous to canonical human choice experiments.

Our first set of experiments investigate the purchasing behavior of capuchins when they are asked to allocate a budget of tokens among a set of possible foods. In response to both price and wealth shocks, capuchins adjust their purchasing behavior in ways consistent with the generalized axiom of revealed preferences (GARP).[5] In this way, capuchins' choice closely mirrors our own and admits the standard tools of utility analysis and price theory. This closely mirrors the early economic work on rats and pigeons and provides 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. In these experiments, a capuchin must choose between trades in which one amount of food is initially displayed (serving as a frame) but another amount may be delivered. Specifically, in our main experiments, capuchins express a strong preference for gambles in which good outcomes are framed as bonuses (the subject sometimes receives more than was initially displayed) rather than payoff-identical gambles in which bad outcomes are framed as losses.

[4] For similar trading methodologies with capuchins, see Westergaard, Liv, Chavanne, and Suomi (1998), Liv, Westergaard, and Suomi (1999), Brosnan and de Waal (2003, 2004), and Westergaard et al. (2004); for a very early example of primates' capacity to trade, see Wolfe (1936).

[5] Early 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 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 nonsatiated utility function.

Furthermore, capuchins seem to weigh those losses more heavily than comparable gains, displaying not just reference dependence but loss aversion. These experiments also allow us to reject most competing models of naive or unsophisticated choice. In particular, several of our results require a capuchin to overcome the impulse to try to obtain an initially larger food reward and instead to trade with an experimenter who initially displays a smaller food reward. Numerous studies have shown that this type of inhibition problem is difficult for both monkeys and even great apes.[6]

Arguing that loss aversion is not an acquired characteristic of capuchin preferences 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; subjects had no prior trading experience that could have led to the development of these biases. As such, our results suggest that loss-averse behavior is a very general feature of economic choice and extends to some of our closely related neighbors. A large body of research in biology suggests that since humans and capuchins are closely related, any shared cognitive systems are likely to have a common origin. In light of this, our results may suggest that loss aversion is an innate and evolutionarily ancient feature of human preferences, a function of decision-making systems that evolved before the common ancestors of capuchins and humans diverged.

The remainder of the paper is organized as follows. Section II reviews the set of papers in behavioral economics that speak to the origin and scope of behavioral biases. Section III describes capuchins both as a species and as our test subjects, and lays out some of the trade-offs inherent in experimenting on nonhuman primates. Section IV provides details on our laboratory setup and our method for eliciting preferences from capuchin trading behavior. Section V describes our initial compensated price shift experiments and presents those results. Section VI describes our method for inducing gain/loss frames and the setup of our loss-aversion experiments. Section VII presents the results of these experiments, which we discuss before presenting conclusions in Section VIII.

[6] Both new-world monkeys and great apes fail to solve 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) (Boysen and Berntson 1995; Kralik, Hauser, and Zimlicki 2002). However, when the game is modified such that the large reward is presented whenever the agent reaches for a *symbol* of the small reward (and the small reward is presented whenever the agent reaches for a symbol of the large reward), many primates succeed (e.g., Boysen, Mukobi, and Berntson 1999). 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 symbol-mediated choice does.

II. Related Literature

Several recent papers have shed light on the foundations 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 et al. (2004) show that 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.

A few theory papers have also tried to understand what types of evolutionary forces might have selected for organisms that display behavioral biases. These papers typically model preferences as a means by which nature incentivizes an organism to maximize its evolutionary fitness; various constraints on nature's ability to achieve first-best incentives give rise to behavioral biases. Most notably, Rayo and Becker (2005) explore what types of evolutionary pressures would have produced both past payoff and social reference point effects, and Samuelson and Swinkels (2006) explore under what conditions choice set effects might be evolutionarily optimal.

More similar in goals to our approach, Henrich et al. (2001) perform behavioral experiments in 15 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 his 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.

Economic experiments with children and animals.—Harbaugh, Krause, and Berry (2001) also conduct experiments with goals similar to our own, exploiting age instead of cultural or species variation. They conduct numerous simple budgeting experiments on children between the ages of 7 and 11 and find that violations of GARP are relatively rare. Harbaugh, Krause, and Vesterlund (2001), in contrast, find evidence of the endowment effect in children as young as 5 and find no evidence that the effect diminishes with age up through college.[7] This suggests that the endowment effect is not reduced by market exposure, though it leaves open the possibility that children learn this behavior through experience they receive before age 5.

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 their colleagues (Kagel et al. 1975; Battalio, Green, and Kagel 1981; Battalio, Kagel, Rachlin, and Green 1981; Kagel, Battalio, Rachlin, and Green 1981; Battalio, Kagel, and MacDonald 1985; Kagel, MacDonald, and Battalio 1990; Kagel, Battalio, and Green 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 his colleagues then employed a simple revealed preference method in which they examined their subjects' choices when presented with a "budget" of limited lever presses.

Most applicable to our work, Kagel et al. (1975) explore 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 nonhuman psychological choice model, the matching law).[8] In later experiments involving gambles, Kagel and his colleagues observed that, on balance, rats and pigeons obeyed expected utility theory but do display some systematic biases. However, in contrast to results on human (and our capuchin) subjects, Kagel and his colleagues find that prospect theory does not explain the deviations from expected utility theory that are present in rats and pigeons.[9]

We depart from the important work of Kagel and his coauthors in

[7] Closely 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 (the 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 et al. (1991).

[8] For 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).

[9] Instead, they find evidence of nonstandard probability weighting that is best represented by some mix of fanning out and fanning in (see Kagel, MacDonald, and Battalio 1990).

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 presumably emerged after our common evolution. The contribution of Kagel and his 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 carry out only 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 trade-offs. 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 developed in response to similar environments.

III. Subjects: The Tufted Capuchin

The tufted (or brown) capuchin is a new-world monkey native to tropical climates within South America. 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 hard nuts, stripping tree bark, raiding beehives, and even killing small vertebrates. For an excellent survey of the species covering all aspects of their native ecology, see Fragaszy, Visalberghi, and Fedigan (2004).

As a species, tufted capuchin monkeys (*Cebus apella*) have been widely studied in psychology and anthropology. They make excellent subjects since they are relatively quick and adept problem solvers, skilled tool users, and a close evolutionary neighbor to humans.[10]

[10] Within 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, 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 that 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).

Despite this history, conducting economic experiments with capuchins carries with it several trade-offs 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 postexperimental conditions outside the laboratory. With nonhuman 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, 2004).

The experimental subjects.—Our test subjects were all born in captivity and live in a single social group. Six adult capuchins, three male (AG, FL, and NN) and three female (MD, HG, and JM) ranging from 7 to 8 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 but had never before participated in a study involving trade. In all our experiments we used sweet foods, which are highly valued by capuchins, as food rewards. Outside our experiments all subjects have ample access to water and calories (in the form of fruits and monkey chow); what motivates our capuchins is payment of desirable foods, not hunger per se.

IV. 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 12 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 it 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 fig. 1). In each trial, two potential

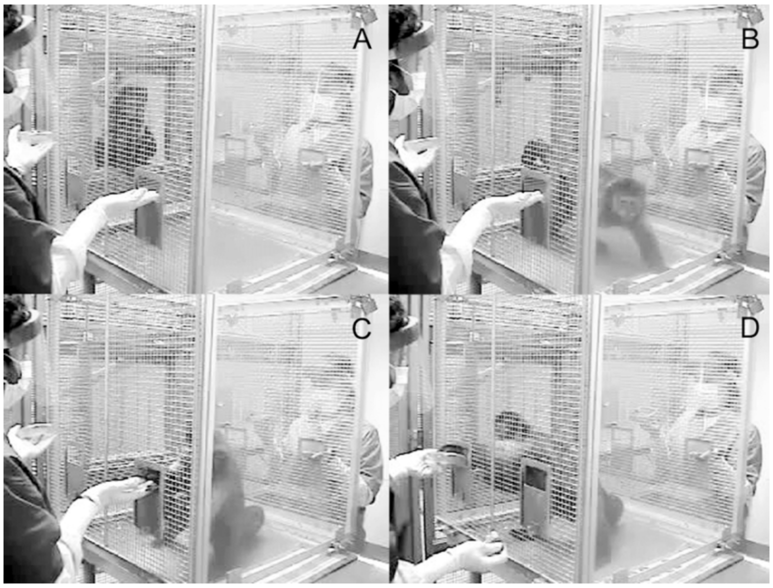

FIG. 1.—A capuchin must decide how to spend a budget of coins. The tray at the front of the testing chamber holds the monkey's budget, and each of the two experimenters displays in one hand a food reward in a small tray; his other hand is empty and outstretched. The subject enters the testing chamber (frame A), takes a token from the tray (frame B), places it in the hand of an experimenter (frame C), and receives a food reward from a tray in his other hand (frame D). The film clip from which these pictures are drawn is available from Chen on request.

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, a research assistant recorded each actor's string of choices. The experimental trading protocol is pictured in figure 1.

General methods.—Before each session of 12 trials, two experimenters (E_1 and E_2, wearing different colors)[11] 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 (frame A in fig. 1). Each experimenter held a dish with a food reward (in clear view of the subjects) approximately 6 inches above the opening closest to the interior of the cage and extended an

[11] For expositional simplicity each experimenter is denoted by E_x, where x is the number of pieces of food the experimenter would initially display before possibly adding or taking away pieces.

empty hand into the other panel opening. If the capuchin took a token from the tray (frame B) and placed it in an experimenter's hand (frame C), then the experimenter would lower his food dish and present the capuchin with the food reward (frame 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 exchanged all 12 of its budgeted tokens for food rewards. Nonstandard 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 it wished to.

Subjects each participated in experiments 1, 2, and 3 in sequence, moving from one experiment to the next when their choices in the previous one had stabilized. In each of our experiments this criterion was set as five consecutive sessions in which a capuchin allocated its tokens in near-constant proportion.[12] We took these final choices to express each actor's preferred split between the choices each experiment affords. In our data analysis we use only the last five sessions for each actor; each capuchin took between six and 12 sessions to stabilize. Once a subject was finished with an experiment, we transitioned it 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 its 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 its new options.

V. Preliminary Experiments: Capuchins Obey Price Theory

Our preliminary experiments closely mirror those of Kagel and his coauthors and allow us to directly test that capuchin choice looks broadly

[12] This meant that the token allocations moved no more than one out of each session's 12 trials, for five consecutive sessions.

rational and admits standard price theory. In order to do so, we first found a set of two goods for the subjects between which they were roughly indifferent and 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 it responded.

A. *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 12 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 10 sessions, it was determined that the subject was roughly indifferent between the two goods offered (these foods ended up being either grapes or gelatin cubes). Until this combination was found, the nonapple experimenter changed the good he or she offered until this interior budgeting condition was satisfied.

Baseline and compensated price shift.—Once an appropriate good was found, the next steps of our price shift experiment were very straightforward. To establish a baseline measurement, each one of three subjects was repeatedly asked to allocate a budget of 12 disks between food 1 (apples) and food 2 (either grapes or gelatin). This was done exactly as described in the general methods above, with each experimenter trading one token for one piece of the respective food reward.

Each capuchin was run on this baseline condition until its choices stabilized, that is, until its choices did not change by more than one token for a span of five sessions. Once an actor had stabilized, its behavior over the next week was averaged into a representative consumption bundle (see the solid budget line in fig. 2). Using this bundle, we assigned a new budget of disks to each actor for use in a compensated price shift.

In this compensated price shift, the experimenter who trades for apples changed the amount he was 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. The budget was reduced from 12 to either nine or 10 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

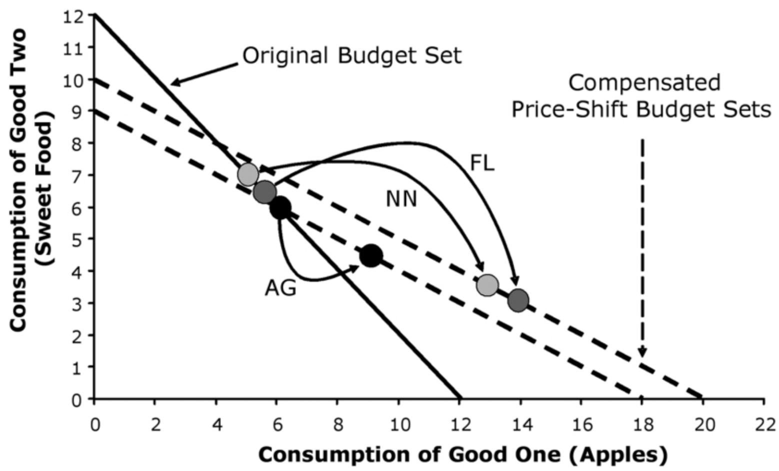

FIG. 2.—Subjects' choices over baseline and compensated price shift regimes. Capuchins satisfy GARP; each shaded point represents the purchasing behavior of a single subject in either a baseline or compensated price shift condition. We represent their average choice behavior from the week after they had learned the trading regime and their choices stabilized.

the dashed budget lines in fig. 2). Each subject's preferences were again allowed to stabilize, and then another week's worth of sessions were elicited under this new price regime.

B. Results: Preliminary Price Theory Experiments

The results of the preliminary pricing experiment are summarized in table 1 and figure 2. In figure 2, the solid line represents the initial budget set each actor was presented in the baseline condition, and the dashed 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 10 sessions, and every actor's choices easily satisfy GARP at the 1 percent level.

Note though that with only one compensated price shift, several naive models of choice will satisfy GARP. Indeed as Becker (1962) points out, many forms of random behavior can satisfy GARP in response to compensated price changes. In this setting, note that since we have experimenters switch sides of the testing chamber after each trial (to eliminate side bias), constant budget rationing could arise from complete inattention. To account for this we also test whether each subject's choices are inelastic enough to have arisen from random behavior, using as our null that neither price nor wealth shocks affect a capuchin's choices.

TABLE 1
PRICE THEORY AND COMPENSATED PRICE SHIFTS

	ACTOR		
	FL	NN	AG
	A. Baseline Experiment[a]		
Food used for good 1	Apples	Apples	Apples
Food used for good 2	Gelatin	Grapes	Gelatin
Amount spent on good 1	47%	42%	51%
Number of trials	132	72	144
	B. Compensated Price Shift Experiment[b]		
New budget of disks	10	10	9
Amount spent on good 1	69%	64%	50%
Good 1 consumed before → after shift	5.6→13.8	5.0→12.8	6.1→9.0
Number of trials	140	100	216
Choices satisfy GARP[c]	$p < .001$	$p < .001$	$p < .001$
Choices respond to shift[d]	$p < .001$	$p < .004$	$p < .966$

NOTE.—Tests of significance are reported as p-values of a two-sided test.
[a] Subjects were given a budget of 12 disks, and both goods have price 1.
[b] Good 1's price falls from 1 to .5, good 2's price stays 1, and the budget shrinks.
[c] Our 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.
[d] This 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 but AG does not.

Accordingly, in table 1 we also examine the percentage of trials each capuchin spends trading for apples before and after the price shift and test whether these percentages differ. Two out of three subjects showed a significant response to this test at the 1 percent level. Note that our third subject, AG, does not pass this test, spending 50 percent of its time (and, consequently, its budget) trading for apples both before and after the price shift. AG's behavior suggests that it either is maximizing Cobb-Douglas preferences (with a price elasticity of minus one) or is inattentive to prices when choosing which experimenter to trade with.

VI. Main Experiments: Are Capuchins Reference-Dependent?

Once our original subjects completed these initial experiments, an additional three subjects were recruited for our main set of experiments; AG was dropped from the study for becoming unresponsive. In this set of experiments the same budgeting procedure was used to elicit choices, with each session composed of 12 opportunities to trade a token for one of two possible food rewards. In contrast to the initial experiment, though, only apples were used as a food reward, and the experimenters no longer automatically presented the capuchin with the apples displayed in their tray when given a token. Now, experimenters sometimes 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, which sometimes consisted of a gamble.

A. *Methods: Experiment 1, Stochastic Dominance*

In experiment 1, a capuchin could trade its 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. Experimenters E_1 and E_2 also differed in how many pieces they initially showed the capuchin: E_2 displayed two squares of apple, and E_1 displayed only one square of apple. This experiment tests whether capuchins' 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 deliver both pieces only 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 an 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 subject.

B. *Methods: Experiment 2, Reference Dependence*

In experiment 2, 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, either E_2 would present the subject with the two apple pieces he had displayed or he 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 50-50 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 versus penalties conditions.

C. *Methods: Experiment 3, Loss Aversion*

In experiment 3, 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 2, 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 identical payoffs. However, on all trades E_1 gave exactly the quantity of apple he displayed, whereas E_2 displayed a quantity of apple that was always reduced from two to one before it was made available to the subject.

VII. Results: Main Experiments

The results of all our experiments are reported in tables 2, 3, and 4, 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 that allows more precise analysis.

The results of experiment 1 are summarized in table 2, which shows how the subjects behaved over five sessions (60 trials), after an initial set of sessions in which their choices stabilized as they learned about the experimental choices. The capuchins express a clear preference (87 percent 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.

TABLE 2
Experiment 1: Gambles and Stochastic Dominance

	Subject Name				
	FL	HG	JM	MD	NN
Trials E_1 chosen*	10%	18%	12%	22%	5%
Sessions till stable	7	6	7	6	11

Note.—Experimenter E_1 shows one and gives one; E_2 shows two and then gives one or two with probability one-half.
* All subjects' choices are different from 50 percent at the 1 percent level in a two-sided p-test.

TABLE 3
Experiment 2: Reference Dependence in Gambles

	Subject Name				
	FL	HG	JM	MD	NN
Trials E_1 chosen*	68%	70%	70%	70%	78%
Sessions till stable	11	9	9	9	13

Note.—Experimenter E_1 shows one and E_2 shows two. Then both give one or two with probability one-half.
* All subjects' choices are different from 50 percent at the 1 percent level in a two-sided p-test.

Given this result, though, the results of experiment 2 are quite surprising. All subjects left experiment 1 conditioned to favor E_2, the experimenter who displays two pieces of food. Despite this, in experiment 2 the capuchins quickly reverse this preference and trade much more with the experimenter who displays only one piece of food. Table 3 summarizes these results. 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 percent chance of a bonus rather that a 50 percent chance of a loss. Pooled, subjects traded with E_1 in 71 percent 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 1 but from random (50-50) behavior.

Note that any theory of choice that 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 its preferred experimenter from side to side, expending both time and attention.

Experiment 3 shows that this effect is not confined to risky choices and, when combined with experiment 2, suggests that capuchins are not just reference-dependent but loss-averse. The results are summarized in table 4. 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 provided the same, sure

TABLE 4
EXPERIMENT 3: RISKLESS REFERENCE DEPENDENCE

	Subject Name				
	FL	HG	JM	MD	NN
Trials E_1 chosen*	73%	75%	80%	82%	87%
Sessions till stable	9	10	10	8	10
Trials greater than in experiment 2	$p<.27$	$p<.27$	$p<.10$	$p<.07$	$p<.11$
Pooled, trials greater than in experiment 2	$p<.023$	$p<.023$	$p<.023$	$p<.023$	$p<.023$

NOTE.—Experimenter E_1 shows one and E_2 shows two. Then both give one.
* All subjects' choices are different from 50 percent at the 1 percent level in a two-sided p-test.

TABLE 5
EXPECTED GAINS, LOSSES, AND VALUES FOR EACH EXPERIMENTAL CHOICE

	Experiment 1		Experiment 2		Experiment 3	
	E_1	E_2	E_1	E_2	E_1	E_2
Gamble offered	(1, 1, 1)	(2, 1, 2)	(1, 1, 2)	(2, 1, 2)	(1, 1, 1)	(2, 1, 1)
Gains	0	0	.5	0	0	0
Losses	0	.5	0	.5	0	1
Expected value	1	1.5	1.5	1.5	1	1
Trials chosen	13%	87%	71%	29%	79%	21%

NOTE.—The table is constructed pooling all subjects' last five sessions after choices stabilize (60 trials).

payoff of one apple piece. For all subjects this preference (percentage of trials trading with E_1) was stronger than in experiment 2, suggesting that 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 that our capuchins are reference-dependent even in riskless choice settings.

To investigate whether capuchins display loss aversion, we now compare their behavior across experiments. Table 5 summarizes the aggregate behavior of our subjects in all three experiments and summarizes what each available gamble represented with respect to expected gains, losses, and food rewards. Looking at experiments 2 and 3, note that in experiment 2, 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 choose E_1. In experiment 3, however, E_1 always gave exactly what he showed whereas E_2 delivered a *sure loss*. An interpretation of the fact that subjects show a stronger preference in experiment 3 than in 2 (79 percent to 71 percent) 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 a subject's choices more than gains. In other words,

$$|\text{loss}| > \tfrac{1}{2}|\text{loss}| + \tfrac{1}{2}|\text{gain}| \Leftrightarrow |\text{loss}| > |\text{gain}|.$$

Thus the pooled p-test of experiment 3 being stronger than in experiment 2 ($p < .023$ in a two-sided p-test) can be taken to confirm the presence of loss aversion in capuchin choice.

A similar intuition suggests that expected value calculations weigh more heavily than framing effects on capuchin choice. Comparing experiments 1 and 2, note that in both choices, a half chance of a loss pushes the subject to prefer E_2 to E_1. However, these experiments differ in what force would attract a subject to E_1: in experiment 2, it is the half chance of a possible gain, whereas in experiment 1, it is the expectation of an extra half piece of apple. That the capuchins strongly prefer E_2 in experiment 1 and E_1 in experiment 2 suggests that increases in expected rewards greatly outweigh the effects of potential losses.

VIII. Discussion and Conclusion

Taken together our results suggest that capuchins' choice both is very sensitive to changes in prices, budgets, and expected payoffs and, to a lesser degree, displays both reference dependence and loss aversion. That expected rewards carry a much larger effect than gain/loss frames in our experiments is perhaps not surprising; our analysis examines only 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.

It is tempting to ask whether capuchins' loss aversion resembles human aversion in magnitude; for example, in much of the work on human loss aversion, the ratio of the coefficients on losses and gains is remarkably stable and is commonly used as a measure of loss aversion, representing how "kinked" the utility function is at the reference point. Tversky and Kahneman (1991) summarize a large body of survey evidence on minimally acceptable gambles and find a ratio of roughly 2.5 : 1. Average ratios of willingness to pay to willingness to accept found in most endowment effect experiments (see, e.g., Kahneman et al. 1990) yield a ratio of around 2.7 : 1. Benartzi and Thaler (1995) calibrate a ratio of approximately 2.3 from the aggregate risk preferences of stock investors. Although the amount of price variation we achieve with only three experiments is minimal, imposing additional structure on capuchin preferences allows the estimation of a comparable but extremely rough measure of losses to gains for our capuchins.[13]

We have argued that finding behavioral biases in capuchin choice

[13] In a preliminary analysis, we fit a simple linear utility specification to our capuchins' behavior. While parametric results are always to be taken with caution, the relative strength of losses to gains in capuchin decisions (the coefficient of loss aversion) seems similar to human estimates. For details, see the working version of this paper on Chen's Web site (http://www.som.yale.edu/Faculty/keith.chen/).

can suggest an early-evolutionary origin for these biases in humans. This relies heavily on the fact that on 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 that common traits evolved once during our common evolutionary heritage. While our results are by no means definitive proof that loss aversion is innate in humans, to the degree that they make us more likely to believe that some amount of this behavior has a biological component, they may have implications for how we treat loss-averse tendencies in human behavior.

For example, if these biases are innate, we may be more inclined to believe that they will persist in both common and novel settings, will be stable across time and cultures, and may endure even in the face of large individual costs, ample feedback, or repeated market disciplining. This would greatly constrain both the potential for successful policy intervention and the types of remedies available. In contrast, while a learned, noninnate heuristic may arise in many (if not all) cultures, we may not expect it to persist in settings in which it was highly suboptimal 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 that increase feedback or learning may eliminate what losses do exist.

Short of this attributional leap, however, that loss-averse behavior is not confined to humans can be seen as adding scope to the growing evidence of loss-averse behavior in many aspects of human economic behavior. Understanding how broadly these biases manifest themselves may influence how we should incorporate them into an adequate model of individual decision making. It is hoped that 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

Afriat, Sydney N. 1967. "The Construction of Utility Functions from Expenditure Data." *Internat. Econ. Rev.* 8 (February): 67–77.

Battalio, Raymond C., Leonard Green, and John H. Kagel. 1981. "Income-Leisure Tradeoffs of Animal Workers." *A.E.R.* 71 (September): 621–32.

Battalio, Raymond C., John H. Kagel, and Don N. MacDonald. 1985. "Animals'

Choices over Uncertain Outcomes: Some Initial Experimental Results." *A.E.R.* 75 (September): 597–613.

Battalio, Raymond C., John H. Kagel, Howard Rachlin, and Leonard Green. 1981. "Commodity-Choice Behavior with Pigeons as Subjects." *J.P.E.* 89 (February): 67–91.

Becker, Gary S. 1962. "Irrational Behavior and Economic Theory." *J.P.E.* 70 (February): 1–13.

———. 1976. *The Economic Approach to Human Behavior.* Chicago: Univ. Chicago Press.

Benartzi, Shlomo, and Richard H. Thaler. 1995. "Myopic Loss Aversion and the Equity Premium Puzzle." *Q.J.E.* 110 (February): 73–92.

Boysen, Sarah T., and Gary G. Berntson. 1995. "Responses to Quantity: Perceptual versus Cognitive Mechanisms in Chimpanzees (Pan troglodytes)." *J. Experimental Psychology: Animal Behavior Processes* 21 (January): 82–86.

Boysen, Sarah T., Kimberly L. Mukobi, and Gary G. Berntson. 1999. "Overcoming Response Bias Using Symbolic Representations of Number by Chimpanzees (Pan troglodytes)." *Animal Learning and Behavior* 27 (May): 229–35.

Brosnan, Sarah F., and Frans B. M. de Waal. 2003. "Monkeys Reject Unequal Pay." *Nature* 425 (September 18): 297–99.

———. 2004. "Socially Learned Preferences for Differentially Rewarded Tokens in the Brown Capuchin Monkey (Cebus apella)." *J. Comparative Psychology* 118 (June): 133–39.

Camerer, Colin F. 2000. "Prospect Theory in the Wild: Evidence from the Field." In *Choices, Values, and Frames,* edited by Daniel Kahneman and Amos Tversky. Cambridge: Cambridge Univ. Press.

Fragaszy, Dorothy M., Elisabetta Visalberghi, and Linda Marie Fedigan. 2004. *The Complete Capuchin: The Biology of the Genus Cebus.* Cambridge: Cambridge Univ. Press.

Genesove, David, and Christopher Mayer. 2001. "Loss Aversion and Seller Behavior: Evidence from the Housing Market." *Q.J.E.* 116 (November): 1233–60.

Harbaugh, William T., Kate Krause, and Timothy R. Berry. 2001. "GARP for Kids: On the Development of Rational Choice Behavior." *A.E.R.* 91 (December): 1539–45.

Harbaugh, William T., Kate Krause, and Lise Vesterlund. 2001. "Are Adults Better Behaved than Children? Age, Experience, and the Endowment Effect." *Econ. Letters* 70 (February): 175–81.

Henrich, Joseph, et al. 2001. "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies." *A.E.R. Papers and Proc.* 91 (May): 73–78.

Herrnstein, Richard J., and Drazen Prelec. 1991. "Melioration: A Theory of Distributed Choice." *J. Econ. Perspectives* 5 (Summer): 137–56.

Houthakker, Hendrik S. 1950. "Revealed Preference and the Utility Function." *Economica* 17 (May): 159–74.

Kagel, John H., et al. 1975. "Experimental Studies of Consumer Demand Behavior Using Laboratory Animals." *Econ. Inquiry* 13 (March): 22–38.

Kagel, John H., Raymond C. Battalio, and Leonard Green. 1995. *Economic Choice Theory: An Experimental Analysis of Animal Behavior.* Cambridge: Cambridge Univ. Press.

Kagel, John H., Raymond C. Battalio, Howard Rachlin, and Leonard Green. 1981. "Demand Curves for Animal Consumers." *Q.J.E.* 96 (February): 1–16.

Kagel, John H., Don N. MacDonald, and Raymond C. Battalio. 1990. "Tests of

'Fanning Out' of Indifference Curves: Results from Animal and Human Experiments." *A.E.R.* 80 (September): 912–21.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1990. "Experimental Tests of the Endowment Effect and the Coase Theorem." *J.P.E.* 98 (December): 1325–48.

———. 1991. "The Endowment Effect, Loss Aversion, and Status Quo Bias: Anomalies." *J. Econ. Perspectives* 5 (Winter): 193–206.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (March): 263–91.

Kralik, Jerald D., Marc D. Hauser, and Rasa Zimlicki. 2002. "The Relationship between Problem Solving and Inhibitory Control: Cotton-Top Tamarin (Saguinus oedipus) Performance on a Reversed Contingency Task." *J. Comparative Psychology* 116 (March): 39–50.

Liv, Chanya, Gregory C. Westergaard, and Stephen J. Suomi. 1999. "Exchange and Value in *Cebus apella*." *American J. Primatology* 49 (1): 74–75.

McClure, Samuel M., David I. Laibson, George Loewenstein, and Jonathan D. Cohen. 2004. "Separate Neural Systems Value Immediate and Delayed Monetary Rewards." *Science* 306 (October 15): 503–7.

Odean, Terrance. 1998. "Are Investors Reluctant to Realize Their Losses?" *J. Finance* 53 (October): 1775–98.

Rayo, Luis, and Gary S. Becker. 2005. "Evolutionary Efficiency and Happiness." Working paper, Univ. Chicago.

Samuelson, Larry, and Jeroen M. Swinkels. 2006. "Information, Evolution and Utility." *Theoretical Econ.* 1 (March): 119–42.

Samuelson, Paul A. 1938. "A Note on the Pure Theory of Consumer's Behavior." *Economica* 5 (February): 61–71.

Schneider, Horácio, et al. 2001. "Can Molecular Data Place Each Neotropical Monkey in Its Own Branch?" *Chromosoma* 109 (February): 515–23.

Tversky, Amos, and Daniel Kahneman. 1991. "Loss Aversion in Riskless Choice: A Reference-Dependent Model." *Q.J.E.* 106 (November): 1039–61.

Varian, Hal R. 1982. "The Nonparametric Approach to Demand Analysis." *Econometrica* 50 (July): 945–73.

Westergaard, Gregory C., Chanya Liv, Tara J. Chavanne, and Stephen J. Suomi. 1998. "Token Mediated Tool-Use by a Tufted Capuchin Monkey (*Cebus apella*)." *Animal Cognition* 1 (2): 101–6.

Westergaard, Gregory C., Chanya Liv, Andrea Rocca, Allison Cleveland, and Stephen J. Suomi. 2004. "Tufted Capuchins (*Cebus apella*) Attribute Value to Foods and Tools during Voluntary Exchanges with Humans." *Animal Cognition* 7 (1): 19–24.

Wolfe, John Bascom. 1936. *Effectiveness of Token-Rewards for Chimpanzees.* Baltimore: Johns Hopkins Press.