ISSN 0143-4543

University of Aberdeen Business School

Payoff levels, loss avoidance, and equilibrium selection in the Stag Hunt: an experimental study

By

Nick Feltovich, Atsushi Iwasaki and Sobei H. Oda

Discussion Paper 2008-3
February 2008

Editor: Dr W David McCausland
www.abdn.ac.uk/business/

Payoff levels, loss avoidance, and equilibrium selection in the Stag Hunt: an experimental study

Nick Feltovich*
University of Aberdeen Business School
Edward Wright Building, Dunbar Street
Aberdeen AB24 3QY, UK
n.feltovich@abdn.ac.uk

Atsushi Iwasaki
Department of Intelligent Systems
Graduate School of Information Science and Electrical Engineering, Kyushu University
Motooka 744, Nishi–ku, Fukuoka, 812–8581, Japan

Sobei H. Oda
Faculty of Economics, Kyoto Sangyo University
Kamigamo, Kita–ku, Kyoto 603–8555, Japan

February 26, 2008

Abstract

Game theorists typically assume that changing a game's payoff levels—by adding the same constant to, or subtracting it from, all payoffs—should not affect behavior. While this invariance is an implication of the theory when payoffs mirror expected utilities, it is an empirical question when the "payoffs" are actually money amounts. In particular, if individuals treat monetary gains and losses differently, then payoff–level changes may matter when they result in positive payoffs becoming negative, or vice versa. We report the results of a human–subjects experiment designed to test for two types of loss avoidance: certain–loss avoidance (avoiding a strategy leading to a sure loss, in favor of an alternative that might lead to a gain) and possible–loss avoidance (avoiding a strategy leading to a possible loss, in favor of an alternative that leads to a sure gain). Subjects in the experiment play three versions of Stag Hunt, which are identical up to the level of payoffs, under a variety of treatments. We find differences in behavior across the three versions of Stag Hunt; these differences are hard to detect in the first round of play, but grow over time. When significant, the differences we find are in the direction predicted by certain– and possible–loss avoidance. Our results carry implications for games with multiple equilibria, and for theories that attempt to select among equilibria in such games.

Journal of Economic Literature classifications: D83, C72, C73.
Keywords: experiment, game theory, behavioral economics, stag hunt, learning.

*Corresponding author. Oda and Iwasaki are supported in part by Grants–in–Aid of Scientific Research (#17310029 and #17700155) and by the Ministry of Education, Culture, Sport, Science and Technology (Japan). We thank Alexander Elbittar, Andreas Ortmann, John Van Huyck, Nathaniel Wilcox, Erte Xiao, two anonymous referees, and participants at several conferences for helpful suggestions and comments.

1 Introduction

Game–theoretic solution concepts carry the implication that modifying a game by adding the same (positive or negative) constant to all payoffs should have no effect on behavior.[1] Such changes do not affect a player's rank–ordering of outcomes, so that pure–strategy best responses, and hence pure–strategy Nash equilibria, are unaffected. Furthermore, they do not affect expected–payoff differences between strategies under any given set of beliefs about opponents' pure or mixed strategies, so that best–response correspondences, and therefore mixed–strategy Nash equilibria, are also unaffected. When the game in question has multiple Nash equilibria, game theory does not rule out the possibility that payoff–level changes affect which one of the equilibria is played, but it does not predict when such sensitivity will be present, nor how it will be manifested. When the game has a unique equilibrium, game theory specifically predicts that changing payoff levels can have no effect. So, game theory is at best silent as to whether changes in payoff levels affect players' choices, and at worst explicitly rules such an effect out.

This implication of game theory, however, relies on the game's payoffs reflecting the actual preferences of the players—that is, the equivalence between payoffs and players' expected utilities. When the game's payoffs are simply monetary gains or losses (as is often the case in economics experiments with human subjects), this equivalence cannot be taken for granted, and the robustness of behavior to changes in payoff levels becomes an empirical question. Indeed, there is some empirical evidence suggesting that changing payoff levels may affect behavior. In economics, this discussion dates back at least to Kahneman and Tversky (1979), who devised their "prospect theory" as an alternative to the standard version of expected–utility theory to account for the systematic violations of the predictions of the latter they observed in a large number of decision–making tasks.[2] (See also Tversky and Kahneman (1991, 1992).) Although they did not look specifically at the effects on decision–making behavior of payoff levels, their prospect theory allows for such sensitivity to payoff levels to be observed. More recently, Cachon and Camerer (1996) found that players' decisions in coordination–game experiments could be sensitive to changes in payoff levels, when these changes affected the *signs* of payoffs: previously–positive payoffs were now negative, or vice versa. (Section 2.2 has more discussion of the literature on payoff levels and behavior.) They speculated that subjects in the experiments were exhibiting "loss avoidance", which they defined to be a tendency to avoid choices that with certainty yield negative payoffs in favor of alternative choices that could yield positive payoffs.[3] Based on this finding, they conjectured that loss avoidance could be used as a criterion for equilibrium selection.[4]

Cachon and Camerer's notion of loss avoidance is one of several ways in which individuals could treat losses and gains differently. For the remainder of the paper, we will distinguish between their notion of loss avoidance, which we will call "certain–loss avoidance", and another type of loss avoidance. This latter type, which we call "possible–loss avoidance", is defined to be a tendency to avoid strategies

[1]For conciseness, we will refer to such a modification of a game as a "change in payoff levels".

[2]It should be noted that their experiment was conducted using hypothetical payments, not monetary ones, so their results should be viewed as suggestive, not conclusive. Some research has compared behavior under real incentives versus hypothetical incentives; for example, Holt and Laury (2002) found that when payments are small, behavior in decision–making tasks is similar under either monetary or hypothetical incentives, but as payments are scaled up, subjects become more risk averse under monetary incentives but not under hypothetical incentives. (See also Laury and Holt (2002).)

[3]More precisely, they speculated that subjects believed their opponents were exhibiting loss avoidance; see Section 2.2.

[4]In Appendix B, we discuss in some detail the extent to which loss avoidance and Kahneman and Tversky's (1979) prospect theory are connected. For now, we merely note that loss avoidance is consistent with, but neither entailing nor implied by, prospect theory, and that despite an unfortunate similarity in their names, loss avoidance is distinct from the component of prospect theory known as *loss aversion*.

that give a *possible* negative payoff, in favor of one that gives a certain positive payoff. Throughout this paper, we will use the term "loss avoidance" to encompass both certain– and possible–loss avoidance.[5]

The goal of this paper is to test for certain– and possible–loss avoidance in strategic behavior. To do so, we design and run an experiment using a stage game well–suited for such a test: Stag Hunt (Rousseau (1973)), a symmetric two–player game with two available actions, a risky action and a safe action. Three versions of Stag Hunt are shown in Figure 1; here, the risky and safe actions are labeled R and S respectively. These three games are constructed so that they differ only in payoff

Figure 1: The Stag Hunt games

High payoffs (SHH)

		Player 2: R	Player 2: S
Player 1	R	7,7	1,5
	S	5,1	5,5

Medium payoffs (SHM)

		Player 2: R	Player 2: S
Player 1	R	5,5	–1,3
	S	3,–1	3,3

Low payoffs (SHL)

		Player 2: R	Player 2: S
Player 1	R	1,1	–5,–1
	S	–1,–5	–1,–1

level—any one of the games could be obtained from either of the others by addition of a constant to, or subtraction of a constant from, all payoffs—so they are equivalent from a game–theoretic standpoint, as long as these payoffs really do represent the preferences of the players.

When these payoffs represent monetary gains and losses, then despite this apparent equivalence, our two notions of loss avoidance can be used to make directional predictions regarding how individuals' choices change across these three games. In the high–payoff game (SHH), all payoffs are positive, so neither certain– nor possible–loss avoidance will have an impact. In the medium–payoff game (SHM), on the other hand, choosing S leads to a certain positive payoff, but choosing R could lead to a negative payoff (if the opposing player chooses S). If an individual exhibits possible–loss avoidance in her decision making, R should therefore be less attractive here than it would have been in SHH. So, possible–loss avoidance implies that R is *less* likely to be chosen in SHM than in SHH, or equivalently (from an observational standpoint), that the individual will act as though she is *more* risk averse in SHM than in SHH.

By the same token, in the low–payoff game (SHL), choosing S leads to a certain negative payoff, but choosing R could lead to a positive payoff (if the opposing player also chooses R). An individual exhibiting certain–loss avoidance will find S less attractive here than it would have been in SHH. So, certain–loss avoidance implies that S is less likely to be chosen in SHL than in SHH, so that R is *more* likely to be chosen in SHL than in SHH, or equivalently, that the individual will act as though she is *less* risk averse in SHL than in SHH.

In our experiment, subjects play all three of these games. We manipulate various aspects of the strategic environment, including the number of times a game is played (once or forty times), how players are matched to each other (repeated play against the same opponent or random matching over all possible opponents), and how much information players are given about the payoffs of the game (complete information or limited information). We find that in all treatments, there are significant

[5]Cachon and Camerer (1996) speculate about yet another type of loss avoidance, in which individuals will avoid a choice that yields a negative *equilibrium* payoff (p. 167, note 2). They call this "losing–equilibrium avoidance". We do not look at this third type of loss avoidance in this paper, noting only that it is indeed different from both certain– and possible–loss avoidance.

differences in behavior across versions of the game—that is, behavior is indeed sensitive to payoff levels. The effect of payoff levels on behavior is small (and often statistically insignificant) in early rounds, but grows over time, typically peaking between the midpoint and the end of the experimental session. This effect, whenever significant, is in the direction predicted by loss avoidance. In some treatments, subjects exhibit certain–loss avoidance. In some treatments, they exhibit possible–loss avoidance. Sometimes, subjects exhibit both certain– and possible–loss avoidance, but in none of our treatments do they exhibit neither.

The rest of the paper is structured as follows. In Section 2, we describe the concept of loss avoidance, as well as the design of the experiment we use to test for it, including the games subjects play, and we discuss some of the relevant work done by others. In Section 3, we present several aspects of the experimental results, including tests of statistical significance, and compare our results to those of others. Section 4 concludes.

2 The experiment

The experiment involves the three versions of Stag Hunt shown in Figure 1 and described above, played under various experimental manipulations.

2.1 The Stag Hunt games

As mentioned already, each player in a Stag Hunt game chooses between two actions, which we call Risky (R) and Safe (S). Safe leads to a payoff that does not depend on what the other player does, while Risky earns a variable payoff: higher than the certain payoff from Safe if the other player also chooses Risky, but lower if the other player chooses Safe. Notice that our medium–payoff game (SHM) is obtained by subtracting 2 from every payoff of our high–payoff game (SHH), and our low–payoff game (SHL) is obtained by subtracting 6 from every payoff of SHH. So, from a game–theoretic standpoint, these games are identical. They have the same best–reply correspondences, and the same three Nash equilibria: (R,R), (S,S), and a mixed–strategy equilibrium in which both players choose R with probability 2/3. Not only do these games have the same Nash equilibria, but criteria for equilibrium selection that are based only on payoff differences between outcomes—such as Harsanyi and Selten's (1988), which predicts the payoff–dominant (R,R) outcome (see also Selten (1995)), and Carlsson and van Damme's (1993), which predicts the risk–dominant (S,S) outcome—will make the same prediction for each of the games.

Stag Hunt games are particularly well–suited for studying equilibrium selection in general, and its sensitivity to payoff–level changes in particular. Both R and S belong to strict pure–strategy Nash equilibria, so either choice can be justified, given appropriate beliefs about the behavior of one's opponent; indeed, both (R,R) and (S,S) survive all standard equilibrium refinements. Furthermore, since all three equilibria (pure– or mixed–strategy) are symmetric, the coordination problems implicit in symmetric games with asymmetric equilibria are not an issue. On a related note, Stag Hunt games are games with strategic complementarities—the more likely a player believes her opponent is to choose a particular strategy, the stronger the attraction to that strategy is for her. As a result, a change to the environment that affects all players should be self–reinforcing: a change that made, say, the risky action more appealing to a player should also raise her perceived likelihood that her opponent would also choose the risky action (to the extent that the player viewed her opponent as having a similar thought process to hers), making the risky action more appealing still to her.[6]

Stag Hunt is a simple member of a class of coordination games called "order–statistic games" (Van Huyck, Battalio, and Beil (1990)). All of these games have the advantageous features of Stag Hunt for studying equilibrium selection: multiple strict symmetric Nash equilibria and strategic complementarities.[7] Typically in these games, any pure–strategy choice is justifiable, so which strategies decision makers actually choose in situations like this is an empirical question (as Schelling (1960, p. 162) pointed out). Early experiments involving these games showed that behavior in these games is difficult to predict—with the implication that it is quite possible that subjects will coordinate on an inefficient Nash equilibrium, or fail to coordinate on any pure–strategy equilibrium at all.[8] For example, Van Huyck, Battalio, and Beil (1990) found that subjects' choices in two multi–player seven–strategy minimum–effort games always converged to the least efficient pure–strategy equilibrium, though the speed of convergence depended on differences between the payoffs from different actions. On the other hand, they also found that when only two players played minimum–effort games, and were matched to each other repeatedly (a "fixed–pairs" matching), the opposite happened: choices tended to converge to the most efficient equilibrium. (In two–player games with random rematching of opponents, subject behavior did not appear to converge at all.) Van Huyck, Battalio, and Beil (1991) looked at a multi–player median–effort game, and found that choices after the first round tended to converge to whatever the first–round median was. Rankin, Van Huyck, and Battalio (2000) looked at Stag Hunt games with payoffs that varied across rounds, so that the payoff–dominant outcome was sometimes also risk dominant, and sometimes not. They found that population behavior always moved toward the payoff–dominant outcome, but these changes were faster when this outcome was also risk dominant than when it was not. Battalio, Samuelson, and Van Huyck (2001) looked more precisely at differences in payoffs from different actions by using three versions of Stag Hunt that had exactly the same best–reply correspondence, but varying penalties for not playing best replies; each subject would play one of these games 75 times. They found that behavior in the three games was similar initially, but diverged over time; in later rounds, play of the safe action became much more likely as the penalty for not playing a best reply increased.

Because early research failed to find strong evidence in favor of a single type of strategy, much recent research has focused on factors, besides the payoffs of the game, that make subjects more likely to choose one strategy over others. Stahl and Van Huyck (2002) found that giving subjects experience in playing different versions of Stag Hunt increased the likelihood of their playing the risky action. Berninghaus and Ehrhart (2001) found that giving subjects in a minimum–effort game more

[6]Indeed, the positive feedback doesn't stop with this one step. A player sophisticated enough to reason that her opponent will perceive her likelihood of choosing the risky strategy to increase, thus raising his likelihood of choosing it, will find the risky strategy to be still more appealing. In the extreme case where the effect of a payoff–level change were common knowledge amongst the players, there would be an infinite number of these chains of reasoning.

[7]In many order–statistic and related games, players' strategies are ordered in such a way that higher–numbered strategies are more socially beneficial than lower–numbered strategies. For example, symmetric pure–strategy Nash equilibria might be Pareto ranked from lower–numbered to higher–numbered strategies, or more strongly, one player's moving from lower–numbered to higher–numbered strategies might weakly increase all other players' payoffs. As a result, strategies in these games are often called "effort levels". A player's payoff is a function of the distance between her choice of effort level and some order statistic of all choices of the players in the group. Typically, either the minimum or the median is the order statistic used, leading to a "minimum effort" or "median effort" game. Our Stag Hunt games are minimum–effort games with only two strategies (R="high effort", S="low effort").

[8]There have been many experiments examining behavior in Stag Hunt and other order–statistic games. We do not attempt an exhaustive literature review here; instead, we concentrate on the literature closest in spirit to our experiment. For a more comprehensive review, see Ochs (1995), Camerer (2003, Chapter 7), and Devetag and Ortmann (2007).

information about others' choices at the end of the round increased the likelihood of their playing the payoff–dominant outcome. Devetag (2003) looked at the effect of varying end–of–round feedback in a different type of order–statistic game: multi–player "critical–mass" games. She found that giving subjects the entire distribution of choices made the payoff–dominant outcome more likely than when subjects were given less information (only the median choice), or no information at all (other than their own choice and payoff). Several researchers have looked at the effects of "cheap talk"—costless, nonbinding sending and receiving of messages—on behavior in these games (see, for example, Clark, Kay, and Sefton (2001), Manzini, Sadrieh, and Vriend (2002), Duffy and Feltovich (2003), Charness and Grosskopf (2004), Duffy and Feltovich (2006), and Blume and Ortmann (2007)). A typical result in these experiments is that allowing cheap talk increases the amount of risky–strategy play, and thus the likelihood of the payoff–dominant outcome, though such successful coordination is still far from complete. Charness and Grosskopf (2004) and Duffy and Feltovich (2006) showed additionally that coordination on the payoff–dominant outcome becomes still more likely if receivers of cheap–talk messages are given other information that can verify the truthfulness of the messages, such as the sender's previous–round action (and possibly the previous–round message as well). Clark and Sefton (2001) compared repeated Stag Hunt games (ten rounds) under two different matching protocols: fixed pairs and rotation (one play against each of ten opponents). They found that the payoff–dominant outcome was more likely under fixed pairs than under rotation matching, even in the first round; Clark and Sefton attributed this to subjects in the fixed–pairs treatment using the first round to signal to their opponents that they would choose the risky action. On the other hand, Schmidt et al. (2003) found no difference in aggregate play between treatments with fixed–pairs matching and treatments with random rematching in eight–round Stag Hunt games (with payoff matrices that varied across sessions), though they did find differences at the session level. Specifically, under fixed opponents, there was little or no change in aggregate behavior over time, while under changing opponents, play moved toward the payoff–dominant outcome in two games and the secure outcome in another (in the remaining game, the outcome varied across sessions). They also looked at one–shot versions of these games, and found that behavior in a given one–shot game was comparable to that in the first round under changing opponents. Heinemann, Nagel, and Ockenfels (2004) found that risky–choice frequencies in Stag Hunt are positively correlated with those in related lottery–choice decision problems as well as, in some cases, age and a measure of "experience seeking" determined by a personality test.

None of the research discussed in this section looked into the effect of payoff levels on behavior; research in this area will be discussed in the next section.

2.2 Payoff levels—literature and predictions

Previous experimental research into the effects of changing payoff levels can be classified according to whether subjects make their choices within the context of an individual decision problem or a game, whether they play one time or repeatedly, and how much information they have about the payoffs they can receive. As mentioned in the introduction, Kahneman and Tversky (1979) used results from a series of one–shot individual decision problems to guide the construction of their "prospect theory". Prospect theory does imply that changing payoff levels can affect choices in situations where there is uncertainty, though the exact nature of this effect may be difficult to predict a priori, as it can be sensitive to framing effects.[9]

[9]For example, Thaler and Johnson (1990) suggest the effect of past lump–sum gains or losses on future choices might vary, depending on whether the gains or losses have been internalized by the decision maker already.

Most of the research on payoff–level effects in repeated individual decision problems has been done by Ido Erev and colleagues.[10] Their experiments discussed here involved large numbers of repetitions of a decision problem by individuals who initially have no information about payoffs (though the instructions given to subjects imply that the environment is stationary, so that payoff information can be learned over time). Barkan, Zohar, and Erev (1998) considered a set of six decision problems, in each of which subjects chose between a safe and a risky alternative. Three of the problems were presented under a "gains" frame, where the safe choice gave a sure payoff of 0 and the risky choice gave—depending on the outcome of a random draw—either a payoff of +1 or a negative number that ranged between –10 and –109 across treatments. The probability of the bad outcome varied along with the resulting payoff across treatments, so that the expected payoff from the risky choice was the same in each of the three treatments, and less than the payoff to the safe choice. The other three problems were presented under a "losses" frame; each was formed from one of the gains treatments by subtracting 1 from every payoff (so that all payoffs were nonpositive). Each subject played one of these six decision problems 600 times. In all six problems, the frequency of safe choices rose over time. The results regarding payoff levels were mixed; in two of the pairs of problems, learning was not affected by changing the level of payoffs, while in the third pair (the one in which the bad outcome was the worst and least likely), learning the safe choice was substantially slower in the losses frame than in the gains frame.

Bereby–Meyer and Erev (1998) looked at individual choices in simpler decision–making tasks. Subjects were asked in each of 500 rounds to predict which of two events would occur. The probabilities of the two events were 0.7 and 0.3—i.i.d. across rounds. A subject who predicted correctly in a round would receive a high number of points for that round; a subject who predicted incorrectly would receive a low number of points. Their three treatments differed only in these rewards: the payoff for a correct prediction was 4, 2, and 0 in the three treatments, and the payoff for an incorrect prediction was always 4 points fewer; they called these their (4,0), (2,–2), and (0,–4) treatments, respectively. Since one choice in this situation (predicting the more likely event) stochastically dominates the other, Bereby–Meyer and Erev were primarily interested in how changing payoff levels affected the speed of convergence toward making the optimal choice with probability one. They found that convergence was slowest when no negative payments were possible, but there was little difference between the other two treatments.[11] Erev, Bereby–Meyer, and Roth (1999) extended this work by adding two new treatments, (6,2) and (–2,–6), so that the payoff for an incorrect prediction was still 4 points fewer than that for a correct prediction. They found that the speed of learning in these new treatments was similar to that in Bereby–Meyer and Erev's (4,0) treatment; thus, combining this result with the results of Bereby–Meyer and Erev's, convergence toward always choosing the dominant action is faster when losses were possible but not certain (the (2,–2) and (0,–4) treatments) than otherwise (the other

[10]With the exception of the aforementioned study by Kahneman and Tversky, all of the experiments discussed in this section involved actual—not hypothetical—monetary payments to subjects.

[11]As far as we know, Erev and colleagues were the first to look at the effects of changing payoff levels. However, earlier work by others gave suggestive results. Siegel and Goldstein (1959) studied decision making in tasks similar to those considered by Bereby–Meyer and Erev (1998) (see also Siegel, Siegel, and Andrews (1964)). In one experiment they report, they find that learning is faster—and asymptotic behavior closer to optimal—in a (5,–5) treatment than in a (5,0) treatment. Since these treatments differ not only in payoff levels, but also in the scale of payoffs, it's not possible to determine which of these is driving these results. However, as Bereby–Meyer and Erev point out, other studies (e.g., Myers et al. (1963)) found that quite substantial changes in the scale of payoffs had little effect on learning speed, suggesting that the differences found by Siegel and Goldstein may indeed have been due to payoff–level changes: specifically, the possibility of losses in one treatment but not the other.

three treatments).

Erev, Bereby–Meyer, and Roth (1999) also presented the results of an experiment involving two versions of a 2x2 constant–sum game with a unique mixed–strategy Nash equilibrium (also played 500 times), where payoffs were in units of probability of "winning" rather than "losing" (see the left game in Figure 2). The versions differed only in the sizes of the prizes—in one treatment, the prize

Figure 2: Constant–sum games used in previous payoff–level experiments

		Player 2	
		A2	B2
Player 1	A1	0.3, 0.7	0.8, 0.2
	B1	0.4, 0.6	0.1, 0.9

Erev, Bereby–Meyer, and Roth (1999)

		Player 2				
		C	L	F	I	O
	C	*W*	*L*	*L*	*L*	*L*
Player	L	*L*	*L*	*W*	*W*	*W*
1	F	*L*	*W*	*L*	*L*	*W*
	I	*L*	*W*	*L*	*W*	*L*
	O	*L*	*W*	*W*	*L*	*L*

(W = "winning" payoff, L = "losing" payoff)

Rapoport and Boebel (1992)

for "winning" was either 0.5 or 1, and the prize for "losing" was –0.5 or 0—so again, the payoffs in one treatment could be found by adding a constant to those in the other. Here, they found that in the treatment where losses were possible, subjects were more likely to choose the action that had yielded the higher average payoff over all previous rounds; they concluded that learning was faster when losses were possible than when they weren't. Earlier, Rapoport and Boebel (1992) had also looked at two versions of a repeated 5x5 constant–sum game with a unique mixed–strategy equilibrium (see the right game in Figure 2), played 120 times. The payoffs in their game were also framed in terms of "winning" and "losing", though theirs were certain wins or losses rather than the probabilistic payoffs used by Erev, Bereby–Meyer, and Roth. In one treatment, the payoff for "winning" was +\$10, while the payoff for "losing" was –\$6; in the other treatment, the payoffs were +\$15 and –\$1, respectively. Unlike Erev, Bereby–Meyer, and Roth (1999), Rapoport and Boebel found only insignificant differences in play between the two treatments. Note that one difference between their experiment and Erev, Bereby–Meyer, and Roth's is that in theirs, both gains and losses were possible in both treatments, while Erev, Bereby–Meyer, and Roth also had treatments without gains and treatments without losses.

A few researchers have looked at changes in payoff levels in the context of market experiments, in which subjects played the role of firms (and sometimes buyers as well), and changing payoff levels was accomplished by varying firms' sunk costs. Kachelmeier (1996) and Waller, Shapiro, and Sevcik (1999) considered double–auction markets with many buyers and sellers with varying valuations and costs. Both found that changes in the level of firms' sunk costs had no effect on market prices, which tended to be close to the market–clearing prices. (However, Kachelmeier found that changes of accounting report format are associated with changes in bids and asks when sunk costs are positive, but not when they aren't.) Offerman and Potters (2003) conducted experimental markets under several conditions. In one condition, subjects played price–setting firms with capacity constraints (Bertrand–Edgeworth duopolists); in each round, two firms were chosen randomly to enter the market, and those chosen were required to pay a fixed entry fee. Another condition was identical, except that those chosen were not charged any fee. Offerman and Potters found that subjects' price choices were higher when subjects had to pay the entry fee than when entry was free, though this was only a change in sunk–cost level, so that nothing changed but payoff levels. Buchheit and Feltovich (2008) also looked at repeated Bertrand–Edgeworth duopoly games, with a sunk cost that was varied across sessions over several values, and found a nonlinear effect of sunk–cost level on behavior: at low levels, increasing the sunk–cost level tended to result in lower average price choices, but after a point, increasing it tended to lead to higher average prices.

The studies most similar to ours have involved changes in payoff levels in coordination games. Cachon and Camerer (1996) looked at the effects of changing the level of payoffs in a median–effort game (see Figure 3) and a minimum–effort game with similar payoffs. Rather than actually changing

Figure 3: Cachon and Camerer's (1996) 7–player median–effort game

		Median effort level						
		1	2	3	4	5	6	7
	1	140	150	140	110	60	–10	–100
	2	130	160	170	160	130	80	10
Player's	3	100	150	180	190	180	150	100
effort level	4	50	120	170	200	210	200	170
	5	–20	70	140	190	220	230	220
	6	–110	0	90	160	210	240	250
	7	–220	–90	20	110	180	230	260

(Payoffs reduced by 225 units in "Must Play" treatment)

the numbers in the payoff matrices, they included a "Must Play" treatment in which they imposed a mandatory fee of 225 on all players, and made this fee common knowledge among the players.[12] They found that subjects tended to avoid strategies that yielded a certain negative payoff, which they termed "loss avoidance", and which we refer to as certain–loss avoidance. In another "Private Cost" treatment, the mandatory fee was not common knowledge; subjects were told their own fee but not their opponents'. In this treatment, certain–loss avoidance was not observed. Cachon and Camerer concluded that the subjects in the experiment themselves did not exhibit certain–loss avoidance; rather, the subjects believed that their opponents exhibited certain–loss avoidance, so acted in a way consistent with certain–loss avoidance due to the strategic complementarities in the game.

Cachon and Camerer conjectured that certain–loss avoidance could be used more generally for equilibrium selection in games: Nash equilibria involving strategies leading to certain losses ought not to be selected. Rydval and Ortmann (2005) performed a first test of this conjecture, examining behavior in several versions of Stag Hunt, including two pairs within which only payoff levels were changed. These four games are shown in Figure 4; note that Games 2 and 3 are identical except for the level of payoffs, as are Games 4 and 5 (Game 3 is obtained from Game 2 by subtracting 60 from all payoffs, and Game 5 is obtained from Game 4 by subtracting 180). Note also that within each pair, one game has only positive payoffs—like our SHH game—and the other had all negative payoffs except for the (R,R) profile—like our SHL game.[13] Subjects played each game once, with opponents

[12]They had yet another treatment ("Opt Out") in which the fee was optional in the sense that players could earn a payoff of zero by staying out of the game. When incurring the fee is optional, forward–induction arguments can be used to refine away some of the low–payoff Nash equilibria; this cannot be done when the fee is mandatory.

[13] Note also that Games 4 and 5 have triple the payoffs of Games 2 and 3, respectively, so that Rydval and Ortmann

Figure 4: Two high–payoff and two low–payoff Stag Hunt games from Rydval and Ortmann (2005)

Game 2

		Player 2: C	Player 2: D
Player 1	A	80,80	10,50
	B	50,10	50,50

Game 3

		Player 2: C	Player 2: D
Player 1	A	20,20	–50,–10
	B	–10,–50	–10,–10

Game 4

		Player 2: C	Player 2: D
Player 1	A	240,240	30,150
	B	150,30	150,150

Game 5

		Player 2: C	Player 2: D
Player 1	A	60,60	–150,–30
	B	–30,–150	–30,–30

changing from game to game, but with no feedback between games. In Games 2 and 3, subjects' choices were consistent with certain–loss avoidance; they were significantly more likely to choose the risky strategy when the safe strategy led to a sure negative payoff (Game 3) than when all payoffs were positive (Game 2). However, in Games 4 and 5, there was no difference between them in the frequency with which subjects played the risky strategy; thus the overall evidence for certain–loss avoidance in their experiment was weak.[14]

Our experiment is designed to detect not only certain–loss avoidance, but also possible–loss avoidance—the desire to avoid strategies that *might* yield a negative payoff (but also might yield a positive payoff), in favor of other available strategies that yield a sure positive payoff. Both types of loss avoidance imply predictions for our experiment. In SHH, all payoffs are positive, while in SHL, the safe action guarantees a negative payoff while the risky action might lead to a positive payoff. If individuals exhibit certain–loss avoidance, they will tend to avoid the safe action in SHL relative to SHH; thus, certain–loss avoidance implies:

Hypothesis 1 *(Certain–loss avoidance) Play of the risky action should be* more *likely in the low–payoff version of Stag Hunt than in the high–payoff version.*

On the other hand, in SHM, the safe action guarantees a positive payoff while the risky action might lead to a negative payoff. If individuals exhibit possible–loss avoidance, they will tend to avoid the risky action in SHM relative to SHH; thus, possible–loss avoidance implies:

Hypothesis 2 *(Possible–loss avoidance) Play of the risky action should be* less *likely in the medium–payoff version of Stag Hunt than in the high–payoff version.*

2.3 Experimental design and procedures

As already mentioned, our design uses three versions of Stag Hunt, so that we are able to look for possible–loss avoidance as well as certain–loss avoidance. We used a within–subjects design, so that each subject played all three versions. One concern we had was that, if subjects recognized the close

were also able to look for changes in behavior due to multiplication of payoffs by a constant: payoff–scale effects.

[14]Rydval and Ortmann suggest that because of the higher scale of payoffs in Games 4 and 5, the lack of difference in behavior between Games 4 and 5 may be due to the effects of loss aversion countering those of loss avoidance.

similarity of these games, they might view the experiment as a "consistency check" and choose the same actions in all three. To limit this possibility, we tried to make it less transparent that the three games were identical from a game–theoretic standpoint, without using any deception. This was done by adding three more games, shown in Figure 5: a version of the Prisoners' Dilemma (which we abbreviate PD), a version of Battle of the Sexes (BoS), and a coordination game (CG).

Figure 5: The other games used in the experiment

		Player 2	
		R	S
Player 1	R	2,2	0,0
	S	0,0	1,1

Coordination Game (CG)

		Player 2	
		R	S
Player 1	R	0,0	3,5
	S	5,3	0,0

Battle of the Sexes (BoS)

		Player 2	
		R	S
Player 1	R	7,7	1,8
	S	8,1	4,4

Prisoners' Dilemma (PD)

Subjects played all six games in the following order: CG-SH-BoS-SH-PD-SH. Another concern in our within–subjects design was that the results might be sensitive to the order in which the games were played. (Other researchers have shown that the order of the games can matter; see, for example, Falkinger et al. (2000) or Duffy and Feltovich (2004).) In order to ameliorate this problem, we varied (across experimental sessions) the order in which the three versions are played. We used three of the six possible Stag Hunt game orderings: SHH-SHM-SHL, SHM-SHL-SHH, and SHL-SHH-SHM.

We used four design treatments in an effort to examine several ways in which the two types of loss avoidance might manifest themselves. In our O (one–shot) treatment, subjects played each game once and the payoff matrices were commonly known. This treatment allows us to see how individuals behave in each game before acquiring any experience in playing that game, so that their decisions are the result of their own deductive reasoning, along with the understanding that their opponents have the same information. This treatment also allows the most direct comparison between our experimental results and Rydval and Ortmann's. Our C (complete–information) treatment was similar to our O treatment, except subjects played each game repeatedly—40 times—before moving on to the next game. Subjects in this treatment were matched randomly to opponents in every round. The results from this treatment allow us to see the consequences of subjects' deductive reasoning, like the O treatment, but also allow us to see whether, and how, choices change over time in response to the experience they receive in playing a game—in particular, whether loss avoidance develops or decays over time.

In our other two treatments, subjects were not told the payoff matrices of the games they were playing, though they were given information at the end of each round that would enable them to infer the payoff matrices eventually. (Specifically, a subject would be told her opponent's choice in the just–completed round, and her own payoff.) Our R (random–matching, limited–information) treatment was otherwise similar to the C treatment: subjects played each of the six games 40 times against randomly–chosen opponents. Our F (fixed–pairs, limited–information) treatment was similar to the R treatment, except that subjects were not rematched to new opponents after each round; instead, they played against the same opponent for all 40 rounds of a game.[15] These two treatments

[15] Van Huyck, Battalio, and Beil (1990) and Schmidt et al. (2003) (discussed in Section 2.1 above) also varied the matching mechanism between (in Schmidt et al.'s nomenclature) "fixed match" and "random match." However, subjects in all treatments of both of these experiments had complete information about payoffs; thus, the "random match"

allow us to see the effects of the experience subjects receive in playing each game, without the benefit of any initial introspection—as Erev and colleagues did with the decision–making experiments discussed in Section 2.2. In this way, we can determine whether loss avoidance might be learned, even when it is not possible to act in such a way through deductive reasoning. By comparing the results of these two treatments, we can also determine whether any learning of loss avoidance depends on whether subjects play repeatedly against the same opponent, or against changing opponents.[16]

In the two limited–information treatments, we were concerned that subjects' choices might be sensitive to the order in which the actions (Risky or Safe) appeared in the payoff matrix. (For example, it is possible that there is a bias in favor of the upper or left action at the expense of the lower or right action, purely because of their locations). To limit this possibility, we varied (across sessions) the order in which the actions appeared. The action orderings were R-S (risky action on top or left) and S-R (risky action on bottom or right). Table 1 summarizes our four treatments. When necessary to avoid confusion with our manipulations of game and action ordering, we will refer to them as "information treatments"—although in some cases (such as between F and R) they differ only in ways other than the information given to subjects.

Table 1: Information treatments used in the experiment

Information treatment	Number of rounds	Payoff information	Matching mechanism	Game orderings	Action orderings
O	1	Full	—	H-M-L, M-L-H, L-H-M	R-S
C	40	Full	Random	H-M-L, M-L-H, L-H-M	R-S
R	40	Limited	Random	H-M-L, M-L-H, L-H-M	R-S, S-R
F	40	Limited	Fixed pairs	H-M-L, M-L-H, L-H-M	R-S, S-R

Each experimental session involved subjects playing each of the six games under a single information treatment, with a single game ordering and single action ordering (thus, our variation of information treatments, game ordering, and action ordering are across subjects, while our variation in payoff levels is within subjects). Experimental sessions took place at the Economic Science Laboratory at Kyoto Sangyo University (Kyoto, Japan). Subjects were primarily undergraduate students at Kyoto Sangyo University, recruited via a database of participants in other experiments and via advertisements posted on campus. About a third of subjects were economics majors; the rest were relatively representatively distributed across the student population. No subject participated in more than one session. At the beginning of a session, each subject was seated at a computer and given written instructions. These instructions were also read aloud in an effort to make the rules of the experiment common knowledge. Partitions prevented subjects from seeing other subjects' computer

treatment is comparable to our C treatment. We have no treatment analogous to the "fixed match" treatment; our expectation was that given 40 rounds of a game with the same opponent and complete information about the payoff matrix, there would be quick convergence to the payoff–dominant outcome in all three of our versions of Stag Hunt.

[16]We note here that in all treatments other than our O (one–shot) treatment, the number of rounds subjects play is several times larger than the pool of potential opponents; as a result, subjects in our C and R treatments do play against the same opponent more than once. This is in contrast to other experiments using random rematching—such as Clark and Sefton (1999) and Schmidt et al. (2003), who have subjects playing against a given opponent at most once. We acknowledge that our design does open up the possibility of repeated–game behavior such as signaling or reputation–building, but we expect that these effects are limited, since matchings are completely anonymous: subjects receive no identifying information about their opponents (even their subject numbers).

screens, and subjects were asked not to communicate with each other during the session. The experiment itself was programmed in the Japanese version of the z–Tree experimental software package (Fischbacher (2007)), and all interaction took place via the computer network. Subjects were asked not to write down any results or other information.

At the beginning of an experimental session, subjects were told how many rounds of each game they would be playing. Prior to the first round of each game, they were reminded that they were to begin a new game. In this round, and in each round thereafter, subjects were prompted to choose one of two possible stage–game strategies. To reduce demand effects, actions had generic names (R1 and R2 for row players, C1 and C2 for column players), rather than being called Risky and Safe. After all subjects had chosen strategies for a round, each subject was told her own choice, her opponent's choice, and her own payoff. In the O, C, and R treatments, subjects were matched to a new opponent after each round; in the F treatment, subjects were matched to a new opponent when the 40 rounds of one game ended and they began a new game.

At the end of an experimental session, one round was randomly chosen, and each subject was paid 200 yen (at the time of the experiment, equivalent to roughly $1.90) for each point earned in that round. In addition, subjects in the C, R, and F treatments were paid a showup fee of 3000 yen, from which negative payoffs were subtracted, if necessary. The O treatment was tacked on to an asset–market experiment, in which subjects earned an average of 2700 yen, so no additional showup fee was paid.

3 Results

A total of 19 sessions were conducted: 3 of the O treatment, 3 of the C treatment, 6 of the R treatment, and 7 of the F treatment. The number of subjects in a session varied from 6 to 28. Table 2 gives some information about the individual sessions of the experiment.[17]

3.1 Initial choices under complete information

We first look at initial behavior of subjects playing the Stag Hunt games under complete information; this comprises the O treatment (in which each game was played only once) and the first round of each game in the C treatment.[18] Table 3 shows the corresponding relative frequencies of risky–action choices in each of the three Stag Hunt games. Several results are apparent. First, play of the risky action in all three games is more likely in Round 1 of the C treatment than in the O treatment, though the size of the difference varies across games: negligible in the low–payoff Stag Hunt, nonnegligible but insignificant ($\chi^2 \approx 1.42$, d.f.=1, $p > .10$) in the medium–payoff Stag Hunt, significant ($\chi^2 \approx 3.93$, d.f.=1, $p < .05$) in the high–payoff Stag Hunt.

Next, consistent with our Hypothesis 1 (certain–loss avoidance), risky–action play is substantially more prevalent in the low–payoff version of Stag Hunt than in the high–payoff version. Subjects play the risky action in the low–payoff version 92.4% of the time overall—91.7% of the time in the O treatment, and 93.1% in the first round of the C treatment—versus only 77.1% of the time in the high–payoff version (69.4% of the time in the O treatment, 84.7% in the first round of the C

[17]Our F4 session had only 6 subjects, so we ran an additional session (called F4a) with the same experiment parameters. The raw data from our experiment are available upon request.

[18]We do not analyze the data from the other three games (CG, BoS, PD) in this paper. A preliminary analysis of some aspects of these data, as well as some aspects of the Stag Hunt results unconnected with loss avoidance, can be found in Feltovich, Iwasaki, and Oda (2008).

Table 2: Session information

Session	Information treatment	Game ordering	Action ordering	N
F1	F	H-M-L	R-S	28
O1	O	H-M-L	R-S	22
C1	C	H-M-L	R-S	28
R1	R	H-M-L	R-S	20
O2	O	M-L-H	R-S	26
O3	O	L-H-M	R-S	24
F2	F	H-M-L	S-R	22
R2	R	H-M-L	S-R	20
F3	F	M-L-H	R-S	16
F4	F	M-L-H	S-R	6
R3	R	M-L-H	R-S	14
R4	R	M-L-H	S-R	10
C2	C	M-L-H	R-S	20
F5	F	L-H-M	R-S	10
C3	C	L-H-M	R-S	24
F6	F	L-H-M	S-R	18
R5	R	L-H-M	R-S	26
R6	R	L-H-M	S-R	18
F4a	F	M-L-H	S-R	26

treatment). The difference in risky–action play between SHL and SHH is significant for the O data alone (McNemar change test with correction for continuity, $p < 0.001$) but misses being significant for the first round of the C treatment by itself ($p \approx 0.15$)—though even here, the difference is in the direction predicted by certain–loss avoidance.[19] If we pool the O data and the first–round data from the C treatment, risky–action play is significantly more frequent in SHL than in SHH (McNemar test, $p < 0.001$).

On the other hand, initial behavior shows no evidence in favor of our Hypothesis 2 (possible–loss avoidance). When the O treatment and the first round of the C treatment are pooled, the relative frequency of risky–action choices in the high– and medium–payoff versions of Stag Hunt are exactly the same. Even if we consider the first round of the C treatment by itself, the difference in risky–action play between the two games—84.7% in SHH and 81.9% in SHM—is not significant at conventional levels (McNemar test, $p > 0.10$).

3.2 Aggregate behavior over all rounds

We next look at aggregate behavior over all 40 rounds of the three repeated–game treatments (C, F, and R), in order to see whether certain–loss avoidance persists after subjects have received some experience in the game, and understanding of the way other subjects are playing, and whether possible–

[19] Risky–action play in the first round of the C treatment is significantly more likely in SHL than in SHM (McNemar test, $p \approx 0.08$), suggesting weak support for the combination of certain– and possible–loss avoidance in this case. See Siegel and Castellan (1988) for descriptions of the nonparametric statistical tests used in this paper.

Table 3: Initial frequencies of risky action (complete information treatments, all sessions)

Game	O treatment		Round 1 of C treatment		Combined	
SHH	$.694^{a}$	(50/72)	$.847^{ab}$	(61/72)	$.771^{a}$	(111/144)
SHM	$.722^{a}$	(52/72)	$.819^{a}$	(59/72)	$.771^{a}$	(111/144)
SHL	$.917^{b}$	(66/72)	$.931^{b}$	(67/72)	$.924^{b}$	(133/144)

a,b: Within a treatment, frequencies with superscripts having no letters in common are significantly different at the 10% level or better, with the higher frequency corresponding to the letter coming later in the alphabet. Frequencies with superscripts having a letter in common are not significantly different. See text for details.

loss avoidance eventually develops. Table 4 shows some aspects of play in the three repeated–game treatments: overall (40–round) frequencies of risky–action choices, and 40–round frequencies broken down by ordering of games (H-M-L, M-L-H, or L-H-M) and by ordering of actions (R-S or S-R). These relative frequencies suggest that, while there is some variation across treatment, game ordering, and action ordering, differences in play across games is consistent with both certain–loss avoidance and possible–loss avoidance.

Table 4: Aggregate frequency of risky action (repeated–game sessions)

		C treatment			R treatment			F treatment		
		SHH	SHM	SHL	SHH	SHM	SHL	SHH	SHM	SHL
Game ordering:	H-M-L	.995	.983	.973	.304	.139	.643	.461	.493	.780
	M-L-H	.649	.738	.694	.145	.098	.170	.334	.292	.713
	L-H-M	.818	.607	.954	.141	.127	.548	.421	.625	.745
Action ordering:	R-S	.840	.790	.889	.224	.118	.455	.462	.469	.801
	S-R	—	—	—	.176	.133	.555	.360	.428	.706
Overall		.840	.790	.889	$.203^{a}$	$.125^{a}$	$.499^{b}$	$.404^{a}$	$.446^{a}$	$.746^{b}$

a,b: Within a treatment, overall frequencies without superscripts and overall frequencies with superscripts having no letters in common are significantly different at the 5% level or better, with the higher frequency corresponding to the letter coming later in the alphabet. Overall frequencies with superscripts having a letter in common are not significantly different. See text for details.

Consider first the C treatment. As in the first round alone, we see that play of the risky action is more likely overall in the low–payoff Stag Hunt (88.9%) than in the high–payoff Stag Hunt (84.0%), which in turn is more likely than in the medium–payoff Stag Hunt (79.0%). These two order relationships are consistent with certain– and possible–loss avoidance respectively, though there is some variation across game orderings.[20]

We next look at the two limited–information treatments: the R (random matching) treatment and the F (fixed pairs) treatment. The results from both of these treatments are also consistent with

[20]Interestingly, within each game ordering in the C treatment, risky–action choices are always most frequent in the version of Stag Hunt that is played first and least frequent in the version played last; this is the opposite result from that found by Stahl and Van Huyck (2002).

certain–loss avoidance. Risky–action frequency in the R treatment is 49.9% in the low–payoff Stag Hunt versus 20.3% in the high–payoff version, and in the F treatment, the frequencies are 74.6% and 44.6%, respectively. The higher frequency of risky–action play in SHL than in SHH remains when we disaggregate the data in either of these treatments by game ordering or by action ordering. These data also show mixed evidence for possible–loss avoidance. In the R treatment, risky–action play is somewhat higher in the high–payoff Stag Hunt than in the medium payoff Stag Hunt (20.3% versus 12.5%), and this difference is robust to disaggregating the data by game ordering or action ordering. However, risky–action play in the F treatment is actually somewhat lower in SHH than in SHM (40.4% versus 44.6%), though this is not robust to disaggregating the data.

In order to test for significance in the differences found here, we use nonparametric statistical tests. (We will consider parametric statistics in Section 3.4.) Power is an issue here, especially in the C and R treatments, where there is interaction among all subjects in a given session, so that the smallest independent observation is at the session level. In the F treatment, on the other hand, each subject plays an entire game against the same opponent, so we can consider each matched pair of subjects to be an independent observation. Using data from individual pairs in this treatment, we fail to find evidence of possible–loss avoidance, as a one–tailed Wilcoxon summed–ranks test for matched samples is not significant in the direction predicted by possible–loss avoidance ($T^+ = 686.5$, $N = 57$, $p > 0.10$). On the other hand, the difference between SHL and SHH observed in Table 4 is significant ($T^+ = 426$, $N = 59$, $p < 0.001$). Using session–level instead of pair–level data gives results that are broadly similar, despite the small number of sessions: the difference in risky action choices between SHH and SHM is still not significant ($T^+ = 11$, $N = 7$, $p > 0.10$), while the difference between SHL and SHH still is ($T^+ = 0$, $N = 7$, $p \approx 0.007$).

In the C and R treatments, we are forced to use session–level data, as individual pairs of players are not independent observations. There are only 3 sessions of the C treatment, so it is not possible for differences to be significant at conventional levels. In the R treatment, the difference between SHL and SHH is significant ($T^+ = 1$, $N = 6$, $p \approx 0.031$), while the difference between SHH and SHM just misses being significant ($T^+ = 17$, $N = 6$, $p \approx 0.11$). Thus, these results for the R treatment are consistent with both certain– and possible–loss aversion, though the evidence for certain–loss aversion is much stronger.

3.3 Behavior dynamics

We next look at how subject behavior changes over time. Figures 6 and 7 show round–by–round frequencies of risky–action choices in each of the three Stag Hunt games by subjects in the C treatment and R and F treatments, respectively. Figure 6 confirms what we saw in Table 4. There are differences across the three games, and they are in the directions predicted by certain– and possible–loss avoidance, but they are small. We also see that play does not change much over time, though there does seem to be a slight decline in risky–action choices over time in all three games, and a somewhat larger decline in the last ten rounds of SHM.

In the R and F treatments, subjects are not initially told the payoffs to the two strategies, so it is not surprising that in the first few rounds, they play the risky action roughly half the time in all three games in both treatments. Over time, there is some divergence in play across games. In the R treatment, risky–action play in the low–payoff game stays roughly constant (on average) over time, while in the other two games, subjects gradually learn to choose the safe action almost exclusively by the end of the session. This learning is quicker in SHM than SHH, so that over the middle twenty

Figure 6: Frequency of risky action (complete–information treatment, all sessions)

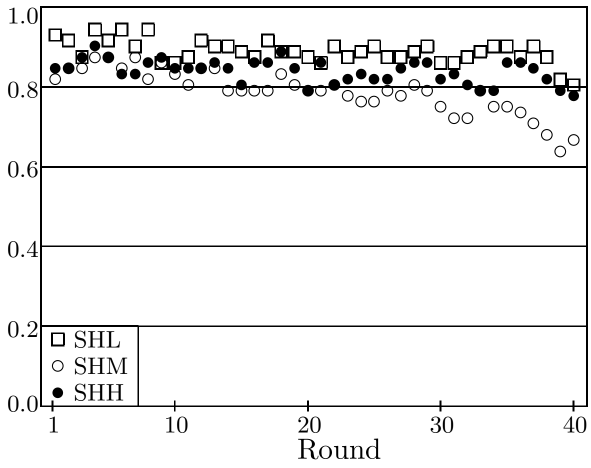

rounds, there is a visible difference in risky–action frequency between these games. In the F treatment, most of the changes in aggregate frequencies occur over the first ten rounds or so. In the high– and medium–payoff games, risky–action frequencies fall slightly—to about 40% for SHH and about 45% for SHM—and remain there for the remainder of the session; in the low–payoff game, the risky–action frequency rises to about 80%, then remains roughly constant.[21]

Figure 7: Frequency of risky action (limited–information treatments, all sessions)

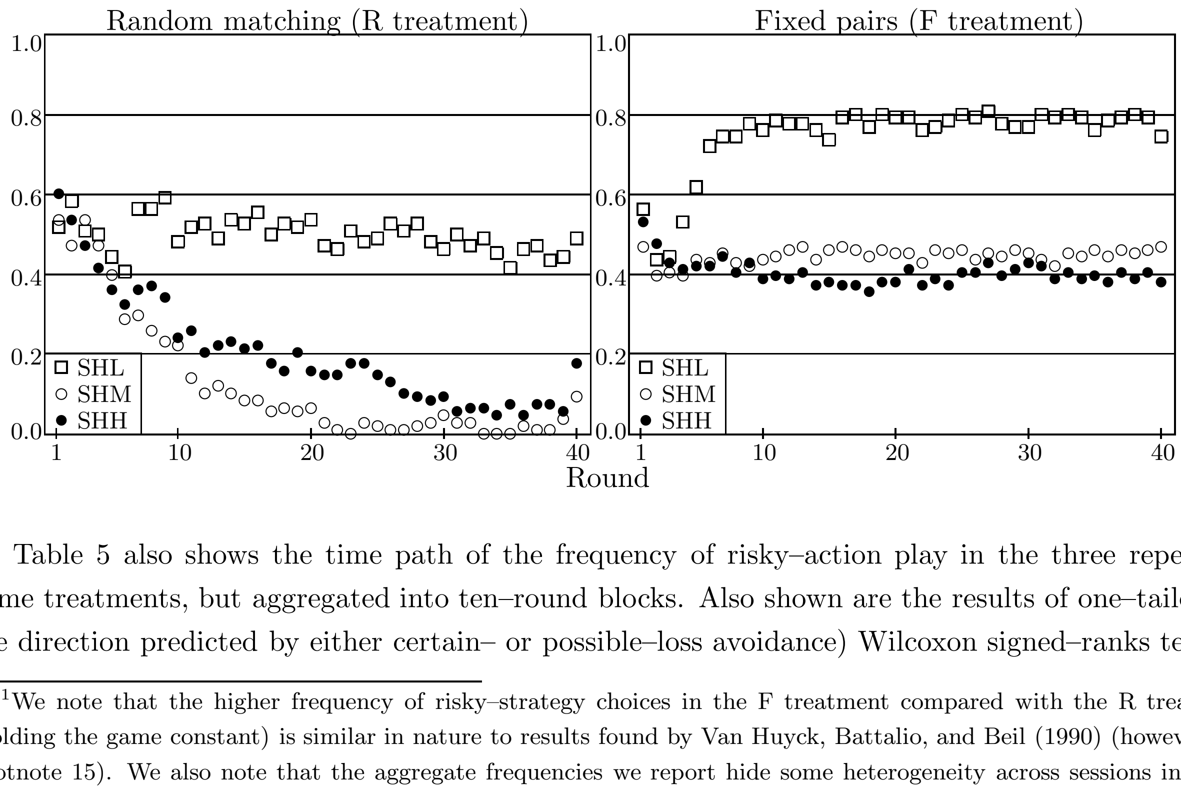

Table 5 also shows the time path of the frequency of risky–action play in the three repeated–game treatments, but aggregated into ten–round blocks. Also shown are the results of one–tailed (in the direction predicted by either certain– or possible–loss avoidance) Wilcoxon signed–ranks tests of

[21]We note that the higher frequency of risky–strategy choices in the F treatment compared with the R treatment (holding the game constant) is similar in nature to results found by Van Huyck, Battalio, and Beil (1990) (however, see Footnote 15). We also note that the aggregate frequencies we report hide some heterogeneity across sessions in the R treatment and individual pairs in the F treatment since, as one might expect, some sessions or pairs tend toward the (R,R) action pair while others tend toward the (S,S) pair. In each of the three versions of Stag Hunt in the R treatment, the frequency of *coordination* on either of these pure–strategy Nash equilibria rises steadily from slightly over one–half in the first ten rounds to 80% or higher in the last ten rounds. In the F treatment, coordination in all three games rises quickly over the first ten rounds from about one–half to over 80%, and is steadily above 90% throughout most of the last twenty rounds of each game.

significant differences between SHH and either SHM or SHL. These results back up what we saw in

Table 5: Frequency of risky action by 10–round blocks (repeated–game treatments)

Rounds	C treatment			R treatment			F treatment		
	SHH	SHM	SHL	SHH	SHM	SHL	SHH	SHM	SHL
1–10	.860	.850	.910	.403	.371*	.517	.436	.427	.635***
11–20	.846	.810	.893	.205	.087**	.524**	.381	.456	.780***
21–30	.835	.782	.883	.130	.019**	.493**	.402	.450	.783***
31–40	.818	.717	.871	.073	.022	.464**	.396	.450	.787***

* (**,***): *Significantly different from SHH at the 10% (5%, 1%) level (one–tailed Wilcoxon signed–ranks test, session–level data for C and R treatments, individual–pairs data for F treatment).*

Figures 6 and 7. As in the previous section, we have too few C treatment sessions for us to find significance. In the R treatment, we also must use session–level data. Differences between SHH and SHM are significant in each ten–round block except for the last one, while differences between SHH and SHL are significant in all but the first. In the F treatment, we are able to use data from individual pairs of subjects. No differences between SHH and SHM are significant, but differences between SHH and SHL are significant—and in the direction predicted by certain–loss avoidance—for all four ten–round blocks.

3.4 Parametric statistics

To this point, we have seen reasonably strong evidence consistent with certain–loss avoidance, with somewhat weaker evidence for possible–loss avoidance. However, the statistical tests we have been using have thus far been nonparametric tests, which tend to err on the conservative side (particularly when session–level data are used). In order to gain some power in our tests, as well as to make sure we've controlled for other factors that may be affecting behavior, we next look at results from several probit regressions. (Logit regressions—not reported here—gave similar results.) In each of our regressions, we take the dependent variable to be an indicator for current–round action choice (1=risky choice, 0=safe choice). In order to determine the effect of payoff levels (and thus test for certain– and possible–loss avoidance), we use as our primary explanatory variables indicators for the SHM and SHL games; in addition, we include as variables the products of these indicators with the round number and the square of the round number, to allow for the possibility that the effects of payoff–level changes develop (or, possibly, disappear) over time. As controls, we include variables for the round number itself and its square, indicators for two of the three game orderings (M-L-H and L-M-H) and one of the two action orderings (R-S). The probit regressions were performed using Stata (version 10). We estimate coefficients separately for each of the three multiround treatments (C, R, and F), and in each case we use individual–subject random effects.

Coefficients and standard errors (in parentheses) are shown in Table 6; this figure also shows log likelihoods and pseudo–R^2 values for both models.[22] Before examining their implications for loss avoidance, we mention some of their other interesting implications. Subject behavior is nonstationary

[22]The pseudo–R^2 values were computed by rescaling the log–likelihoods into [0,1], such that a model with no right–hand–side variables other than the constant term maps to zero, and a perfect fit maps to one.

Table 6: Results of probit regressions with random effects (std. errors in parentheses)

Dependent variable: risky–action choice (round t)	C treatment ($N = 8640$)	R treatment ($N = 12960$)	F treatment ($N = 15120$)
constant	2.769***	0.570***	–0.130
	(0.240)	(0.114)	(0.153)
SHM	–0.088	0.280***	–0.161*
	(0.159)	(0.097)	(0.086)
SHM · Round	0.000	–0.087***	0.031***
	(0.018)	(0.013)	(0.010)
SHM · Round^2	–0.0004	0.0017***	–0.0006***
	(0.0004)	(0.0003)	(0.0002)
p–value (joint significance of three SHM variables)	< 0.001***	< 0.001***	< 0.001***
SHL	0.348	–0.152*	0.064
	(0.170)	(0.091)	(0.086)
SHL · Round	0.002	0.098***	0.113***
	(0.019)	(0.010)	(0.010)
SHL · Round^2	–0.0001	–0.0012***	–0.0022***
	(0.0004)	(0.0002)	(0.0002)
p–value (joint significance of three SHL variables)	< 0.001***	< 0.001***	< 0.001***
Round	0.000	–0.094***	–0.024***
	(0.013)	(0.008)	(0.007)
Round^2	–0.0002	0.0012***	–0.0005***
	(0.0003)	(0.0002)	(0.0002)
M-L-H game ordering	–1.977***	–0.911***	–0.583***
	(0.303)	(0.132)	(0.164)
L-H-M game ordering	–1.430***	–0.354***	0.113
	(0.290)	(0.111)	(0.185)
R-S action ordering	—	–0.071	0.382**
		(0.098)	(0.158)
–ln(L)	2362.048	5441.483	7767.587
pseudo–R^2	0.053	0.233	0.119

* (**,***): *Coefficient significantly different from zero at the 10% (5%, 1%) level.*

in all treatments; the coefficients for the six variables containing either the round number or its square are jointly significant in all cases (likelihood–ratio test, $p < 0.001$ for all three treatments). In addition, behavior is sensitive to the order in which the games are played: in each of the three treatments, coefficients for the game–ordering indicator variables M-L-H and L-H-M are jointly significant ($p < 0.001$ for all three treatments), and taken individually, they are nearly always significantly different from zero (as shown in the table) and are always significantly different from each other ($p < 0.10$ for the C treatment, $p < 0.001$ for the R and F treatments). There is also some evidence that the order in which the actions are shown in the payoff matrix matters, as the coefficient for R-S is significant (at the 5% level) in the F treatment, though it is not significant in the R treatment.

We next look at the results connected with loss avoidance. In these results, we continue to find solid support for certain–loss avoidance, but we now find substantial support for possible–loss avoidance as well. Consider first certain–loss avoidance, which implies that the coefficients of the three SHL variables (SHL, SHL · Round, and SHL · Round^2) will be different from zero. In the C treatment, the individual coefficients are not significantly different from zero, but the three variables are jointly significant ($\chi^2_3 = 36.46, p < 0.001$). In both the R and F treatments, the results are even stronger, as either two or all three of the individual variables are significant at the 1% level, and the three variables are jointly significant ($\chi^2_3 = 1101.94, p < 0.001$ for the R treatment, $\chi^2_3 = 1590.14, p < 0.001$ for the F treatment).

Similarly, possible–loss avoidance implies that the coefficients of the three SHM variables (SHM,

SHM $\cdot$ Round, and SHM $\cdot$ Round^2) will be different from zero. As with the SHL variables, the SHM variables in the C treatment are individually not significant, but jointly are significant ($\chi^2_3 = 50.64, p < 0.001$). In the R and F treatments, two of the three SHM variables are individually significant, and the three are jointly significant ($\chi^2_3 = 182.67, p < 0.001$ for the R treatment, $\chi^2_3 = 36.16, p < 0.001$ for the F treatment).

However, these joint significance tests are unsatisfactory for two reasons. First, these sets of variables contain interaction terms, and it is well known that for nonlinear regression models (such as probits), the marginal effect of the interaction between two variables is not equal to the coefficient of the interaction term. (See Ai and Norton (2003) and Norton, Wang, and Ai (2004).) Second, even if we find that a particular set of variables is jointly significant, this alone does not tell us the actual effect of these variables; for example, we have seen that the variables SHM, SHM $\cdot$ Round, and SHM $\cdot$ Round^2 are jointly significant in the R treatment, but we do not yet know whether C choices Invest choices are more likely, less likely, or even equally likely in SHM compared to SHH. The total effect in round t of the SHM game, rather than the SHH game, on the argument of the normal c.d.f. used in the probit model is given by $\beta_{\text{SHM}} + \beta_{\text{SHM}\cdot\text{Round}} \cdot t + \beta_{\text{SHM}\cdot\text{Round}^2} \cdot t^2$ (where β_Y is the coefficient of the variable Y). So, the incremental effect (the analog to a marginal effect, for a discrete variable) of the SHM game rather than the SHH game in round t has the form

$$\Phi\left(\bar{X} \cdot B + \beta_{\text{SHM}} + \beta_{\text{SHM}\cdot\text{Round}} \cdot t + \beta_{\text{SHM}\cdot\text{Round}^2} \cdot t^2\right) - \Phi\left(\bar{X} \cdot B\right), \tag{1}$$

where $\bar{X}$ is the row vector of the other right–hand–side variables' values, and B is the column vector of their coefficients. If subjects exhibit possible–loss avoidance in a particular round, the sign of this expression should be negative for that value of t; if they exhibit certain–loss avoidance, the sign of a corresponding expression using the SHL variables should be positive for that value of t. In the following, we will use the mean conditional on an appropriate subsample as values for $\bar{X}$.

In Figure 8, we graph all six of these expressions (versions of Equation 1 for SHL and SHM variables and for C, R, and F treatments). This figure shows, for each expression, the corresponding point estimates and 90% confidence intervals for each value of t (that is, the round number) from 1 to 40—the number of rounds of each game in the experiment.[23] Again, we see strong evidence for certain–loss avoidance. In all rounds of all treatments, except for Round 1 of the R treatment, the point estimate for the marginal effect of SHL is positive, and the 90% confidence interval lies entirely above zero (so that a one–sided hypothesis test at the 5% level would yield significance) for Rounds 3–38 of the C treatment, Rounds 3–40 of the R treatment, and all rounds of the F treatment. We also see evidence for possible–loss avoidance, though not in the F treatment, where the point estimate of the marginal effect of SHM has the wrong sign in all but the first five rounds, and is never significant in the predicted direction. (This is consistent with the effect apparent in Figure 7.) In the C treatment, the marginal effect of the SHM variables is negative and significant in Rounds 11–40, and in the R treatment, the marginal effect of SHM is negative and significant in Rounds 6–40.

These time paths show a consistent pattern. Each one starts close to zero in Round 1 and, from there, moves away from zero. Five of the six time paths move in the direction predicted by certain– or possible–loss avoidance (the exception being the SHM variables in the F treatment). In one case (SHM in the C treatment), the time path continues to move away from zero, while in the other four

[23]Note that we use 90% confidence intervals rather than the usual 95% confidence intervals here. Since certain– and possible–loss avoidance make directional predictions, our rejection regions are one–tailed. Use of 90% confidence intervals gives us rejection regions of 5% on each side.

Figure 8: Estimates of the effect of game on risky–strategy choice (vs. SHH) based on Table 6 results

(Circles represent point estimates; line segments represent 90% confidence intervals)

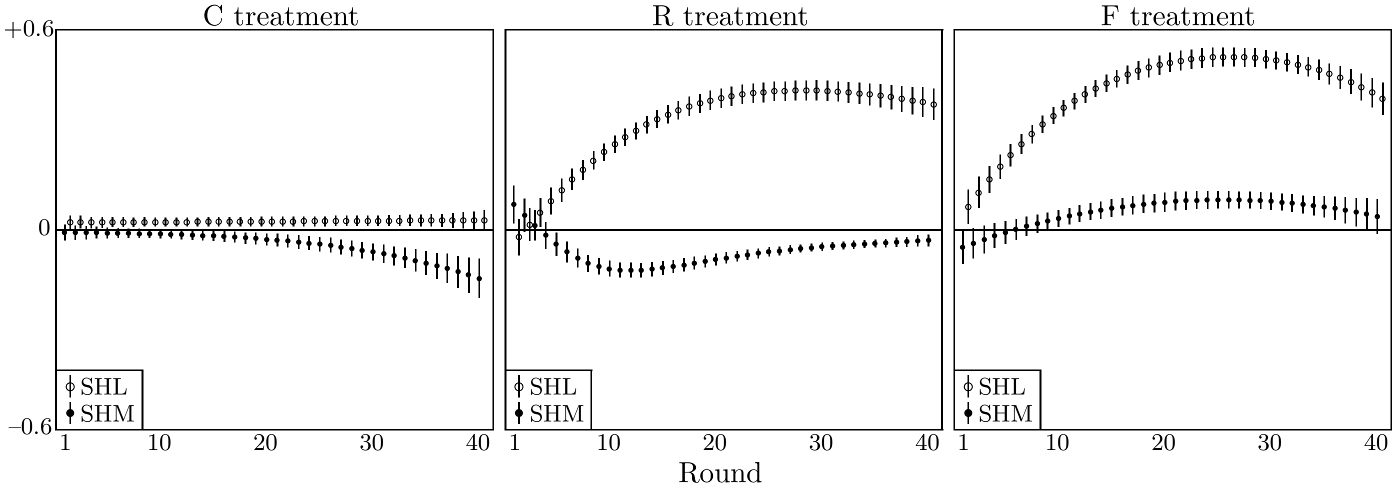

cases it turns back toward zero, though never crossing the horizontal axis.[24] In other words, the effects of loss avoidance are greatest after subjects have some experience in playing a game—not at the beginning—and this is true both when they know the game they're playing (C treatment) and when they have to learn it (F and R treatments).

3.5 Summary and discussion

In the sections above, we saw that subjects' choices are consistent with certain–loss avoidance in all four of our treatments. We also found that behavior is consistent with possible–loss avoidance in the R treatment and overall in the C treatment, but not in the F treatment or initially in either the C or the O treatment. Compiling the significant results we found, then, we see that there is evidence for both certain– and possible–loss avoidance in our data, though the evidence for certain–loss avoidance is somewhat stronger—it is present in more treatments, and the size of its effect seems to be larger.

Several points are worth making at this stage. First, we note that some of our results can be put into the context of previous results from the literature discussed in Sections 2.1 and 2.2. Consistent with many others' results, a comparison of our results in the C and R treatments (Figures 6 and 7) suggests that giving subjects more information leads to more play of the risky action, and hence greater likelihood of the payoff–dominant outcome. We also found that play of the risky action was more likely when subjects played repeatedly with the same opponent (our F treatment) than when opponents changed from round to round (our R treatment). This is similar to the result found by Clark and Sefton (2001), but not that of Schmidt et al. (2003) who found no systematic differences between the two. We hasten to point out that both Clark and Sefton (2001) and Schmidt et al. (2003) implemented these matching mechanisms in ways different enough from ours (their subjects had complete payoff information, while ours had limited information; their games lasted either 10 or 8 rounds, while ours lasted 40; etc.) that our results should not be seen as a replication of the former or a contradiction of the latter. However, the increase in risky–action play we saw in the first round of our

[24] It is possible that the shapes of the trajectories in Figure 8 could be a direct consequence of the quadratic effect of the round number we've been assuming in this section. In order to evaluate this possibility, we also considered probit models that included the cubic terms SHM $\cdot$ Round^3, SHL $\cdot$ Round^3, and Round^3, and reproduced this figure using cubic expressions rather than quadratics. Results were very similar to those in Figure 8; in particular, the apparent critical points were not substantially affected when the cubic terms were present.

complete–information treatment (which corresponds most closely to Schmidt et al.'s random–match treatment), compared with our one–shot treatment, does stand in contrast to Schmidt et al.'s finding of similar behavior in the two treatments.

In our R and F treatments, our results have a similar flavor to those of Erev and colleagues; in their individual–decision problems under conditions with even less information, they also found that behavior is sensitive to changes in payoff levels that affect the signs of payoffs (though not so much to payoff–level changes in general). Additionally, we note that the results from our O and C treatments are in line with those of Rydval and Ortmann (2005). As mentioned previously, they tested for certain–loss avoidance and found weak evidence in favor of it: significant in one pair of games, not significant in the other pair. The SHH and SHL games in our O treatment represent the part of our design closest to theirs, and we found in these results a significant effect consistent with certain–loss avoidance. Arguably, the SHH and SHL games in the first–round of our C treatment are the next–closest to their design, and we found in those results a small but insignificant effect in the direction predicted by certain–loss avoidance. Finally, Cachon and Camerer (1996) found that loss avoidance is more observable in subjects' beliefs about their opponents than in their own behavior. We did not specifically test for this, but consistent with their conclusion, our result that loss avoidance is much stronger after subjects have gained experience than initially suggests that at least some of the effect we found might be due to subjects' reacting to differences across games in opponent strategies, however these differences originally came about. (An alternative explanation is that subjects gradually show their loss avoidance as they begin to understand the strategic structure of the game they are playing.)

Second, we speculate a bit more as to what our results say about the nature of loss avoidance. One major difference between our experiment and Rydval and Ortmann's is that we include treatments with repeated play of the same game. By and large, differences in play across games start out small, but grow over time (at least up to a point) in the directions predicted by certain– and possible–loss avoidance. This suggests that loss avoidance due to introspection is relatively minor (indeed, in the case of possible–loss avoidance, we found no evidence of it whatsoever in first–round choices); rather, it is typically learned. This "learning to avoid losses" was found in all three of our repeated–game treatments. In the C treatment, it could be argued that many of the subjects were not learning loss avoidance per se, but perhaps that a small portion of them exhibit loss avoidance initially, and others are simply learning to play as loss avoidance predicts, due to the strategic complementarities in these games. (Recall the discussion of this issue in Section 2.1.) However, in the R and F treatments, there can be no loss avoidance due to introspection (since subjects are not initially told the payoffs of the game, but must infer them), but differences across games develop anyway, so it would be harder to argue against the claim that at least some subjects are learning to avoid losses—even if others are simply learning to choose best replies to them.

Next, it is worthwhile to mention the significance of our results for equilibrium selection. Much of the literature on Stag Hunt games, and order–statistic games in general, has been motivated by questions about equilibrium selection—more specifically, under what circumstances it is reasonable to expect subjects to successfully coordinate on the payoff–dominant Nash equilibrium ((R,R) in our games) rather than the inefficient secure (and in our games, risk–dominant) equilibrium ((S,S) in our games). Nearly all experimental studies of these games have used versions with only positive payoffs, and have found substantial play of both risky and safe actions. Our results show that using games with some, or mostly, negative payoffs would have an impact on the equilibrium we should expect to see. In any order–statistic game for which certain–loss avoidance has some predictive value (i.e., at least one, but not every, action guarantees a negative payoff) it will rule out safer actions that are part of lower–payoff Nash equilibria, in favor of riskier actions that are part of higher–payoff equilibria. In the extreme case where certain–loss avoidance makes a unique prediction, this prediction will be the payoff–dominant Nash equilibrium. So, to the extent that individuals actually exhibit certain–loss avoidance, this will tend to make the payoff–dominant outcome more likely. On the other hand, in any order–statistic game for which possible–loss avoidance has some predictive value (i.e., at least one, but not every, action earns a possible negative payoff), it will do the opposite: rule out riskier actions that are part of higher–payoff equilibria, in favor of safer actions that are part of lower–payoff equilibria. So, to the extent that individuals actually exhibit possible–loss avoidance, this will tend to make the payoff–dominant outcome less likely. In the extreme case where possible–loss avoidance makes a unique prediction, this prediction will be the lowest–payoff pure–strategy Nash equilibrium.

Finally, we wish to point out some limitations of our results. First, while our results are consistent with loss avoidance in nearly all cases, we can't be certain that loss avoidance is the cause of the effects we have seen. (For example, in Appendix B we show that an individual whose decision making is governed by prospect theory could make a similar pattern of choices, though different—or even opposite—patterns could also be observed.) Further research is necessary to confirm that loss avoidance is actually what is driving these results. Second, the size of the effect we have found is relatively small. Notice in Table 4, for example, that the variation in risky–action frequencies is higher across treatments than across games. Similarly, the probit results in Table 6 and Figure 8 imply that the size of the effect of changing the game from SHH to SHM or SHL is rather less than that of, for example, changing the ordering of the games. Third, our results are not uniformly substantial (or significant) across treatments; it is not clear why there should be no evidence of possible–loss avoidance in the F treatment, in which the evidence for certain–loss avoidance is strongest. Fourth, as we have already mentioned, the case for possible–loss avoidance is a bit weaker than that for certain–loss avoidance.

However, set against these limitations, we point out that our experiment is actually a conservative test of loss avoidance in several ways. First, even though the SHM and SHL payoff matrices show the potential for losses, it was impossible for subjects in the experiment to actually incur losses overall, once the showup fee and other payments are taken into account. Thus, any losses are the result of framing: it is hoped that subjects internalize the showup fee (in the C, R, and F treatments) or other payments (in the O treatment). To the extent that they do not do so, even subjects who do exhibit loss avoidance will not show it in the experiment, so that the actual prevalence of loss avoidance could be higher than we found. Second, it is worth mentioning that an individual who exhibits loss avoidance may not show it for these particular games; for example, if such an individual chooses the risky action over the safe action in SHM, that person would prefer the risky action even more in SHH and SHL, but the data would show no difference across the three games for that person. Without a way of showing intensity of preferences, the observed evidence of loss avoidance will understate the true level.

4 Conclusion

According to game theory, changing the level of payoffs in a game (by adding a constant to each of them) does not change any strategic aspect of the game, and hence does not change equilibrium predictions—as long as payoffs represent players' actual expected utilities. However, there is a substantial body of evidence from individual decision–making experiments that changes in the level of *money* payoffs can affect behavior. There has been less research into the effect of changing payoff levels in games, but what evidence there is—along with inference from the decision–making studies that have been conducted—suggests that it is quite possible for there to be an effect in games too. In order to ascertain the effects of changing payoff levels in games, we have run a human–subjects experiment with three versions of Stag Hunt. These games are similar in that each can be derived from either of the others by addition of a constant to all of the payoffs. Importantly, however, the three games differ in the signs of the payoffs. In the high–payoff (SHH) game, both safe and risky actions can only lead to positive payoffs. In the medium–payoff (SHM) game, the safe action leads to a positive payoff, but the risky action can lead to either a positive or a negative payoff. In the low–payoff (SHL) game, the risky action again can lead to either a positive or a negative payoff, but the safe action leads to a negative payoff.

We consider two particular ways in which changes in payoff levels—specifically, the sign of payoffs—can affect action choices in these games. "Certain–loss avoidance" refers to the tendency to avoid an action yielding a certain loss in favor of another available action that might yield a gain. Subjects exhibiting certain–loss avoidance will be *more* likely to choose the risky action in SHL than in SHH. "Possible–loss avoidance" refers to the tendency to avoid an action that might yield a loss in favor of another available action that yields a certain gain. Subjects exhibiting possible–loss avoidance will be *less* likely to choose the risky action in SHM than in SHH. Both types of loss avoidance are consistent with, but neither entailing nor implied by prospect theory; in particular, loss avoidance is distinct from the component of prospect theory known as loss aversion.

The experiment we conduct allows us to look for certain– and possible–loss avoidance manifested in several ways. In one of our treatments, subjects play each game only once, with complete infomation about the games' payoffs. In a second treatment, they again have complete information, but play each game repeatedly (forty times) against changing opponents. In a third treatment, they still play repeatedly against changing opponents, but are given only limited information about the strategic environment they are in. In a fourth treatment, they play repeatedly with the same limited information, but against the same opponent in all rounds of a game. This design allows us to detect loss avoidance arising from introspection, based on responses to the play of others, and due to the combination of the two.

Our results show support for both certain– and possible–loss avoidance. These phenomena manifest themselves in two ways. We find that when subjects play games only once—but with complete information about the payoffs of the games—their choices are consistent with certain–loss avoidance (that is, significantly more risky–action play in SHL than SHH), but not possible–loss avoidance (that is, not significantly less risky–action play in SHM than SHH). When subjects play games repeatedly, they tend to exhibit certain–loss avoidance, possible–loss avoidance, or both. Moreover, the magnitude of their loss avoidance doesn't die out as they gain more experience; rather, it increases substantially over time for at least the first half or so of the experimental session (though it sometimes decreases in later rounds). We conclude that both types of loss avoidance may be real factors in decision making, and though more research in this area is needed, we should begin to take loss avoidance into account when considering situations in which both gains and losses are possible.

References

Ai, C.R. and E.C. Norton (2003), "Interaction terms in logit and probit models," *Economics Letters* 80, pp. 123–129.

Barkan, R., D. Zohar, and I. Erev (1998), "Accidents and decision making under uncertainty: a comparison of four models," *Organizational Behavior and Human Decision Processes* 74, pp.

118–144.

Battalio, R.C., L. Samuelson, and J.B. Van Huyck (2001), "Optimization incentives and coordination failure in laboratory stag–hunt games," *Econometrica* 69, pp. 749–764.

Bereby–Meyer, Y. and I. Erev (1998), "On learning to become a successful loser: a comparison of alternative abstractions of learning processes in the loss domain," *Journal of Mathematical Psychology* 42, pp. 266–286.

Berninghaus, S.K. and K.-M. Ehrhart (2001), "Coordination and information: recent experimental evidence," *Economics Letters* 73, pp. 345–351.

Blume, A. and A. Ortmann (2007), "The effects of costless pre–play communication: experimental evidence from games with Pareto–ranked equilibria," *Journal of Economic Theory* 132, pp. 274–290.

Buchheit, S. and N. Feltovich (2008), "Experimental evidence of a sunk-cost paradox: a study of pricing behavior in Bertrand–Edgeworth duopoly," working paper, University of Aberdeen.

Cachon, G.P. and C.F. Camerer (1996), "Loss avoidance and forward induction in experimental coordination games," *Quarterly Journal of Economics* 111, pp. 165–194.

Camerer, C. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, Princeton, NJ.

Carlsson, H. and E. van Damme (1993), "Global games and equilibrium selection," *Econometrica* 61, pp. 989–1018.

Charness, G. and B. Grosskopf (2004), "What makes cheap talk effective? Experimental evidence," *Economics Letters* 83, pp. 383–389.

Clark, K., S. Kay, and M. Sefton (2001), "When are Nash equilibria self–enforcing? An experimental analysis," *International Journal of Game Theory* 29, pp. 495–515.

Clark, K. and M. Sefton (2001), "Repetition and signalling: experimental evidence from games with efficient equilibria," *Economics Letters* 70, pp. 357–362.

Cox, J.C. and V. Sadiraj (2006), "Implications of small– and large–stakes risk aversion for decision theory," *Games and Economic Behavior* 56, pp. 45–60.

Devetag, G. (2003), "Coordination and information in critical mass games: an experimental study," *Experimental Economics* 6, pp. 53–73.

Devetag, G. and A. Ortmann (2007), "When and why? A critical survey on coordination failure in the laboratory," *Experimental Economics* 10, pp. 331–344.

Duffy, J. and N. Feltovich (2002), "Do actions speak louder than words? Observation vs. cheap talk as coordination devices," *Games and Economic Behavior* 39, pp. 1–27.

Duffy, J. and N. Feltovich (2006), "Words, deeds and lies: strategic behavior in games with multiple signals," *The Review of Economic Studies*, 73, pp. 669–688.

Erev, I., Y. Bereby–Meyer, and A.E. Roth (1999), "The effect of adding a constant to all payoffs: experimental investigation and implications for reinforcement learning models," *Journal of Economic Behavior and Organization* 39, pp. 111–128.

Falkinger, J., E. Fehr, S. Gächter, and R. Winter–Ebmer (2000), "A simple mechanism for the efficient provision of public goods: experimental evidence," *American Economic Review* 90, pp. 247–264.

Feltovich, N., A. Iwasaki, and S.H. Oda (2008), "The effect of matching mechanism on learning in repeated 2x2 games with limited information," working paper, University of Aberdeen.

Fischbacher, U. (2007), "z–Tree: Zurich toolbox for ready–made economic experiments," *Experimental Economics* 10, pp. 171–178.

Harsayni, J.C. and R. Selten (1988), *A general theory of equilibrium selection in games*, MIT Press, Cambridge, MA.

Haruvy, E. and D.O. Stahl (2004), "Deductive versus inductive equilibrium selection: experimental results," *Journal of Economic Behavior and Organization* 53, pp. 319–331.

Heinemann, F., R. Nagel, and P. Ockenfels (2004), "Measuring strategic uncertainty in coordination games," working paper #804, Universitat Pompeu Fabra.

Holt, C.A. and S.K. Laury (2002), "Risk aversion and incentive effects," *American Economic Review* 92, pp. 1644–1655.

Kachelmeier, S.J. (1996), "Do cosmetic reporting variations affect market behavior?" *Review of Accounting Studies*, 1, pp. 115–140.

Kahneman, D. and A. Tversky (1979), "Prospect theory: an analysis of decision making under uncertainty," *Econometrica* 47, pp. 273–297.

Manzini, P., A. Sadrieh, and N.J. Vriend (2002), "On smiles, winks and handshakes as coordination devices," Dept. of Economics Working Paper No. 456, Queen Mary, University of London.

Myers, J.I., J.G. Fort, L. Katz, and M.M. Suydam (1963), "Differential monetary gains and losses and event probability in a two–choice situation," *Journal of Experimental Psychology* 66, pp. 521–522.

Norton, E.C., H. Wang, and C.R. Ai (2004), "Computing interaction effects and standard errors in logit and probit models," *STATA Journal* 4, pp. 154–167.

Ochs, J. (1995), "Coordination games," in *Handbook of Experimental Economics*, J.H. Kagel and A.E. Roth, eds., Princeton University Press, Princeton, NJ., pp. 195–252.

Rabin, M. and R.H. Thaler (2001), "Anomolies: risk aversion," *Journal of Economic Perspectives* 15, pp. 219–232.

Rankin, F., J.B. Van Huyck, and R.C. Battalio (2000), "Strategic similarity and emergency conventions: evidence from payoff perturbed Stag Hunt games," *Games and Economic Behavior* 32, pp. 315–337.

Rapoport, A. and R.B. Boebel (1992), "Mixed strategies in strictly competitive games: a further test of the minmax hypothesis," *Games and Economic Behavior* 4, pp. 261–283.

Rousseau, J.-J. (1973), "A discourse on the origin of inequality," in *The Social Contract and The Discourses* (translated by G.D.H. Cole), J.M. Dent & Sons, London, pp. 27–113.

Rydval, O. and A. Ortmann (2005), "Loss avoidance as selection principle: evidence from simple stag–hunt games," *Economics Letters* 88, pp. 101–107.

Schelling, T. (1960), *The Strategy of Conflict*, Harvard University Press, Cambridge, MA.

Schmidt, D. R. Shupp, J.M. Walker, and E. Ostrom (2003), "Playing safe in coordination games: the roles of risk dominance, payoff dominance, and history of play," *Games and Economic Behavior* 42, pp. 281–299.

Selten, R. (1995), "An axiomatic theory of a risk dominance measure for bipolar games with linear incentives," *Games and Economic Behavior* 8, pp. 213–263.

Siegel, S. and N.J. Castellan, Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences*, McGraw–Hill, New York.

Siegel, S., and D.A. Goldstein (1959), "Decision–making behavior in a two choice uncertain outcome situation," *Journal of Experimental Psychology* 57, pp. 37–42.

Siegel, S., A.E. Siegel, and J.M. Andrews (1964), *Choice, Strategy, and Utility*, McGraw–Hill, New York.

Stahl, D.O. and J. Van Huyck (2002), "Learning conditional behavior in similar Stag Hunt games," working paper.

Thaler, R.H. and E.J. Johnson (1991), "Gambling with the house money and trying to break even: the effects of prior outcomes on risky choice," *Management Science* 36, pp. 643–660.

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

Tversky, A. and D. Kahneman (1992), "Advances in prospect theory: cumulative representation of uncertainty," *Journal of Risk and Uncertainty* 5, pp. 297–323.

Van Huyck, J.B., R.C. Battalio, and R.O. Beil (1990), "Tacit coordination games, strategic uncertainty, and coordination failure," *American Economic Review* 80, pp. 234–248.

Van Huyck, J.B., R.C. Battalio, and R.O. Beil (1991), "Strategic uncertainty, equilibrium selection, and coordination failure in average–opinion games," *Quarterly Journal of Economics* 106, pp. 885–910.

Waller, W.S., B. Shapiro, and G. Sevcik (1999), "Do cost–based pricing biases persist in laboratory markets?" *Accounting, Organizations, and Society* 24, pp. 717–739.

Appendix A: instructions given to subjects

Below is a annotated translation of the instructions given to subjects in this experiment. Copies of the actual instructions are available from the corresponding author upon request.

1 Introduction

Thank you for your participation in an experiment in the economics of decision-making. If you follow these instructions carefully and make good decisions you might earn a considerable amount of money that will be paid to you in cash at the end of the session.

2 Sequence of Play in a Round

This experimental session consists of six different games. Each game consists of one round [forty rounds in the C, R, and F treatments]. At the start of each game, you will be randomly assigned a player type, either "row player" or "column player." Your type will not change during the course of game. In each game of this experiment you will be randomly matched to a player of the opposite type. You will be matched with a different player in every round [every forty rounds in the F treatment] of a game. We will refer to the person you are paired with in a round as your "partner." Your score in each round will depend on your choice and the choice of your partner in that round. You will not know the identity of your partner in any round, even after the end of the session.

- At the beginning of each game, the computer program randomly matches each player to a partner.
- You and your partner play the game. Figure 1 is displayed on your screen. If you are a row player, you choose which row of the payoff table to play, R1 or R2. If you are a column player, you choose which column of the payoff table to play, C1 or C2.
- After all players have chosen actions, your action, your partner's action, and your payoff or score are displayed. Your score is determined by your action and the action of your partner according to the given payoff table.
- *[This part is replaced according to the treatment.]*
  - **O treatment:** Provided that the last game has not been reached, a new game will then begin, that is, the payoff table will change. You will be matched with a different partner in the new round.
  - **C and R treatment:** Provided that the last round of the game has not been reached, a new round of the same game will then begin. You will be matched with a different partner in the new round.
  - **F treatment:** Provided that the last round of the game has not been reached, a new round of the same game will then begin. You will be matched with the same partner in the new round. Nevertheless, at every forty rounds, your partner and the payoff table will be changed.
- **Notice that you must not record any results of the games.** If the experimenter find you are recording them, you cannot continue your experiment. In that case, you will not be paid for this experiment.

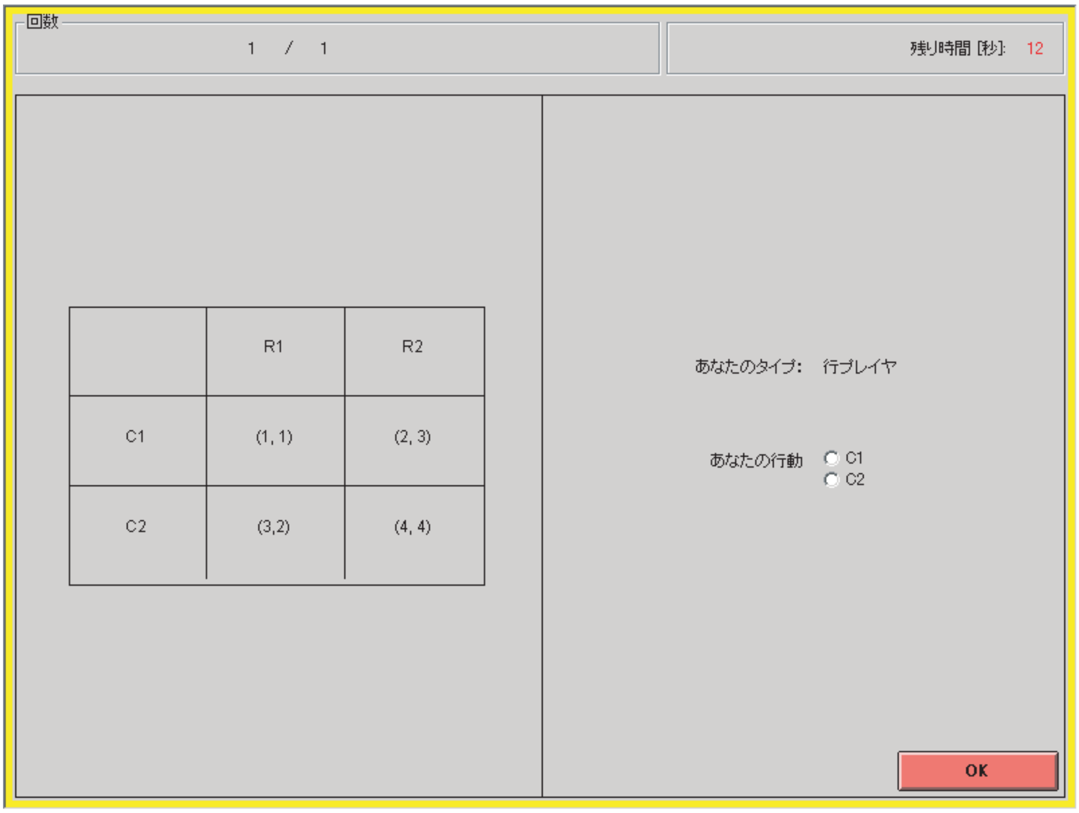

Figure 1: The screen when you are a column player [The payoff table on the left side will not be shown in the R and F treatments].

3 The payoff tables

The payoff table for each game you play will be shown on your computer screen. [This part is replaced with the following sentences in the R and F treatments: The payoff table for each game you play will not be shown on your computer screen. However, let us explain the payoff table to support your decision-making.] In every round of a game, both you and your partner have a choice between two possible actions. If you are designated as the row player, you must choose between actions R1 and R2. If you are designated as the column player, you must choose between actions C1 and C2. Your action, together with the action chosen by your partner, determines one of the four boxes in the payoff table. In each box, the first number represents your score and the second number represents your partner's score.

4 Payments

If you complete this experiment, the computer screen will reveal your score, the round that you got the score, and the payment you obtain. The round will be randomly chosen from all rounds you played. The payment will be calculated from your score as 200 yen for each point in that round. [The following sentence is added in the C, R, and F treatment: In addition, you will be paid a show-up fee of 3000 yen.]

Are there any questions before we begin?

Appendix B: a note on the distinction between loss avoidance and prospect theory

At this point, it is worthwhile to take a moment to look at the similarities and differences between loss avoidance and Kahneman and Tversky's (1979) prospect theory, not least because one component of prospect theory—loss aversion—has a name so similar to loss avoidance as to possibly lead to confusion. Prospect theory comprises several components, one of which is a value function that characterizes the relationship between changes in an individual's wealth from some reference point (which can be sensitive to framing effects, so it may, but need not, be the individual's current wealth level) and changes in that individual's well–being. A sample value function is shown in Figure 9.

Figure 9: A sample prospect–theory value function (adapted from Kahneman and Tversky (1979))

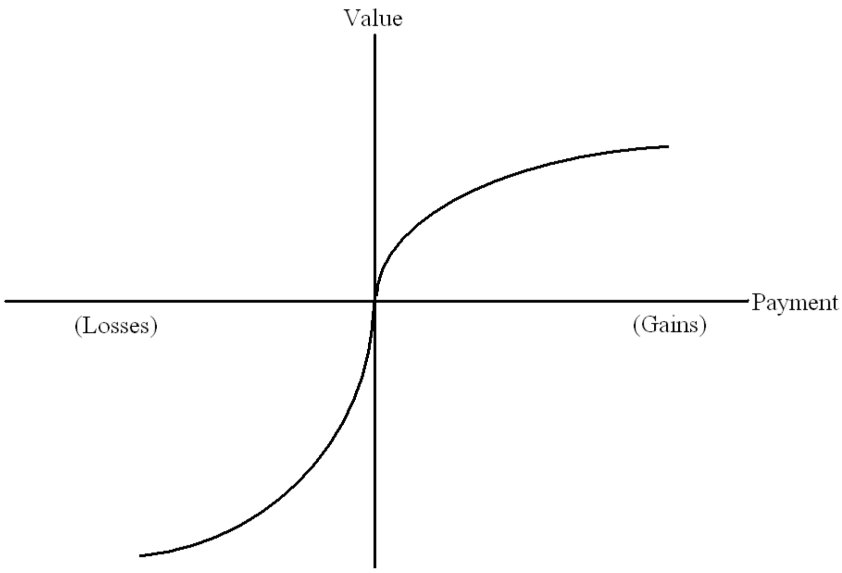

Notice that it has the following properties: (a) increasing everywhere, (b) concave over gains (relative to the reference point), (c) convex over losses, and (d) kinked at the reference point, so that it is steeper for losses than for equal–sized gains. These are the four properties Kahneman and Tversky (1979, p. 279) ascribe to value functions in their description of prospect theory. Properties (c) and (d) are departures from the usual model of expected utility maximization with risk aversion, and (d) is referred to as "loss aversion".

The combination of properties (c) and (d) makes it possible—though not necessary—that an individual who makes decisions according to prospect theory can exhibit both certain–loss avoidance and possible–loss avoidance. In particular, consider an individual decision–making problem that is simply a choice between a safe payment S and a lottery with two possible payments, a low payment L and a high payment H, such that $0 < L < S < H$ (that is, the safe payment and both outcomes to the lottery are gains).[25] Let p be the probability that the high payment is won in the lottery (so that the probability of the low payment is $1 - p$). To save space, we write such a lottery as (H, p, L). In choosing between S and (H, p, L), an individual with the value function shown in Figure 9 would

[25] We consider individual decision problems in this section, rather than strategic games, to keep the example relatively simple. It is straightforward to show that the patterns of behavior we discuss here could also be found in games with suitable payoffs. Also, we use the terms "gains" and "losses" to mean gains and losses relative to the individual's reference point, which as already mentioned, may or may not be her current wealth level.

exhibit risk–averse behavior, as the value function is strictly concave in gains. Its strict convexity over losses, on the other hand, implies that in choosing between the safe payment $S - k$ and the lottery $(H - k, p, L - k)$, where $k > H$—so that all three possible payments became losses—the individual would exhibit risk–seeking behavior. What happens if that constant k is not so large—so that some payments become losses but others remain gains—is less clear cut, and essentially depends on which has the larger effect: the kink in the value function at the reference point, or the convexity of the value function to the left of the reference point. If the kink in the value function dominates, the overall effect will be to make the individual seem *more* risk averse relative to the case of all gains, thus preferring the safe payment relatively more and the risky lottery relatively less. If the convexity of the value function dominates, the overall effect is the opposite: the individual will seem *less* risk averse (again, relative to the case of all gains), as the risky lottery will look relatively better, and the safe payment relatively worse.

In Figure 10, we see a case in which the overall effect leads to the individual's exhibiting possible–loss avoidance. First, consider the left panel of this figure. Here, an individual with a prospect–theory

Figure 10: Possible–loss avoidance with a prospect–theory value function

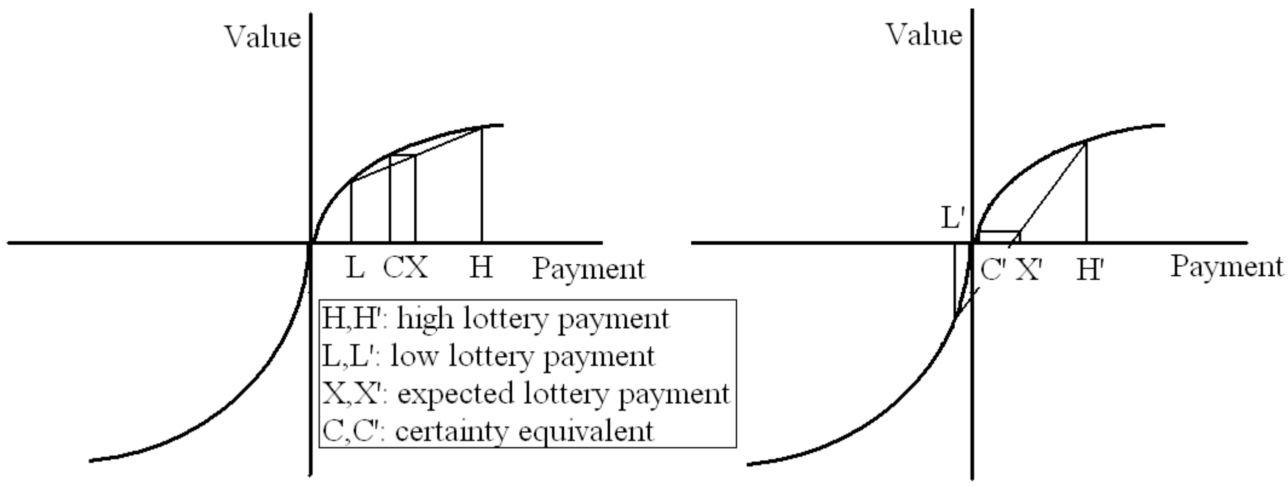

value function is confronting the lottery (H, p, L)—with L and H both positive. The parameter p is such that the expected payment of the lottery is X; X is also shown in the left panel. Vertical line segments connect L and H with their values to that individual, and the diagonal line segment connecting these points shows the values of all possible binary lotteries with these two payments, including the particular lottery (H, p, L). Lastly, this panel shows the certainty equivalent C of this lottery; this is the sure payment that gives the individual the same value as the lottery. (This is shown by the horizontal segment connecting the vertical line segments at C and X.) Any sure payment larger than C will be strictly preferred to the lottery (H, p, L), while the lottery will be strictly preferred to any sure payment smaller than C. Note that if the certainty equivalent is to the left (resp., right) of the expected payment, the individual will be exhibiting risk–averse (resp., risk–seeking) behavior, and the distance between the certainty equivalent and the lottery's expected payment (equal to the length of the line segment between C and X) is a measure of the magnitude of this risk aversion or risk seeking.

Next, consider the right panel of Figure 10. This panel shows the lottery (H', p, L') that results from subtracting $k' \in (L, H)$ from the payments of the original lottery (with p unchanged). We then have $L' < 0 < H'$. Since p is unchanged, the expected payment of this new lottery is $X - k'$, which we denote X'. As in the left panel, the values of the two possible payments are shown, as is the value of this lottery. Finally, the certainty equivalent C' of this new lottery is shown.

In both panels of Figure 10, the certainty equivalent is to the left of the expected payment, so that the individual will exhibit risk–averse behavior in both situations. However, the length of the segment connecting C' and X' in the right panel is visibly greater than that of the segment connecting C and X in the left. This implies that $C' < C - k'$, so that $(C' + k', C)$ is a nonempty interval. Then, for any $S \in (C' + k', C)$, according to the left panel, the individual would choose the lottery (H, p, L) over the sure payment S, but according to the right panel, she would choose the sure payment S' over the lottery (H', p, L') (that is, she would choose $S - k'$ over $(H - k', p, L - k')$). Since $S' > C' > 0$, the sure payment is a gain, so this pattern of choices is consistent with possible–loss avoidance.

In Figure 11, we see a case in which the net effect of the kink and curvature of the value function leads to the individual's exhibiting certain–loss avoidance. The left panel in this figure is identical to

Figure 11: Certain–loss avoidance with a prospect–theory value function

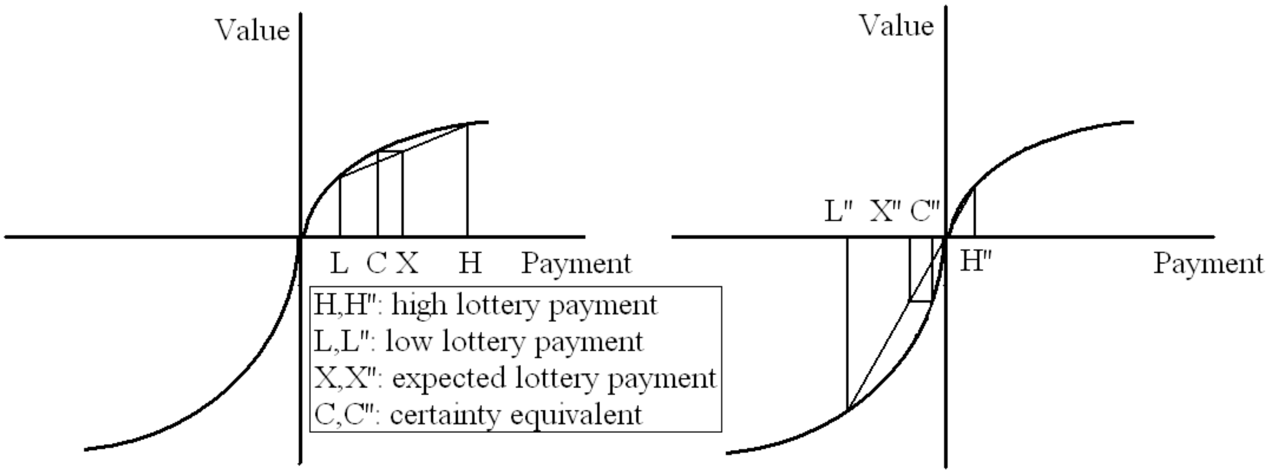

its counterpart in Figure 10, with all payoffs positive. In the right panel, the lottery's payoffs have been reduced by $k'' \in (L, H)$ (with $k'' > k'$) from those in the left panel, so that the new lottery can be written (H'', p, L''), with $L'' < 0 < H''$. As before, both panels show the expected payments X and X'' of the two lotteries (and p is unchanged, so $X'' = X - k''$), as well as their certainty equivalents C and C''. In contrast to the previous case, the certainty equivalent of the lottery in the right panel of Figure 11 is to the right of the expected payment, so that the individual will exhibit risk–seeking behavior when faced with that lottery. This implies that $C'' > C - k''$, so that $(C, C'' + k'')$ is a nonempty interval. Then, for any $S \in (C, C'' + k'')$, according to the left panel, the individual would choose the sure payment S over the lottery (H, p, L), but according to the right panel, she would choose the lottery (H'', p, L'') over the sure payment S'' (that is, she would choose $(H - k'', p, L - k'')$ over $S - k''$). Since $S'' < C'' < 0$ (and hence the sure payment is a loss), this pattern of choices is consistent with certain–loss avoidance.

These two cases show that an individual with a prospect–theory value function can exhibit both certain– and possible–loss avoidance. However, we must emphasize that these examples do not prove that prospect theory implies loss avoidance—merely that it is possible to construct a safe–versus–risky decision in which the safe payment is a gain and the risky decision takes place on a relatively concave section of the value function, due to the kink at the reference point (to observe possible–loss avoidance), and another safe–versus–risky decision in which the safe payment is a loss and the risky decision takes place on a relatively convex section of the value function (to observe certain–loss avoidance).[26] The attitude of the individual to the lotteries in Figures 10 and 11 arises from the shape of the value function, and is only weakly connected to whether the safe payment is positive or negative. Since the value function is continuous, it may well be possible for the same individual to avoid (certain or possible) losses in a particular choice problem, while seeking losses when the choice problem is slightly different.

For an illustration of an individual with a prospect theory value function who *seeks* certain losses, consider Figure 12, in which the individual has a choice problem similar in some ways to that in Figure 11. In this case, however, the effect of the kink in the value function is overshadowed by that of

Figure 12: Certain–loss seeking with a prospect–theory value function

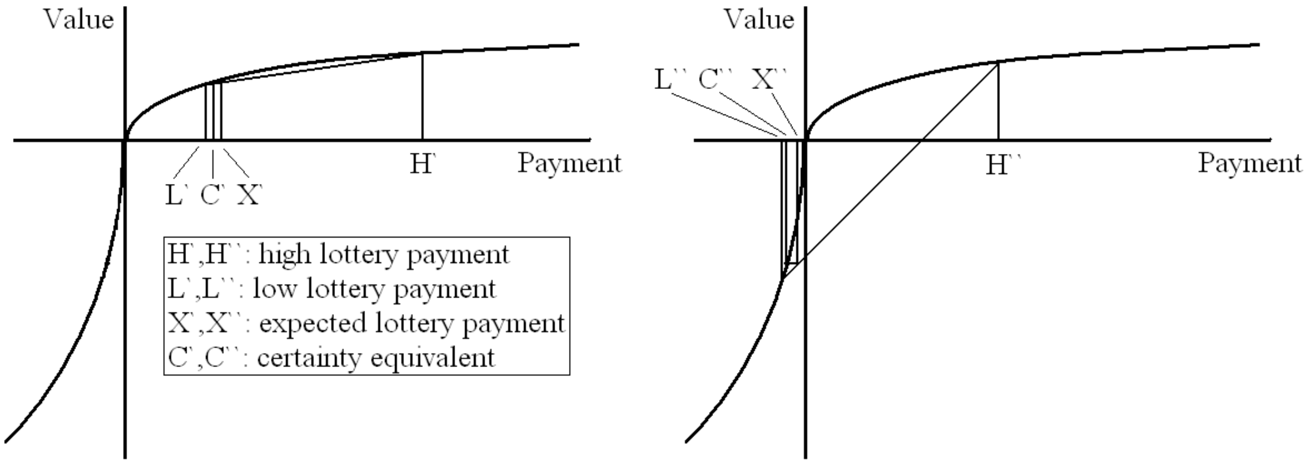

the convexity of the value function. The left panel, as did its counterpart in Figure 11, shows a lottery with only positive payments; we denote this lottery (H', p', L'). In the right panel, the payments have been reduced by a constant k' to obtain the lottery (H'', p'', L''), with $L'' < 0 < H''$. As shown in this panel, the certainty equivalent C'' of the new lottery is not only below the lottery's expected payment X'' (in contrast to the corresponding panel in Figure 11), but is also more than k' below its counterpart C' in the left panel: $C'' < C' - k''$. Therefore, if $S' \in (C'' + k', C')$ (which is a nonempty interval), then based on the left panel, the individual would choose the lottery (H', p', L') over the sure payment S', but based on the right panel, she would choose the sure payment $S'' = S' - k'$ over the lottery (H'', p'', L''). Since $S'' < C' - k' < 0$, S'' is a loss, so this pattern of choices is the opposite of what certain–loss avoidance would predict (we could call this "certain–loss seeking").[27]

Not only is it not true that certain– and possible–loss avoidance are implied by prospect theory,

[26] An implication of this last point is that Kahneman and Tversky's *loss aversion* alone is not enough to imply certain–loss avoidance. An individual whose value function had the prospect–theory kink at the reference point, but not the convexity to the left of the reference point, would exhibit increased risk aversion whenever a lottery had both gains and losses possible, compared to when only gains or only losses were possible. Such an individual would exhibit possible–loss avoidance, but certain–loss *seeking*.

[27] In case these graphical examples are not convincing, consider the value function given by $v(x) = \sqrt{x}$ for $x \geq 0$, and $v(x) = -3\sqrt{-x}$ for $x < 0$; it is easy to show that v satisfies the four properties of a value function mentioned earlier (for payments x). Suppose an individual with this value function, and a reference point equal to current wealth, is faced with the lottery $(H, p, L) = (4, \frac{1}{2}, 1)$. In a choice between this lottery and a sure payment of $S = 2.1$, the individual would prefer the lottery (with a value of 1.5 versus approximately 1.449). Subtracting 1.6 from all payments yields a choice between the lottery $(2.4, \frac{1}{2}, -0.6)$ and a sure payment of 0.5; here, the individual would prefer the sure payment (with a value of approximately 0.707 versus approximately –0.387)—that is, she would prefer the sure gain over the possible loss, giving us an example of possible–loss avoidance. For an example of certain–loss avoidance, consider the original

but the converse is also not true; it is possible that an individual exhibits either or both forms of loss avoidance, but her behavior is not governed by prospect theory. First of all, prospect theory entails more than just a value function like that shown in Figure 9—for example, it includes a probability "weighting function" (Kahneman and Tversky (1979, p. 282)), which models defects in the way individuals assess probabilities—it's possible that someone with a prospect–theory value function does not have these other features. Second, other theories besides prospect theory may give rise to a similar type of value function.[28] Because certain– and possible–loss avoidance are consistent with, but neither entailing nor implied by, prospect theory, we stress that this paper should not be viewed as a test of prospect theory, or indeed any particular model of decision making. Rather, our aim is simply to examine whether individuals in certain situations exhibit certain– and/or possible–loss avoidance.

---

lottery versus a sure payment of 2.4; the individual would prefer the sure payment (with a value of approximately 1.549 versus 1.5). Subtracting 3.4 from all payments yields a choice between the lottery $(0.6, \frac{1}{2}, -2.4)$ and a sure payment of –1, in which case the individual would prefer the lottery (with a value of approximately –1.936 versus –3)—that is, she would prefer the possible gain over the certain loss. For an example of certain–loss *seeking*, consider the original lottery versus a sure payment of 2.1, in which case we've already shown the individual will prefer the lottery. Subtracting 2.2 from all payments yields a choice between the lottery $(1.8, \frac{1}{2}, -1.2)$ and a sure payment of –0.1. Here, the individual will prefer the sure payment (with a value of approximately –0.949 versus approximately –0.972)—that is, she would prefer the sure loss over the possible gain.

[28]Even expected–utility maximizers may exhibit behavior sensitive to payoff levels, due to curvature of the utility function, if the gains and losses are money amounts (as is typical in economics experiments). This effect would have to be quite small under the usual assumption that expected–utility maximization takes place over final wealth levels (Rabin and Thaler (2001)), but if expected–utility maximization is considered to take place over gains and losses from a reference point such as current wealth, a larger effect is possible without the implication of a bizarre utility function (Cox and Sadiraj (2007)).