JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 2008, **89,** 145–155 NUMBER 2 (MARCH)

ON LOSS AVERSION IN CAPUCHIN MONKEYS

ALAN SILBERBERG[1], PETER G. ROMA[1], MARY E. HUNTSBERRY[1], FREDERICK R. WARREN-BOULTON[2], TAKAYUKI SAKAGAMI[3], ANGELA M. RUGGIERO[4], AND STEPHEN J. SUOMI[4]

[1]AMERICAN UNIVERSITY
[2]MICRA, INC.
[3]KEIO UNIVERSITY
[4]NATIONAL INSTITUTE OF CHILD HEALTH AND HUMAN DEVELOPMENT

Chen, Lakshminarayanan, and Santos (2006) claim to show in three choice experiments that monkeys react rationally to price and wealth shocks, but, when faced with gambles, display hallmark, human-like biases that include loss aversion. We present three experiments with monkeys and humans consistent with a reinterpretation of their data that attributes their results not to loss aversion, but to differences between choice alternatives in delay of reinforcement.

Key words: loss aversion, token exchange, reversed contingency, delay of reinforcement, prospect theory, choice, monkeys, humans

Loss Aversion in Capuchins

In Experiment 2 of a study by Chen, Lakshminarayanan, and Santos (2006), an experimenter stood in front of each of two opposed walls of a cage that housed a capuchin monkey (*Cebus apella*). There were two holes in each of these wire mesh walls. At the beginning of a trial, each experimenter, one male and one female, displayed a hand, palm upward, near each of the holes in the walls. In one hand was a dish containing either one or two pieces of apple, depending on the experimenter. The other hand was empty. The monkey responded by taking a token from a tray and placing it in the empty hand of one of the experimenters. If the token was given to the male experimenter, he would deliver the two pieces of apple to the monkey, or would remove one piece before food delivery, placing the removed food in a bowl. Alternatively, if the token was given to the female experimenter, she would give the one piece of food to the monkey through the adjacent aperture in the wall of the cage, or occasionally add one piece of food prior to delivery by taking the additional piece from a bowl. For both experimenters, changing the quantity in the food dish prior to delivery occurred randomly on half the trials. A 71% mean preference for the food that was increased on half the trials emerged even though the returns from exchange were the same, on average, from both experimenters.[1]

In Experiment 3, the female experimenter displayed a single piece of apple at the beginning of a trial, while the male experimenter displayed two pieces of apple. If the monkey exchanged with the female experimenter, then the single piece of apple on display was delivered. If, however, the monkey exchanged with the male experimenter, then one piece of apple was removed from the dish before food delivery. In consequence, the subject always received a single piece of apple regardless of locus of exchange. Despite this apparent equivalence of outcomes, a 79% preference for the fixed, single-item alternative was present at the end of training.

Address correspondence to Alan Silberberg, Department of Psychology, American University, Washington, DC 20016-8062; email: asilber@american.edu.

We thank Ruth Woodward, Steve Dinterman, Katalin Kerekes, and the NIH Animal Center veterinary, animal care, and facilities teams for their help and cooperation throughout the monkey portions of the project. This research was supported by intramural funds from the National Institute of Child Health and Human Development/National Institutes of Health/Public Health Service/US Department of Health and Human Services, and by the Japan Society for the Promotion of Science.

doi: 10.1901/jeab.2008-89-145

[1] In Experiment 1, Chen et al. (2006) showed that monkeys' response to a compensated price change was consistent with the generalized axiom of revealed preferences. This result joins demonstrations of income effects, Giffen goods and the like in linking human economic behavior to that seen in nonhuman primates (e.g., Silberberg, Warren-Boulton, & Asano, 1987). Experiment 1 is not discussed further because it makes no claim for demonstrating the phenomenon the present report targets, that of loss aversion.

In both experiments, robust preferences emerged even though the amount of apple delivered to the monkeys was the same regardless of the alternative purchased. These results are clearly incompatible with psychological accounts of choice allocation such as those offered by the matching law (Herrnstein, 1961).

How then are these results to be interpreted? Chen et al. (2006) appeal to the notion of loss aversion, a component of prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1991). Loss aversion refers to the finding that, for people, reductions in wealth cause greater dissatisfaction than the satisfaction produced by equivalent increases in wealth (see Rabin & Thaler, 2001). Therefore, when risk is present in choice, humans will often choose to avoid the prospect of a loss even when assuming the risk has positive expected value.

Chen et al. (2006) ascribe their results to a loss-aversion process: In Experiments 2 and 3, monkeys avoid exchanging a token for the alternative that is reduced by removal of a piece of apple. Presumably, the removal of a piece of food is perceived as a loss by the subjects. If this is human-like loss aversion, then subject preferences for the alternative not reduced are explained, and a probable evolutionary process common to humans (*Homo sapiens*) and monkeys is identified.

An Empirical Problem: Reversed-Contingency Effect

The explanation Chen et al. (2006) offer is consistent with their data, but their results nevertheless surprise those who are aware of the reversed-contingency procedure and the performance it fosters. In this procedure, an animal is given a choice between two alternatives that differ only in amount. If animals reach for the larger amount, they are given the smaller amount; and if they reach for the smaller amount, they receive the larger amount. The maximizing solution to this problem is to reach for the smaller amount of food. Nevertheless, nonhuman primates consistently have been incapable of doing this without remedial or atypically prolonged training (e.g., Boysen & Berntson, 1995; Murray, Kralik, & Wise, 2005; Silberberg & Fujita, 1996). Given these results, how did Chen et al. 's monkeys so readily master exchanging coins in Experiments 2 and 3 for the lesser amount of food?

Mindful of the reversed-contingency effect, Chen et al. (2006) explained their results by appeal to a second manipulation in Boysen and Bernston's (1995) original demonstration of this phenomenon. They noted that Boysen and Bernston found that apes could not maximize food consumption on a reversed-contingency procedure when ''sights of food amounts'' served as the discriminative stimuli. However, these animals were able to maximize when these stimuli were replaced by others (in this case, numerals on placards) that bore a relation to food amounts not burdened by the lifelong overtraining history that sights-of-food-amounts–food-consumption pairings had produced (see Silberberg & Fujita, 1996). How, then, did Chen et al.'s capuchins avoid a reversed-contingency effect? In their view, the tokens used in exchange were akin to the numerals used by Boysen and Bernston. That is, tokens broke the linkage between ''sights of food amounts'' and responding for larger food amounts, uncovering thereby the presence of loss aversion in choice.

From the perspective of a learning theorist, Chen et al.'s (2006) account may be troubling. Classically, the tokens used in their study would be viewed as conditioned reinforcers (see Cowles, 1937; Wolfe, 1936), not as discriminative stimuli. Yet in their report, the tokens must play the same role as numerals in the Boysen and Bernston (1995) study—and that is as discriminative stimuli. One could argue that the roles of discriminative stimulus and conditioned reinforcer are overlapping. However, there are demonstrations that the effects of these stimuli are separable (see Fantino, 1977). This leaves Chen et al. with the problem of explaining why their study's conditioned reinforcers also function as discriminative stimuli. Even if such an explanation were forthcoming, another problem remains because Chen et al.'s task requires two discriminative stimuli, one for each alternative. How can a single stimulus, a token, typically viewed as a conditioned reinforcer, serve as discriminative stimuli distinguishing two outcomes of choice?

These problems with the Chen et al. (2006) analysis call for an experimental test of their thesis that the use of tokens abrogated their study's reversed-contingency effect. Toward this end, we assess the effects of a reversed-contingency procedure on choice when a token is used to exchange for food rewards.

EXPERIMENT 1

Method

Subjects

Three male, tufted capuchin monkeys (*Cebus apella*), ranging in age from 4 to 14 years, served as subjects. The monkeys (Garth, HotRod, and Manuel) were experienced in exchanging chow biscuits and other objects for grapes with experimenters. They lived as part of a stable social group of 19 capuchins in an indoor complex composed of three interconnected enclosures, each sized 5.79 × 6.40 × 2.44 m (width × depth × height). The enclosures were separated from each other by translucent Plexiglas with small floor-level openings allowing passage between them. The front face of each enclosure was made of heavy wire mesh.

Monkey chow and water were freely available throughout the experiment. The facility-established feeding and environmental enrichment programs remained intact throughout the experiment except that late-afternoon fresh-fruit supplements were now given at the conclusion of each day's testing. All experimental procedures were conducted between 1400 and 1800 hr.

All activities involving animals were conducted in compliance with the US Animal Welfare Act and National Research Council guidelines, and were approved by the National Institute of Child Health and Human Development Animal Care and Use Committee.

Materials

Store-bought miniature marshmallows, cut in half, served as reinforcers. The tokens used for exchange sessions were metal fender washers (2.3 cm in diameter × 1 mm thick) painted with light blue enamel.

Procedure

The study, conducted over 9 consecutive days at the rate of one session per day, was overseen by three experimenters (the "Distributor," the "Exchanger," and the "Recorder") arrayed in that order, parallel to, and approximately 70 cm from, the wire-mesh front of the experimental space. The experimenters' location along the front face of the cage varied according to the location of the monkey being tested. All experimenters were aware of the results commonly found in tests of the reversed-contingency effect.

The Distributor presented single tokens to the monkey through the wire mesh, while the Recorder informed the Exchanger of which randomly assigned hand should hold two marshmallow halves. The Exchanger adjusted her position upon the monkey's return from token retrieval so that the monkey's head was equidistant between her palms up, outstretched hands (hands spaced 8–10 cm apart) when she began a trial. She simultaneously presented both food options to the monkey (two marshmallow halves in one open palm, and one marshmallow half in the other) with each option resting in the joint between her palm and fingers. The monkey was required to place the token on the Exchanger's hand between the food and fingertips, at which point the Exchanger withdrew that hand and made the alternate hand and its food available to the subject. The Recorder documented the outcome of each trial. Trials were conducted without a scheduled intertrial interval.

In sessions when either the Distributor or the Recorder could not participate, the other performed both tasks; however, the role of Exchanger was always fulfilled by the same experimenter.

All monkeys in the group were free to move about the housing enclosure at all times during testing, and all their interactions with the experimenters were voluntary. For this reason, data collection for each monkey at each session was opportunistic, and the number of trials for each monkey varied from session to session.

In the pretest condition, each monkey was permitted to select one or two marshmallow halves in a 25-trial free-choice procedure. The choices were presented to the monkeys as described above, except the monkeys were permitted to take the food from the chosen hand at each trial. The Distributor was present, but the monkeys did not use tokens.

The reversed-contingency test followed the pretest condition. For all trials, the monkey first retrieved a token from the Distributor. Then the Exchanger presented the monkey with half a marshmallow in one hand and two

Table 1

Trials per session in Experiment 1 in which a subject deposited a token in the experimenter's hand holding 1 piece of marshmallow in order to get 2. The numbers in parentheses indicate the number of trials in a session. Cases where most choices were consistent with mastery of the reversed-contingency problem are denoted by an asterisk.

	Session							
Subject	1	2	3	4	5	6	7	8
Garth	22 (58)	24 (52)	26 (60)	30 (75)	23 (52)	28 (62)	26 (64)	38 (77)
HotRod	27 (78)	38 (76)	30 (65)	31 (58)*	27 (65)	17 (68)	20 (46)	16 (44)
Manuel	27 (79)	39 (81)	31 (75)	36 (61)*	29 (63)	31 (64)	12 (39)	14 (38)

marshmallow halves in the other, the loci of each alternating randomly from trial to trial. The hand the monkey placed the coin in was immediately withdrawn and the monkey was offered the contents of the alternate hand.

The reversed-contingency procedure consisted of eight sessions. Within each session, the location of the Distributor relative to the Exchanger (i.e., to the left or to the right) was switched at the approximate half-way point for each monkey. The Distributor's location for the first half of each session for each monkey alternated from day to day. A total of 500 trials per monkey was collected. The mean number of trials per session was 63 per monkey, with the combined number of trials per session ranging from 149–215.

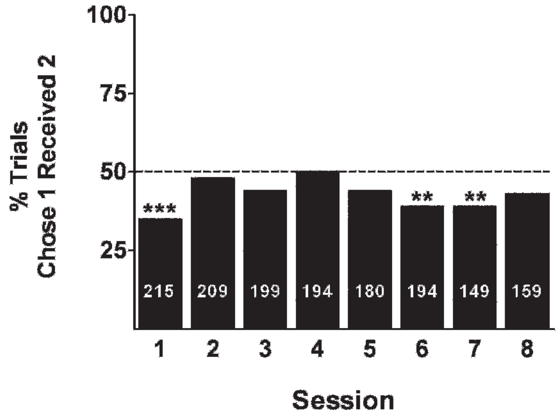

Fig. 1. Percentage of trials where the monkeys purchased one marshmallow piece and received two as a function of session number. The white-number inset within each bar discloses the number of pooled trials for a given session. The broken horizontal line indicates chance-level performance. Statistical significance is denoted by ** (binomial $p < .01$) and *** (binomial $p < .001$).

Results

All data were pooled and the group's preference in each session was evaluated by calculating binomial probabilities under an assumed $p = 0.5$ for choice to each outcome in pretest and reversed-contingency conditions. Statistical significance was set at $\alpha = .05$ for this and all other analyses.

As regards the pretest condition, the monkeys showed a significant preference for two marshmallow halves over one, opting for the larger alternative on 87% of the trials ($p < .001$). During sessions 1, 6, and 7 of the reversed-contingency test, the monkeys placed the token on the hand holding two marshmallow halves on a significant majority of trials (61–65%) ($ps < .01$). As shown in Table 1, preference for the single marshmallow half exceed 50% on only two occasions.

Figure 1 presents the percentage of trials in which the subjects deposited the token in the hand containing one marshmallow half, the behavior consistent with mastery of the reversed-contingency procedure, as a function of session number. As is apparent, subjects failed to master this task.

Discussion

In Experiment 1, monkeys that preferred more over less were unable to learn over 500 trials to exchange a token to the hand containing the smaller food amount in order to obtain the contents of the alternate hand that contained the larger amount of food. These preference data are typical of those seen in the reversed-contingency effect in the absence of token use (e.g., see Silberberg & Fujita, 1996). In consequence, it seems probable that Chen et al. (2006) err in their claim that use of tokens accounts for their failure to obtain a reversed-contingency effect. In addition, it weakens credence in Chen et al.'s loss-aversion effect: If the factors causing the unprecedented failure of the reversed-contingency effect in Chen et al. also are factors in their producing loss aversion, then loss aversion, like their failed reversed-contingency effect, may be caused by unintended features of experimental design.

At first blush, it might seem that this experiment's failed reversed-contingency effect should be followed by an attempted replication of Chen et al. (2006). We demur. After all, a monkey capable of demonstrating loss aversion must show mastery of the reversed-contingency task because the latter task is embedded in the former. That our monkeys fail in this task despite token use leaves no basis for expecting loss aversion in what would be a generalized replication of Chen et al.'s procedure.

The questions of why Chen et al. (2006) fail to find the reversed-contingency effect, and the probably related question of why they find loss aversion, call out for resolution. We attempt to do this in Experiment 3. However, the next questions we address are not of Chen et al.'s procedures and how they relate to extant data sets, but of the evolutionary linkages Chen et al. draw between capuchins and humans, and their implications for interpreting their results.

EXPERIMENT 2

In Chen et al. (2006; also see Chen & Santos, 2006), the search for loss aversion in capuchins is based on the evolutionary argument that humans and capuchins likely share an ancestor in which loss aversion was a heritable characteristic. Therefore, they believe loss aversion exists in both of these species and can be demonstrated given the appropriate experimental design.

As noted earlier, the appropriate design for humans is based on tests from prospect theory. A representative problem might be giving a person a choice between, say, winning \$100 or losing \$50 with equal probability versus winning \$40 for sure. A frequent finding is that subjects opt for the sure outcome despite its lower expected value. Such a result is the hallmark of loss aversion—the subject chooses to avoid loss rather than maximize returns.

This example of human loss aversion differs from that used by Chen et al. (2006) in several ways, but the difference we underscore is that the prospect-theory gamble for humans is unique whereas that of Chen et al. is repeated. The reason for this difference is obvious. Humans can understand differing prospects through language, whereas a capuchin can only learn the nature of prospects through experience. In consequence, a gamble can be ''one shot'' in an assay of human loss aversion, but must be repeated again and again before it can be fairly evaluated by a monkey.

Is this difference consequential? One way to answer this question is to assess risky choice under conditions where this difference is eliminated. Since unique gambles cannot be meaningfully administered to a monkey, using a common risky-choice procedure in humans and monkeys means using a repeated-gambles design. Fortunately, the literature provides guidance on the effects of repeated gambles on human risky choice, and they can be compared to the findings in Experiments 2 and 3 of Chen et al. (2006).

When a risky prospect is presented repeatedly to humans, the loss aversion that so dominates one-shot data tends to disappear. An example of this comes from Samuelson (1963). He offered a colleague a bet based on a flip of a coin that would either earn the bettor \$200 or cost him \$100. Samuelson's subject turned him down (he was loss averse), but said he would accept the bet if Samuelson offered it 100 times in succession. The results of Samuelson's thought experiment have been corroborated by studies showing that repeating a prospect causes human loss aversion to diminish and maximizing to appear (Keren & Wagenaar, 1987; Thaler, Tversky, Kahneman, & Schwartz, 1997; see Rabin & Thaler, 2001, for a discussion).

If, as Chen et al. (2006) suggest, people and monkeys share the process of loss aversion as a genetic predisposition, then a repeated test for loss aversion such as the one they used to uncover loss aversion in their report should produce a similar loss-aversion effect when adapted to humans. But, as noted above, humans typically fail to show loss aversion when prospects are presented repeatedly. If the usual result obtains—human loss aversion diminishing when prospects repeat — then Chen et al.'s results with capuchins would, ironically, point not to a process similarity with humans, but to a process difference.

Experiment 2 addressed this issue by giving a loss-aversion test to humans akin to Chen et al.'s (2006) with monkeys (see their Experiment 3). In particular, humans were given one-yen coins with which to buy one of two money amounts: two 10-yen coins that were reduced to one when the subject purchased them, or one 10-yen coin that was delivered without addition or reduction. If the data conform to Chen et al.'s, then humans would prefer the outcome that is unchanged over the outcome that is reduced even though the expected value of each is the same. Alternatively, if the human data conform to those in the literature, then loss aversion should be minimal; and since the expected value of each outcome is the same, a human utility maximizer should be indifferent between them.

Method

Subjects

Two female and three male psychology students at Keio University in Tokyo, Japan served as subjects.

Procedure

Each subject was individually tested by an experimenter (Alan Silberberg) who sat at a table opposite from the subject. Four paper cups were arrayed on the table in front of the subject to form a diamond shape with each vertex serving as a cup location. The cup closest to and directly in front of the subject contained 1-yen coins (''1 yen'' cup). The cup directly in front of, but farthest from the subject was an empty ''10 yen'' cup. It was approximately 10 cm in front of the 1-yen cup. Finally, the left and right ''choice'' cups were displaced approximately 5 cm to the left and right of the line bisecting the 1-yen and 10-yen cups, each cup equidistant from the 1-yen and 10-yen cups.

Each trial began with the experimenter placing one 10-yen coin in front of one choice cup and two, side-by-side, 10-yen coins in front of the other, each money amount directly in front of the experimenter and approximately 25 cm from its respective choice cup. The subject purchased an alternative by dropping a 1-yen coin into the appropriate cup. Once the selection was made, the experimenter slid a single 10-yen coin to the subject. If the subject chose the single coin, it was pushed across the table top to the subject. If the subject chose the two-coin alternative, one of the two coins was pushed to the choice cup. The subject then placed the coin in the 10-yen cup while the experimenter arranged the choice array for the next trial.

The way to make choices and to store rewards in the 10-yen cup were explained to each subject before the session began. No comment was made about the fact that both alternatives resulted in the same payoff. Each session ended after 25 choices. The position of the one- and two-coin alternatives strictly alternated across trials. There was no scheduled interval between successive trials. A 1-min rest period separated the two sessions comprising the experiment. At the end of the experiment, the subject kept all of the money (500 yen) accumulated in the 10-yen cup.

Results

For each session of the experiment, trials were pooled across subjects, and the proportion of choices made in each session was subjected to binomial probability tests based on the assumption of $p = 0.5$. In addition, mean data from the group were compared across sessions by a paired-sample, one-tailed t-test.

When offered a choice between two coins that were reduced to one versus one coin that remained one, the two-coin alternative was preferred in a significant majority of trials during Sessions 1 (binomial $p < .01$) and 2 (binomial $p < .05$). For Sessions 1 and 2, the mean frequencies and 95% confidence intervals for choice of the two-coin alternative were 15.6 ± 6.38 and 15.4 ± 6.19, respectively. These mean data show no evidence of a between-session shift in choice toward the one-coin alternative ($t(4) = .084$, $p > .90$).

As a percent of trials, these means represent a 62% preference for the two-coin alternative in each session. Even at the end of training, no evidence of a shift in preference toward the one-coin alternative was apparent: Pooled across subjects, the preference for the two-coin option was 72% during the last five trials of Session 2.

Individual subject data are presented in Table 2.

Table 2

Number of trials per 25-trial session in which humans chose two, 10-yen coins over one. Regardless of the choice made, the subject received a single 10-yen coin immediately after responding.

	Session	
Subject	1	2
HA1	15	10
HA2	23	23
HA3	10	16
HA4	12	16
HA5	18	12

Discussion

In the present experiment, humans preferred two 10-yen coins that were reduced to one, over a 10-yen coin delivered without addition or reduction. Moreover, when the second session's data were compared with the first or with the last five trials of the experiment, there was no evidence of a change in preference as the subjects gained experience with the contingencies. This preference is opposite from that in Chen et al.'s (2006) Experiment 3, where they found that monkeys preferred the single-piece-of-food alternative.

These results are not surprising. We noted earlier that loss aversion tends not to appear in multiple-trial gambles. For this reason, we did not expect to replicate Chen et al. (2006). That a significant preference developed for the alternative opposed to Chen et al.'s loss-aversion prediction is not of concern because actually no pattern of choice allocation violates maximizing predictions. The fact that the two-coin alternative was sometimes preferred may represent in humans the overtrained preference for more over less that is seen so often in reversed-contingency studies in other species. In any case, this did not happen for all subjects in both sessions, and no systematic preference is apparent in the data.

In this experiment, an analogue of the Chen et al. (2006) loss-aversion procedure failed to demonstrate loss aversion in humans. Given that Chen et al. claim humans and capuchins share a heritable predisposition to be loss averse, and given that humans are a species that is demonstrably loss averse, the failure of our experiment to uncover a preference for the single, unchanged alternative as seen in Chen et al.'s Experiment 3 raises the possibilities either that their procedure does not test loss aversion, or that capuchins are not loss averse.

Accepting this argument invites a question: If loss aversion did not produce Chen et al.'s (2006) results, what did? A review of their methods provides a possible answer to this question. In their Experiment 3, choice of the unchanging food amount led to delivery of the purchased item, whereas purchase of two pieces of food reduced to one led to: (a) removal of a piece of food, placing it in a bowl out of reach of the monkey; and (b) delivery of the remaining piece of food. Since removing food and placing it in a bowl takes time, the delivery of the two-foods-reduced-to-one outcome must have been delayed relative to the unchanged-food alternative. The literature offers many examples of where even small differences in reinforcer immediacy impact choice (e.g., Chung & Herrnstein, 1967; Rachlin & Green, 1972). By this argument, monkeys preferred the constant food amount in Experiment 3 of Chen et al. not because they were loss averse, but because it was delivered more quickly than the alternative.

In Experiment 3, we attempt to demonstrate the operation of this hypothesis by repeating the prior experiment with one important change: When the subject chose the two-coin alternative, one of the two coins was placed in a cup (experimenter's cup) below the table next to the experimenter. Only then did the experimenter slide the remaining coin to the subject's choice cup. This time-consuming step mimicked the behavior of the experimenter in Chen et al. (2006) who removed a piece of apple, placing it in a bowl that the monkey could not access, before giving the remaining piece of apple to the monkey. Implicit in this arrangement is the assumption that in Experiment 3 of Chen et al., the time it took the experimenter to remove a piece of apple from the dish containing two pieces of food before delivering the single remaining piece of food approximates the time it took our experimenter to place a single coin in the experimenter's cup before he delivered the remaining coin to the subject. If this assumption is correct, then the differential delay of reinforcement in Chen et al. would be successfully simulated in the proposed experiment and, if delay of reinforcement is the potent variable we claim it to be, should result in a clear preference for the unchanged, single, 10-yen coin.

EXPERIMENT 3

Method

Subjects

Four male and six female students from Keio University, none of whom participated in the prior experiment, served as subjects.

Method

The procedure was unchanged from Experiment 2 except that: (a) a cup was placed on a chair adjacent to the experimenter (Alan Silberberg); and (b) whenever the subject chose the two-coin-reduced-to-one alternative, the experimenter removed one of the two coins and placed it in the cup. Thereafter, he slid the remaining coin to the subject's appropriate choice cup.

Results

One subject only chose the two-coin-reduced-to-one alternative. Since he was a native Japanese speaker and the experimenter only spoke English, it is possible he misunderstood the instructions. In any case, this study required that choices be made on a between-outcome comparison of contingencies. Since his performance could not possibly have been based on such a comparison, his data were excluded from analysis. All analyses are based on the remaining 9 subjects.

Pooled across subjects, Session 1's data showed indifference in choice (49.8% choice of the immediate, one-coin alternative; binomial $p = 1.00$). However, during Session 2, a significant majority of trials was to the one-coin option (74.2%, binomial $p < .001$). The mean and 95% confidence intervals of one-coin choices during Sessions 1 and 2 were 12.4 ± 4.56 and 18.6 ± 4.08, respectively. As hypothesized, the mean preference for the one-coin alternative increased significantly across sessions ($t(8) = 2.10$, $p < .05$).

Table 3 presents the individual data for the 9 subjects who completed the experiment. These data show that for 6 of the 9 subjects, choice of the two-coin option decreased from Session 1 to Session 2, whereas choice increased for the other 3 subjects. Also notice that 7 of the 9 subjects chose the immediate, one-coin alternative on more than half the trials in Session 2.

Table 3

Number of trials per 25-trial session in which humans chose two, 10-yen coins over one. Regardless of the choice made, the subject received a single 10-yen coin after responding, either immediately (if single-coin alternative was chosen) or after a delay (if the two-coin alternative was chosen).

	Session	
Subject	1	2
HB1	4	2
HB2	18	6
HB3	8	15
HB4	19	6
HB5	18	4
HB6	13	14
HB7	9	2
HB8	6	9
HB9	18	0

Discussion

When the two-coin alternative was selected in the prior experiment, a coin was pushed to the choice cup with approximately the same latency as were choices of the single-coin option. The present experiment differed in that, instead of leaving the remaining coin from the two-coin alternative on the table as was done previously, it was now removed from the table and placed in a cup before the other coin was pushed to the choice cup. As a consequence, Experiment 3, unlike Experiment 2, had an asymmetry in response–reinforcer latencies: Choices to the two-coin alternative took longer to reward than choices of the one-coin outcome. We hypothesized that this difference would shift preference in the direction of the one-coin option, a hypothesis that was realized in the data.

Lifting the coin off the table and placing it in a cup was intended as a proxy for Chen et al.'s (2006) removal of a piece of apple from the two-piece alternative in their Experiment 3. If the proxy is accepted, then a probable cause of their finding of a strong preference for the single-apple alternative is identified. That option, unlike choices to the two-piece-reduced-to-one alternative, was delivered immediately to the capuchin. Perhaps as a consequence of its greater immediacy of reinforcement, the single-piece-of-apple outcome was preferred over its alternative. This appeal to relative immediacy of reinforcement requires no concession to a loss-aversion effect. Indeed, as noted earlier, the multiple-trial procedure used in Experiment 2 of this report and Experiments 2 and 3 in Chen et al. may, by virtue of their repeated-gambles design, have precluded the operation of a loss-aversion process.

GENERAL DISCUSSION

In their final experiment, Chen et al. (2006) found that capuchins prefer to exchange a token for a single piece of apple over two pieces of apple that are reduced to one. These results surprise, for they violate the reversed-contingency effect—the finding that nonhuman primates have a tendency to reach for more food over less despite the reduced reinforcement this choice provides. Chen et al. explained their violation of the reversed-contingency effect by arguing this effect can be abrogated when choice is mediated by tokens. In Experiment 1, we tested their explanation and found it wanting: Despite many trials of exposure, none of our subjects mastered the reversed-contingency problem even though all reinforcers required an exchange of tokens.

Our second experiment questioned a central premise of Chen et al.'s (2006) work: their assumption that their procedure relates to loss aversion as seen in humans. The methods for studying loss aversion in humans are well established, frequently tested and conclusive in the picture they paint of the reality of this effect. Chen et al. can make no similar claims for their test because it differs in many ways from human loss-aversion tests. In fact, it is by no means clear to us that their test relates to loss, never mind loss aversion. If there is reality to the distinction drawn by the expression ''a bird in the hand is worth two in the bush,'' then one might ask whether watching two apple pieces reduced to one prior to delivery is equivalent, or even related to, an actual loss.

In any case, in Experiment 2 our evaluation of the Chen et al. (2006) test is not dependent on the reader's opinion in this matter. Instead, we focus on Chen et al.'s theory that loss aversion is a heritable characteristic common to humans and capuchins. If so, the loss aversion they claim their test demonstrates in their Experiment 3 should be apparent when their test is given to humans. This prediction follows from the fact that humans are the only species for which the presence of this phenomenon appears to be well established in the literature. As matters turned out, loss aversion did not appear in our Experiment 2 when their procedure was given to humans. That their effect does not generalize to humans raises concern with the claim that this test indexes loss aversion. Further, it jeopardizes their evolutionary view of loss aversion as a heritable characteristic. Rather than pointing to a similarity in process between humans and capuchins, a comparison of Chen et al.'s finding with ours suggests, if anything, a process difference.

We expected our Experiment 2 to fail Chen et al.'s (2006) loss-aversion test because repeated-gamble procedures such as theirs typically fail to produce loss aversion (Keren & Wagenaar, 1987). Given this expectation, why then did Chen et al. get their loss-aversion effect—a preference for the unchanged, single-food alternative in their Experiment 3? We attributed their results to the fact that their design caused the single-apple reinforcer to be delivered sooner than the two-apple-reduced-to-one outcome. Our Experiment 3 illustrated the operation of this difference in reinforcer immediacy by simulating with a coin the removal of a single piece of apple in Chen et al.'s Experiment 3. This step made the delay between choice and reward unequal between alternatives, and explains the finding of preference for the unchanged, single-good outcome in our study and theirs in terms of differences in reinforcer immediacy between choice alternatives.

Although the results of Chen et al.'s (2006) third experiment are readily explained in terms of this delay-of-reinforcement account, its application to the results of their Experiment 2 is problematic. To recall the methods of that experiment, food was occasionally added to one alternative or subtracted from the other. Unlike the outcome in their Experiment 3, we see no a priori reason to anticipate that the reinforcer delay caused by adding food differs systematically from that of removing food. If food was added or removed without producing a between-alternative difference in reinforcer latency, then standard accounts of intertemporal choice argue that the alternative that was occasionally reduced should have been preferred to the alternative that was occasionally increased. This discounting-based prediction follows from the fact that in Chen et al.'s second experiment, the difference in value between one and two pieces of apple delivered immediately should outweigh the value difference in these outcomes when delivered with a delay (e.g., see Frederick, Loewenstein, & O'Donoghue, 2002, and Green & Myerson, 2004, for overviews of this literature). Since the aggregate value of an alternative is a composite of the values of the immediate and the delayed outcomes, a delay-based account of choice in Chen et al. leads to the expectation that capuchins should have preferred the immediate outcome of two pieces of apple that were occasionally later reduced to one. Yet Chen et al. found the opposite result—a preference for the immediate, one-piece alternative that was occasionally later increased to two.

Surprisingly, the failure of our account in accommodating the results of Chen et al.'s (2006) second experiment does not require that we abandon differences in reinforcer delays to explain that experiment's outcomes. A delay-based account can be preserved by noting that Chen et al. made no attempt to counterbalance the experimenters making these changes in food amounts. Hence, perhaps the experimenter who added food to the one-piece outcome did so with greater alacrity than did the experimenter who removed food from the two-piece outcome. If this possibility was realized, a difference in delay of reinforcement would distinguish outcomes, and an animal responding for the more immediate outcome would appear, by Chen et al.'s criteria, to be loss averse when in fact the subject was possibly responding only to differences in reinforcer immediacy.

This argument—that there may have been a between-experimenter difference in the time taken to deliver reinforcement in their Experiment 2—is given additional credence by its persistence as a design flaw in Chen et al. (2006). In all three of their experiments, they failed to counterbalance experimenters: across experiments, the option offered by the female experimenter was preferred to that offered by the male experimenter. Such a result is entirely consistent with the idea that she delivered reinforcement with greater rapidity than the male no matter how the outcome alternatives differed between experiments. If so, differences in delay of reinforcement may be present not just in Experiment 3 from Chen et al., but in every experiment in their report.

Why might this have occurred? One possibility is that the experimenters themselves were aware of the hypotheses under consideration, and the very interesting outcome of showing animal loss aversion unconsciously influenced their results. Of course, this risk was not only present in their work, but in ours, as well. In both cases, the risk of experimenter bias stands as a methodological limitation of this research.

We close by underscoring that while our work does not endorse Chen et al.'s (2006) thesis that humans and capuchins share a loss-aversion process, it also does nothing to prove this thesis incorrect. Indeed, the notion of a linkage in loss aversion between capuchins and humans is plausible, and may be established in future work. However, until that work is forthcoming, we urge caution in: (a) expanding the claimed domain of loss-aversion effects beyond the species *Homo sapiens*; and (b) using work such as Chen et al. as evidence for a biological substrate to loss-aversion effects in humans.

REFERENCES

Boysen, S. T., & Berntson, G. G. (1995). Responses to quantity: Perceptual vs. cognitive mechanisms in chimpanzees (*Pan troglodytes*). *Journal of Experimental Psychology: Animal Behavior Processes, 21*, 82–86.

Chen, M. K., Lakshminarayanan, V., & Santos, L. R. (2006). How basic are behavioral biases? Evidence from capuchin monkey trading behavior. *Journal of Political Economy, 114*, 517–537.

Chen, M. K., & Santos, L. R. (2006). Some thoughts on the adaptive function of inequity aversion: An alternative to Brosnan's social hypothesis. *Social Justice Research, 19*, 201–207.

Chung, S. H., & Herrnstein, R. J. (1967). Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior, 10*, 67–74.

Cowles, J. T. (1937). Food-tokens as incentives for learning by chimpanzees. *Comparative Psychology Monographs, 23*, 1–96.

Fantino, E. (1977). Conditioned reinforcement: Choice and information. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 313–339). Englewood Cliffs, NJ: Prentice Hall.

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature, 40*, 351–401.

Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin, 130*, 769–792.

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior, 4*, 267–272.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263–292.

Keren, G., & Wagenaar, W. A. (1987). Violation of utility theory in unique and repeated gambles. *Journal of Experimental Psychology: Learning, Memory and Cognition, 13*, 387–391.

Murray, E. A., Kralik, J. D., & Wise, S. P. (2005). Learning to inhibit prepotent responses: Successful performance by rhesus macaques, *Macaca mulatta*, on the reversed-contingency task. *Animal Behavior, 69*, 991–998.

Rabin, M., & Thaler, R. H. (2001). Anomalies: Risk aversion. *Journal of Economic Perspectives, 15*, 219–232.

Rachlin, H., & Green, L. (1972). Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior, 17*, 15–22.

Samuelson, P. A. (1963). Risk and uncertainty: A fallacy of large numbers. *Scientia, 98*, 108–113.

Silberberg, A., & Fujita, K. (1996). Pointing at smaller food amounts in an analogue of Boysen and Berntson's (1995) procedure. *Journal of the Experimental Analysis of Behavior, 66*, 143–147.

Silberberg, A., Warren-Boulton, F. R., & Asano, T. (1987). Inferior-good and Giffen-good effects in monkey choice behavior. *Journal of Experimental Psychology: Animal Behavior Processes, 13*, 292–301.

Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics, 112*, 647–661.

Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *Quarterly Journal of Economics, 106*, 1039–1061.

Wolfe, J. B. (1936). Effectiveness of token-rewards for chimpanzees. *Comparative Psychology Monographs, 12*, 1–72.

Received: July 27, 2007
Final acceptance: October 9, 2007