

Payoff levels, loss avoidance, and equilibrium selection in games with multiple equilibria: 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, 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

May 25, 2010

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 “payoffs” are actually money amounts. *Loss avoidance*, in particular, predicts that payoff-level changes may matter when they result in gains becoming losses 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 strong evidence of behavior consistent with certain-loss avoidance in the experiment. We also find evidence of possible-loss avoidance, though less than we find for certain-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: D81, C72, C73.

Keywords: experiment, game theory, behavioral economics, stag hunt, coordination game, 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). Feltovich acknowledges support from the University of Aberdeen Business School. We thank Alexander Elbittar, Yasuyo Hamaguchi, Andreas Ortmann, John Van Huyck, Nathaniel Wilcox, Erte Xiao, and participants at several conferences and seminars for helpful suggestions and comments.

1 Introduction

Game-theoretic solution concepts carry the implication that a “change in payoff levels”—a modification to a game by adding the same (positive or negative) constant to all payoffs—should have no effect on behavior. Such changes do not affect expected-payoff differences between strategies under any given beliefs about opponents’ actions, so that best-response correspondences and therefore Nash equilibria (pure or mixed) are unaffected. When the game has multiple Nash equilibria, game theory does not rule out the possibility that payoff-level changes will lead to a change in which equilibrium 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 least silent as to whether changes in payoff levels can affect players’ choices, and at most explicitly rules out such an effect.

This implication relies, however, on the game’s payoffs reflecting the true preferences of the players—that is, the equivalence between payoffs and expected utilities. When payoffs are monetary gains or losses (as 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 evidence dating back at least to Kahneman and Tversky (1979) suggesting that payoff-level changes can affect behavior. Although Kahneman and Tversky did not look specifically at payoff levels, their “prospect theory” allows for behavior to be sensitive to them.¹ More recently, Cachon and Camerer (1996) found that decisions in a coordination-game experiment could be sensitive to changes in payoff levels when these changes affected the *signs* of payoffs: positive payoffs became negative, or vice versa. They speculated that subjects showed “loss avoidance”, which they defined to be a tendency to avoid choices that with certainty yield negative payoffs in favor of alternatives that could yield positive payoffs.² Based on this finding, they conjectured that loss avoidance could be used as a criterion for equilibrium selection.³

Cachon and Camerer’s notion of loss avoidance is one of several ways that individuals could treat losses and gains differently. In this paper, we distinguish between their notion, which we call “certain-loss avoidance”, and another type, which we call “possible-loss avoidance” and define as a tendency to avoid strategies that give a *possible* negative payoff, in favor of one that gives a certain positive payoff. We will sometimes use the term “loss avoidance” to encompass both certain- and possible-loss avoidance.

The goal of this paper is to test for loss avoidance in strategic behavior. To do so, we design and conduct an experiment using a stage game well-suited for such a test: Stag Hunt (Rousseau, 1973), a symmetric two-player game with one risky action and one safe action, and with multiple Nash equilibria. Table 1 shows three versions of Stag Hunt, with the risky and safe actions labeled R and S.

These games differ only in payoff level, so they are equivalent from a game-theoretic standpoint, as long as the payoffs represent the preferences of the players. When these payoffs represent monetary gains and losses, then despite this seeming equivalence, our notions of loss avoidance make qualitative predictions regarding how individuals’ choices might change across the games. In the high-payoff game (SHH), all payoffs are gains, so

¹It should be noted that Kahneman and Tversky used hypothetical, not monetary, payments, so their results should be viewed as suggestive rather than conclusive. Some research has compared behavior under real incentives versus hypothetical incentives; see, for example, Holt and Laury (2002, 2005).

²More precisely, they speculated that subjects believed their opponents were exhibiting loss avoidance; see Section 2.2.

³In an earlier version of this paper (Feltovich, Iwasaki, and Oda, 2008), we discuss the relationship between loss avoidance and Kahneman and Tversky’s (1979) prospect theory. Here, we note merely that loss avoidance is consistent with, but neither entailing nor implied by, prospect theory; that is, the two phenomena sometimes, but not always, imply the same choices. In particular, despite an unfortunate similarity in their names, loss avoidance is distinct from the component of prospect theory known as *loss aversion*.

Figure 1: The Stag Hunt games

		Player 2				Player 2				Player 2	
		R	S			R	S			R	S
Player	R	7,7	1,5	Player	R	5,5	-1,3	Player	R	1,1	-5,-1
1	S	5,1	5,5	1	S	3,-1	3,3	1	S	-1,-5	-1,-1
High payoffs (SHH)				Medium payoffs (SHM)				Low payoffs (SHL)			

neither type of loss avoidance will have an effect. In the medium-payoff game (SHM), on the other hand, choosing S leads to a certain gain, but choosing R could lead to a loss (if the opposing player chooses S). If an individual exhibits possible-loss avoidance in her decision making, R should therefore be less attractive in SHM than it would have been in SHH. Similarly, in SHL, choosing S leads to a certain loss, but choosing R could lead to a gain. An individual exhibiting certain-loss avoidance will find S less attractive in SHL than it would have been in SHH. So, possible-loss avoidance implies that R should be chosen less often in SHM than in SHH, while certain-loss avoidance implies that S should be chosen less often (and thus R should be chosen *more* often) in SHL than in SHH.

In our experiment, subjects play all three of these games. We manipulate various aspects of the strategic environment across sessions, including the number of times a game is played (once or forty times), how players are matched (repeated play against the same opponent or random matching over all possible opponents), and how much information players are given about payoffs (complete or limited information). Our experiment is designed as a fairly conservative test for loss avoidance. Nevertheless, we find that in all treatments, there are significant differences in behavior across the games—that is, behavior is indeed sensitive to payoff levels. The effect of payoff levels on choices is small in early rounds, but often 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. A follow-up experiment using different payoff matrices and minor changes to the experimental procedures gives broadly similar results, suggesting that loss avoidance is a fairly robust phenomenon.

2 The experiment

The experiment uses the three versions of Stag Hunt shown in Figure 1. As mentioned already, these games are identical from a game-theoretic standpoint, as long as payoffs reflect players' preferences. Each has three Nash equilibria: (R,R), (S,S), and a mixed-strategy equilibrium in which both players choose R with probability 2/3. Additionally, 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 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. Also, since all three equilibria (pure- or mixed-strategy) are symmetric, the coordination problems implicit in symmetric games with asymmetric equilibria are not an issue. Relatedly, Stag Hunt games

have 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, any change to the strategic environment (such as a change to payoffs) that is the same for 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.⁴

2.1 Previous research on Stag Hunt and related games

Stag Hunt is a simple member of a class of coordination games called “order–statistic games” (Van Huyck, Battalio, and Beil, 1990), which also includes median– and minimum–effort games. These games have the advantageous features of Stag Hunt for studying equilibrium selection: multiple strict symmetric Nash equilibria and strategic complementarities. 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 showed that behavior in these games is difficult to predict—in particular, it is quite possible that subjects will coordinate on an inefficient Nash equilibrium, or fail to coordinate on any pure–strategy equilibrium at all.⁵ As a result, much later research has focused on factors, besides the payoffs of the game, that make subjects more likely to choose one strategy over others. Factors that have been found to have an effect include experience in the game (Stahl and Van Huyck, 2002), additional feedback about others’ choices (Berninghaus and Ehrhart, 2001; Devetag, 2003), and opportunities for sending cheap–talk messages (Clark, Kay, and Sefton, 2001; Manzini, Sadrieh, and Vriend, 2002; Duffy and Feltovich, 2002; Charness and Grosskopf, 2004; Duffy and Feltovich, 2006; Blume and Ortmann, 2007). Clark and Sefton (2001) found that the way subjects were matched affected behavior: the payoff–dominant outcome was more likely under fixed–pairs matching than under rotation (one play against each opponent). (Interestingly, this difference was present 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 matching, though they did find some differences at the session level. They also looked at one–shot versions of these games, and found a lower frequency of risky–action choices in a given one–shot game compared to the first round of the same game under changing opponents; however, these differences were not significant.

2.2 Previous research on payoff–level effects

There has been some experimental research into the effects of changing payoff levels, in both individual decision problems and games. Kahneman and Tversky (1979) used results from a series of one–shot decision problems to

⁴Indeed, the positive feedback doesn’t stop here. 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 limiting case where the effect of a payoff–level change is common knowledge amongst the players, there would be an infinite number of these chains of reasoning.

⁵For example, Van Huyck, Battalio, and Beil (1990) found that behavior in minimum–effort games was sensitive to the number of players and to the matching mechanism: with multiple players, play converged to the least efficient pure–strategy equilibrium; with two players and fixed–pairs matching, play tended to converge to the most efficient equilibrium; and with two players and random matching of opponents, play didn’t appear to converge at all. Van Huyck, Battalio, and Beil (1991) looked at a multi–player median–effort game, and found that behavior tended to converge to whatever the first–round median was. For more comprehensive reviews of Stag Hunt and other order–statistic games, see Ochs (1995), Camerer (2003, Chapter 7), and Devetag and Ortmann (2007).

guide the construction of their “prospect theory”. Prospect theory does allow for changing payoff levels to affect choices, but the exact nature of this effect may be difficult to predict *a priori*, as it can be sensitive to framing effects.⁶ There has also been some research into payoff-level effects in repeated decision problems, mainly by Ido Erev and colleagues. A typical result from these experiments is that the speed of learning an optimal choice under limited information is faster when both gains and losses are possible than when only gains or only losses are possible.⁷

Erev, Bereby-Meyer, and Roth (1999) examined learning in a repeated 2x2 constant-sum game with a unique mixed-strategy Nash equilibrium. Payoffs in the game represented probabilities of “winning” rather than “losing”. In one treatment, the prizes for “winning” and “losing” were 1 and 0 respectively, while in the other treatment, they were +0.5 and -0.5. They found that when losses were possible, subject choices were more consistent with fictitious-play learning (that is, more sensitive to differences in historical payoffs). Earlier, Rapoport and Boebel (1992) had looked at two versions of a repeated 5x5 constant-sum game with a unique mixed-strategy equilibrium. 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 the former, both gains and losses were possible in both treatments, while in the latter, one treatment had both gains and losses while the other had no losses possible.

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 changing payoff levels was accomplished by varying firms’ sunk costs. Results have depended on the market institution used. When the institution is one with a strong tendency toward equilibrium, sunk costs have typically had no effect (see Kachelmeier (1996), who used a double auction, and Waller, Shapiro, and Sevcik (1999), who used a multi-firm posted-price market). On the other hand, when weaker tendencies toward equilibrium prevail, sunk costs can have an effect (see Offerman and Potters (2003) and Buchheit and Feltovich (2010), both of whom used price-setting duopolies with capacity constraints).

The studies most closely related 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 2) and a related minimum-effort game. They implemented the payoff-level changes in the form of a mandatory fee of either 185 or 225 imposed on all players, and publicly announced this fee (in an attempt to make it common knowledge) in their “Must Play” treatments. They found that subjects tended to avoid strategies that yielded a certain negative payoff. However, they found no such effect in their “Private Cost” treatment, where the mandatory fee was not publicly announced, so that subjects knew only their own fees. Cachon and Camerer concluded that the subjects in the experiment themselves did not exhibit certain-loss avoidance; rather, the subjects believed that their opponents were exhibiting certain-loss avoidance, so they themselves acted in a way consistent with certain-loss avoidance due to the strategic complementarities in the game.

Rydval and Ortmann (2005) examined behavior in several versions of Stag Hunt, including two pairs within which only payoff levels were changed. These four games are shown in Figure 3; 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

⁶For 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.

⁷See, for example, Barkan, Zohar, and Erev (1998), Bereby-Meyer and Erev (1998), Erev, Bereby-Meyer, and Roth (1999). There was also earlier work that didn’t look specifically at payoff-level effects, but gave results suggesting that changes in payoff levels might affect behavior; see Siegel and Goldstein (1959) and Siegel, Siegel, and Andrews (1964).

Figure 2: Cachon and Camerer’s (1996) 7-player median–effort game

		Median effort level						
		1	2	3	4	5	6	7
Player’s effort level	1	140	150	140	110	60	–10	–100
	2	130	160	170	160	130	80	10
	3	100	150	180	190	180	150	100
	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 185 or 225 in “Must Play” treatments)

60 from all payoffs, and Game 5 is obtained from Game 4 by subtracting 180). Note also that within each pair,

Figure 3: Two high–payoff and two low–payoff Stag Hunt games from Rydval and Ortmann (2005)

		Player 2		Player 2		Player 2		Player 2	
		C	D	C	D	C	D	C	D
Player 1	A	80,80	10,50	A	20,20	–50,–10	A	240,240	30,150
	B	50,10	50,50	B	–10,–50	–10,–10	B	150,30	150,150
		Game 2		Game 3		Game 4		Game 5	

one game has only positive payoffs—like our SHH game—and the other had all negative payoffs except when both players choose the risky action—like our SHL game. Subjects played each game once, with opponents 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 in risky–strategy choice frequencies; thus the overall evidence for certain–loss avoidance in their experiment was weak.⁸

2.3 Experimental design and procedures

We used a within–subjects design, with subjects playing all three Stag Hunt games. One concern we had was that subjects who recognized the close similarity of these games might view the experiment as a “consistency check” and choose the same actions in all three. To limit this possibility, we tried to make this similarity less transparent without using deception. This was done by adding three games: a Prisoners’ Dilemma (PD), a Battle of the Sexes (BoS), and a coordination game (CG) (see Figure 4).

Subjects played the 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 re-

⁸Rydval 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 counteracting those of loss avoidance.

Figure 4: The other games used in the experiment

		Player 2				Player 2				Player 2	
		R	S			R	S			R	S
Player	R	2,2	0,0	Player	R	0,0	3,5	Player	R	7,7	1,8
1	S	0,0	1,1	1	S	5,3	0,0	1	S	8,1	4,4
Coordination Game (CG)				Battle of the Sexes (BoS)				Prisoners' Dilemma (PD)			

searchers have shown that the order of the games can matter; see, for example, Falkinger et al. (2000) or Duffy and Feltovich (2006).) In order to ameliorate this problem, we varied, across experimental sessions, the order in which the three Stag Hunt games are played. We used three of the six possible 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 publicly announced (in an attempt to induce common knowledge). 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. Our C (complete-information) treatment was similar, except subjects played each game 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 not only the consequences of subjects' deductive reasoning, but also 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 these eventually. (Specifically, a subject would be told her opponent's choice and her own payoff in the just-completed round.) Our R (random-matching, limited-information) treatment was otherwise similar to the C treatment: subjects played each game 40 times against randomly-chosen opponents. Our F (fixed-pairs, limited-information) treatment was similar, except that subjects played against the same opponent for all 40 rounds of a game.⁹ These two treatments allow us to 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 learning of loss avoidance depends on how subjects are matched.¹⁰

In the R and F treatments, we were concerned that subjects' choices might be sensitive to the order in which the

⁹Van Huyck, Battalio, and Beil (1990) and Schmidt et al. (2003) also varied the matching mechanism between (in Schmidt et al.'s nomenclature) "fixed match" and "random match." However, subjects in all treatments of these experiments had complete information about payoffs, so this "random match" treatment is comparable to our C treatment. We have no treatment analogous to the "fixed match" treatment here (however, see our follow-up experiment in Section 4).

¹⁰We 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 matching—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).

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 design. When necessary to avoid confusion with our manipulations of game and action ordering, we will refer to our O, C, R, and F treatments 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 all six games under a single information treatment, with a single game ordering and single action ordering. Sessions took place at the Kyoto Experimental Economics Laboratory (KEEL) at Kyoto Sangyo University. Subjects were primarily undergraduate students at this university, recruited via a database of participants in other experiments and via advertisements posted on campus. About a third were economics majors; the rest were fairly 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; the instructions were also read aloud.¹¹ Partitions prevented subjects from seeing other subjects’ computer screens, and subjects were asked not to communicate with each other during the session. The experiment 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.

Subjects were told in the instructions 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. Each round began by prompting subjects to choose one of two possible actions. To reduce demand effects, the 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. After seeing this information, subjects were prompted to click a button to go on to the next round.

At the end of a 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.

¹¹English translations of the instructions, copies of the instructions for the follow–up experiment discussed in Section 4, and the raw data from both experiments are available from the corresponding author upon request.

2.4 Hypotheses

Our hypotheses follow directly from the implications of certain- and possible-loss avoidance, as discussed in the Introduction.

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.*

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.*

Before moving on to the results, we wish to point out that our experiment is actually a conservative test for loss avoidance for several reasons. First, even though the SHM and SHL payoff matrices show the potential for losses, it was impossible for subjects to actually incur negative payments in the experiment overall, due to institutional rules for human-subjects experiments. As a result, any perception of losses in a game by the subjects is due not just to negative entries in the payoff matrix, but also to framing—that is, our success in getting subjects to internalize the showup fee and gains from other games. To the extent that subjects do not do so—and thus treat “losses” as simply reductions in an overall gain—even subjects who do exhibit loss avoidance will not show it in the experiment, so that the actual prevalence of loss avoidance in subjects’ preferences will be higher than what we are able to observe. Second, an individual who exhibits loss avoidance may not show it for these particular Stag-Hunt games. For example, an individual who is extremely risk seeking might well choose the risky action over the safe action in all three games, though the intensity of preference (not measured in this experiment) for the risky action would be higher in SHL than in the other two games. Similarly, an extremely risk averse individual might choose the safe action in all three games, even if loss avoidance means that the intensity of preference for the safe action is highest in SHM. Without a way of showing intensity of preferences in our experiment, these two individuals will appear in our results not to exhibit loss avoidance, even if they actually do; as a result, our observed evidence of loss avoidance will understate the true level. Finally, the effect of loss avoidance will only be seen when there is genuine strategic uncertainty in a subject’s decision. If instead a subject believes that her opponent is nearly certain of choosing a particular action, then the best response is clearly a choice of the same action, irrespective of loss avoidance. For all of these reasons, the observed results of our experiment will tend to understate the prevalence of loss avoidance.

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, with a total of 378 subjects; the number of subjects in a session varied from 6 to 28.¹²

3.1 First-round behavior under complete information

We first look at initial subject behavior under complete information—that is, the O treatment (in which each game was played only once) and the first round of each game in the C treatment. Table 2 shows the frequencies of risky-action choices in each Stag Hunt game. 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

¹²We ran the seventh F session in order to replicate the design parameters of an earlier F session that had only 6 subjects.

games: negligible in the low-payoff Stag Hunt, nonnegligible but insignificant (chi-square test, $p > .10$) in the medium-payoff Stag Hunt, significant ($p < .05$) in the high-payoff Stag Hunt.¹³

Next, consistent with certain-loss avoidance (Hypothesis 1), 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 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. 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, $p < 0.001$).

Table 2: 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.

On the other hand, we fail to find evidence in favor of possible-loss avoidance here (Hypothesis 2). 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

Table 3 shows some aspects of play in the three repeated-game treatments: overall frequencies of risky-action choices, and these frequencies broken down by ordering of games and of actions. There is variation—sometimes substantial—across treatments, but overall, differences in play across games are typically consistent with both certain-loss avoidance and possible-loss avoidance.

In the C treatment, we see here—as was true in the first round alone—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 order relationships are consistent with certain- and possible-loss avoidance respectively. The results from our limited-information treatments are also consistent with certain-loss avoidance. The 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

¹³See Siegel and Castellan (1988) for descriptions of the nonparametric statistical tests used in this paper.

Table 3: 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 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 without superscripts or with superscripts having a letter in common are not significantly different. See text for details.

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 3.3.) Power is an issue here, especially in the C and R treatments, where 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 in a game 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 robust rank–order test is not significant in the direction predicted by possible–loss avoidance ($\hat{U} = -0.81$, $p > 0.10$). On the other hand, the difference between SHL and SHH observed in Table 3 is significant ($\hat{U} = -3.70$, $p < 0.001$), consistent with certain–loss avoidance. 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 (Wilcoxon signed–ranks test, $T^+ = 11$, $N = 7$, $p > 0.10$), 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


Figures 5 and 6 show round–by–round frequencies of risky–action choices in the C treatment and R and F treatments, respectively. Figure 5 shows once again that in the C treatment, there are differences across the three games, and while small, they are consistent with both types of loss avoidance. We also see that average play does not change

Figure 5: Frequency of risky action (complete-information treatment, all sessions)



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.

Figure 6: Frequency of risky action (limited-information treatments, all sessions)



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 both treatments, changes over time are consistent with certain-loss avoidance, though the nature of the changes varies: in the R treatment, risky-action play in SHL stays roughly constant over time, while falling sharply in the other two games; in the F treatment, risky-action play rises over time in SHL and remains about the same in the other two games.¹⁴ Evidence of possible-loss avoidance is

¹⁴We 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 our Footnote 9 above). 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

visible in the R treatment, where risky–action choices drop more sharply in SHM than in SHH. In the F treatment, by contrast, there are actually slightly more risky–action choices in SHM than in SHH after the first ten or so rounds.

In order to gain some power in our statistical tests, we next look at some probit regression results. (Logits—not reported here—gave similar results.) The dependent variable in each regression is an indicator for the current–round action (1=risky, 0=safe). 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 their products with the round number and its square, to allow for the possibility that loss avoidance develops (or disappears) over time. As controls, we include indicators for two of the three game orderings (M-L-H and L-M-H) and one of the two action orderings (R-S), and variables for the round number itself and its square. The regressions were performed using Stata (version 10) using individual–subject random effects; we estimate coefficients separately for each of the three multiround treatments (C, R, and F).

Table 4: 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*** (0.240)	0.570*** (0.114)	−0.130 (0.153)
SHM	−0.088 (0.159)	0.280*** (0.097)	−0.161* (0.086)
SHM · Round	0.000 (0.018)	−0.087*** (0.013)	0.031*** (0.010)
SHM · Round ²	−0.0004 (0.0004)	0.0017*** (0.0003)	−0.0006*** (0.0002)
p –value (joint significance of three SHM variables)	< 0.001***	< 0.001***	< 0.001***
SHL	0.348 (0.170)	−0.152* (0.091)	0.064 (0.086)
SHL · Round	0.002 (0.019)	0.098*** (0.010)	0.113*** (0.010)
SHL · Round ²	−0.0001 (0.0004)	−0.0012*** (0.0002)	−0.0022*** (0.0002)
p –value (joint significance of three SHL variables)	< 0.001***	< 0.001***	< 0.001***
Round	0.000 (0.013)	−0.094*** (0.008)	−0.024*** (0.007)
Round ²	−0.0002 (0.0003)	0.0012*** (0.0002)	−0.0005*** (0.0002)
M-L-H game ordering	−1.977*** (0.303)	−0.911*** (0.132)	−0.583*** (0.164)
L-H-M game ordering	−1.430*** (0.290)	−0.354*** (0.111)	0.113 (0.185)
R-S action ordering	—	−0.071 (0.098)	0.382** (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.

Table 4 shows coefficients and standard errors for each variable, as well as log likelihoods and pseudo- R^2 values for each regression.¹⁵ The results show that subject behavior is nonstationary; the coefficients for the six 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.

¹⁵The 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.

variables containing the round number or its square are jointly significant in all three treatments (likelihood–ratio test, $p < 0.001$). Behavior is also sensitive to the ordering of the Stag Hunt games: in each of the three treatments, the variables M-L-H and L-H-M are jointly significant ($p < 0.001$), and taken individually, they are nearly always significantly different from zero (as shown in the table), and always significantly different from each other ($p < 0.10$ for the C treatment, $p < 0.001$ for R and F). 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 not in the R treatment.

We now move on to the results connected with loss avoidance. We continue to find solid support for certain–loss avoidance, but we now find support for possible–loss avoidance as well. First, we note that an implication of loss avoidance—nonzero coefficients for the three SHL variables (SHL, SHL · Round, and SHL · Round²) and for the three SHM variables—is seen in the table. Even when these variables are not significant when considered individually, they are always jointly significant. Next, in order to compute the actual effects of loss avoidance, we estimate the incremental effect of switching the game from SHH to SHM (for possible–loss avoidance) or to SHL (for certain–loss avoidance). In round t , the total effect of the SHM variables 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 of the SHM game rather than the SHH game in round t has the form

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

where \bar{X} is the row vector of the other right–hand–side variables’ values (which will be either the unconditional sample mean or the mean conditional on an appropriate subsample), 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 .


Graphs of all six of these expressions (versions of Equation 1 for SHL and SHM variables and for C, R, and F treatments) are displayed in Figure 7, with point estimates for each round shown as circles and corresponding 90% confidence intervals shown as line segments.¹⁶ Again, we see strong evidence for certain–loss avoidance, but also for possible–loss avoidance. The point estimate for the incremental effect of SHL is nearly always positive, and the corresponding 90% confidence interval lies entirely above zero (so that a one–sided hypothesis test would yield significance at the 5% level) for rounds 3–38 of the C treatment, rounds 3–40 of the R treatment, and all rounds of the F treatment. In the F treatment, the point estimate of the incremental effect of SHM usually has the opposite sign from the prediction of possible–loss avoidance, but in the C and R treatments, the point estimate has the predicted negative sign, and is significantly different from zero at the 5% level in most rounds (rounds 11–40 in the C treatment and 6–40 in the R treatment).

Thus, these time paths show a consistent pattern. Each starts close to zero in round 1 and, from there, moves away from zero. Five of the six time paths (all but SHM in the F treatment) move in the direction predicted by loss avoidance, though in one of these (SHL in the C treatment), the movement is very slight. In the C treatment, both time paths continue to move away from zero, while in the other two treatments, they eventually turn back toward zero, though never crossing the horizontal axis.¹⁷ In other words, the effects of loss avoidance are greatest after

¹⁶Note 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.

¹⁷We also estimated probit models like ones discussed here, but with cubic terms (SHM · Round³, SHL · Round³, and Round³) also included.

Figure 7: Estimates of the effect of game on risky–strategy choice (vs. SHH) based on Table 4 results
(Circles represent point estimates; line segments represent 90% confidence intervals)



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).

3.4 Summary and discussion

In the sections above, we have seen 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. Possible–loss avoidance is more likely to be seen in treatments where strategic uncertainty is relatively high: under random matching versus under fixed pairs, and under limited information versus under complete information.¹⁸

Several points are worth making at this stage. First, we attempt to put some of our results into the context of the literature. Consistent with many others’ results, a comparison of our C and R treatments (see Figures 5 and 6) suggests that giving subjects more information leads to more play of the risky action, and hence greater likelihood of the payoff–dominant outcome. Play of the risky action was also more likely under fixed pairs (our F treatment) than under changing opponents (our R treatment), as was observed by Clark and Sefton (2001)—though not by Schmidt

These gave similar results to those shown in Figure 7, suggesting that the shapes of the trajectories in that figure were not simply an artifact of a restriction to quadratic terms.

¹⁸The connection between loss avoidance and strategic uncertainty is fairly intuitive. As noted earlier, the only time loss avoidance should be expected to have an effect is when the level of strategic uncertainty is fairly high, since whenever a player is nearly certain that the opponent will play a particular action, the best response is obvious: choice of that same action. (In such cases, a “possible loss” according to the payoff matrix is not the same as a possible loss from the player’s point of view.) We claim that other things equal, strategic uncertainty should be higher under random matching than under fixed pairs, since in the latter, subjects can build up experience quickly against the same opponent instead of slowly against a group of potential opponents. Also, strategic uncertainty should be higher under complete payoff information—where it is easier to develop a model about how opponents make decisions, and more reasonable to assume that the opponent has a similar model—than under limited payoff information. As a result, strategic uncertainty ought to be higher in the R treatment (both limited information and random matching) than in the C and F treatments (either complete information or fixed pairs), consistent with our result that possible–loss avoidance is most pronounced in the R treatment, and less so in the other two treatments. For certain–loss avoidance, by contrast, the effect is so strong that it is apparent in all treatments, irrespective of differences in the level of strategic uncertainty. We acknowledge that this is an ex post justification of our results, though the fact that the results in our new experiment show a similar pattern (see Section 4) is at least suggestive.

et al. (2003) who found no systematic differences between the two. We note, however, that both of these latter papers implemented these matching mechanisms in ways different enough from ours (their subjects had complete payoff information, while ours had limited information; their games lasted for fewer rounds; etc.) that our results should not be viewed as a replication of the former or a contradiction of the latter. Like Schmidt et al., we found a higher frequency of risky-action play in the first round of our C treatment (which corresponds most closely to Schmidt et al.'s random-match treatment) compared with our O treatment in Table 2, but also like Schmidt et al. the difference is usually insignificant—though we do find significance in one of our games.

The results of our R and F treatments have a similar flavor to those observed by Erev and colleagues, who found in both individual decision tasks and games 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. Finally, Cachon and Camerer (1996) found that loss avoidance is more observable in subjects' beliefs about their opponents than in their own behavior. Our experiment was not designed to test for this, but consistent with their conclusion, our result that loss avoidance in the complete-information treatment gains strength after subjects have gained experience suggests that at least some of the effect 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 repetition of the games. In many cases, differences in play across games start out small, but grow over time (at least up to a point) in the direction predicted by certain- and possible-loss avoidance. This suggests that loss avoidance due purely to introspection may be relatively minor, but that it can be learned. This "learning to avoid losses" was found in all three of our repeated-game treatments. In the C treatment, one might argue 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 simply learn to play as loss avoidance predicts, due to the strategic complementarities in these games. 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 differences across games develop nonetheless, so it would be harder to argue against the claim that at least some subjects are learning to avoid losses themselves.

Finally, 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 rather than the inefficient secure (and in our games, risk-dominant) equilibrium. 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 negative payoffs will 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 least efficient pure–strategy Nash equilibrium.

4 A new experiment

In order to check the robustness of our results, we conducted a new experiment in which we maintained the important aspects of the original experiment, but altered several minor design features. We continue using Stag Hunt games, but with new payoffs; these games are shown in Figure 8. While clearly different from the original games (see Figure 1), they resemble the original games in that they are the same up to the payoff level, and the signs of payoffs in our new high–payoff (NSHH), new medium–payoff (NSHM), and new low–payoff (NSHL) games are the same as in the SHH, SHM, and SHL games respectively. Our hypotheses for these games are thus identical to our previous hypotheses (see Section 2.3).

Figure 8: The Stag Hunt games used in the new experiment

		Player 2		Player 2		Player 2			
		R	S	R	S	R	S		
Player 1	R	360,360	40,260	R	220,220	−100,120	R	80,80	−240,−20
	S	260,40	260,260	S	120,−100	120,120	S	−20,−240	−20,−20
High payoffs (NSHH)				Medium payoffs (NSHM)				Low payoffs (NSHL)	

The new experiment comprises four information treatments. As in the original experiment, we have a C (complete information, random matching) treatment, an R (limited information, random matching) treatment, and an F (limited information, fixed–pairs matching) treatment, but we now include a CF (complete information, fixed–pairs matching) treatment to complete the 2x2 factorial design. As before, subjects play other games in between the Stag Hunt games, in an attempt to make the similarity of the games less obvious; these games are shown in Figure 9. The

Figure 9: The other games used in the new experiment

		Player 2		Player 2	
		R	S	R	S
Player 1	R	160,160	80,200	R	120,120 40,160
	S	200,80	20,20	S	160,40 -20,-20
Chicken-high payoff (CHH)				Chicken-low payoff (CHL)	

games were played in the order NSH-CH-NSH-CH-NSH, and as in the original experiment, we varied the ordering

of the Stag Hunt games (H-M-L, M-L-H, and L-H-M). We also varied the ordering of the Chicken games (H-L and L-H), but we did not vary the action ordering (only the R-S ordering was used). Each game was played only 20 times, instead of 40.

The new experiment took place at the Scottish Experimental Economics Laboratory (SEEL) at University of Aberdeen; subjects were primarily undergraduate students at this university, recruited via advertisements posted on campus and on a university web site. Much of the experimental procedure was the same as in the original experiment, but some features were changed. The most important changes were: the new experiment was conducted in English and programmed in the English-language version of z-Tree (Fischbacher, 2007), subjects' computer screens included a history of previous results, and the payment scheme was somewhat different, with a showup fee of £8 (at the time of the experiment, equivalent to roughly \$15), added to the results of five randomly-chosen rounds (one from each game) at an exchange rate of £1 per 100 points. Also, limited information was implemented in a slightly different way: a subject's end-of-round feedback in the R and F treatments included only own action and own payoff (in contrast with the original experiment, where opponent action was included). Sessions lasted 45–90 minutes, and average payments were roughly £14.

Twenty sessions of the new experiment were conducted, with a total of 286 subjects; some session details are shown in Table 5.¹⁹ Aggregate frequencies of risky-action choices are also shown in this table, along with significance results from Wilcoxon signed-ranks tests of differences between games within a treatment (one-tailed tests, group-level data).

Table 5: New experiment treatments and results

Treatment	Payoff information	Matching mechanism	Number of subjects	Number of groups	Frequency of risky action (all rounds)		
					NSHH	NSHM	NSHL
C	complete	random	90	7	.488	.481	.875**
CF	complete	fixed	78	7	.672	.845	.938**
R	limited	random	46	5	.251	.136**	.545**
F	limited	fixed	72	5	.337	.328	.654*

* (**, ***): Frequency significantly different from corresponding NSHH frequency at the 10% (5%, 1%) level, in direction predicted by loss avoidance (group-level data).

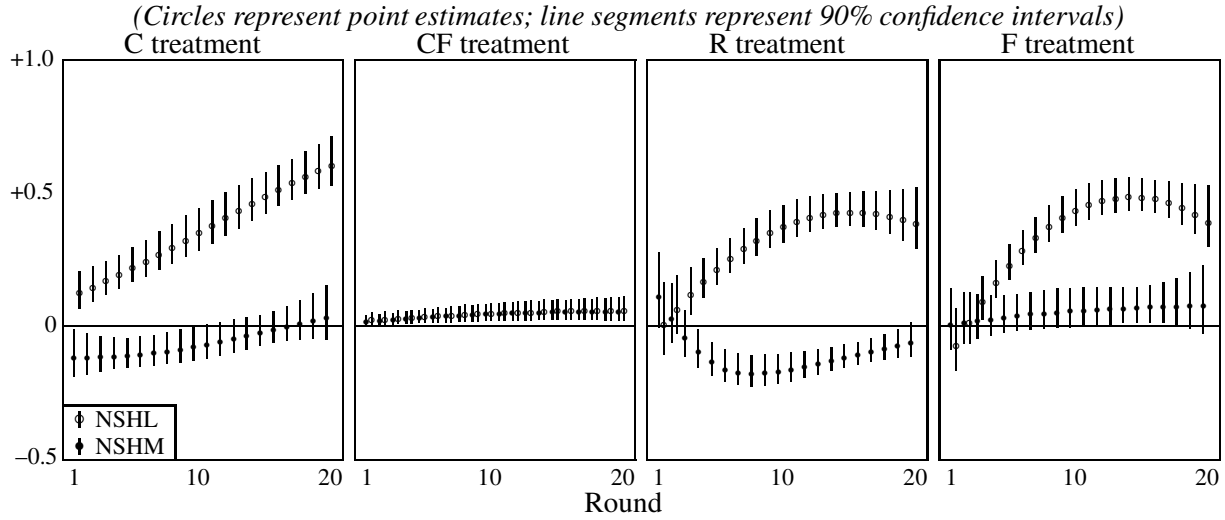
The table shows strong evidence of certain-loss avoidance. Within each treatment, the highest overall frequency of risky-action choices is in the low-payoff NSHL game. In three of the four treatments, the difference in frequencies between NSHL and NSHH is significant at the 5% level; in the fourth, significance is only at the 10% level ($p = 0.0625$). Evidence of possible-loss avoidance in Table 5 is weaker, but as in the original experiment, is strongest in the R treatment. In that treatment, overall risky-action choice frequencies are lower in NSHM than in NSHH (consistent with possible-loss avoidance), and the difference is significant at the 5% level. In the C and F treatments, risky-action frequencies are only slightly lower in the NSHM game than in the NSHH game, and in the CF treatment, the difference is in the other direction.

Figure 10 shows estimates of the effect of the level of certain- and possible-loss avoidance in the new exper-

¹⁹In Table 5, "group" refers to the smallest statistically-independent unit within each session. In most sessions, this is simply the entire session, but a few of the larger sessions were split into two sub-sessions, where subjects in a particular sub-session only interacted with other subjects in the same sub-session.

iment, based on probit results analogous to those used for Figure 7.²⁰ The trajectories in this figure are similar in

Figure 10: Estimates of the effect of game on risky–strategy choice (vs. NSHH) based on probit results



some ways to those in Figure 7. Evidence for certain–loss avoidance is again strong, as the estimated effect of the NSHL game is positive and significant in nearly all rounds of all treatments, though rather small in our new CF treatment relative to the others. Evidence for possible–loss avoidance is again weaker and varying across treatments. As in the original experiment, the strongest evidence for possible–loss avoidance is in the R treatment—as the estimated effect of the NSHM game is negative and significant in most rounds—and to a lesser extent in the C treatment (where the effect declines over time). In the F and CF treatments, as in the F treatment of the original experiment, we find no evidence of possible–loss avoidance, and what effect we do see is in the wrong direction (though small and sometimes insignificant).

To summarize, our new experiment—while differing in several design and procedural aspects from the original experiment—yields broadly similar results. The strong evidence in favor of certain–loss avoidance in all treatments is replicated, as are both the weaker overall support and the variation across treatments for possible–loss avoidance. In particular, we find the same association between observed evidence of possible–loss avoidance and strategic uncertainty as in the original experiment (see Section 3.4).

5 Conclusion

According to game theory, changing payoff levels (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’ expected utilities. However, there is a substantial body of evidence from individual decision–making experiments that changes in the level of *money payments* 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 an effect to exist in games too. In order to ascertain the

²⁰We estimated separate probit models for the four treatments, each with an indicator for a risky–action choice as the dependent variable. As before, our main explanatory variables were indicators for the NSHM and NSHL treatments and their products with the round number and its square. Other right–hand–side variables were the round number, its square, and indicators for the various game orderings. Additional details of the probit results are available from the corresponding author upon request.

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, though, they differ in the signs of the payoffs. In the high–payoff (SHH) game, only positive payoffs are possible. In the medium–payoff (SHM) game, the safe action leads to a positive payoff for sure, 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 for sure.

We consider two particular ways in which changes in payoff levels—specifically, in the sign of payoffs—can affect action choices in these games. *Certain–loss avoidance* refers to a 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 a 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.

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 information about the games’ payoffs. In a second treatment, they again have complete information, but play each game repeatedly 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.

Although our experiment is designed as a conservative test for loss avoidance, our results nonetheless show support for both certain– and possible–loss avoidance. The evidence for certain–loss avoidance is quite strong, and nearly uniform over all treatments—both in our original experiment and in a follow–up experiment we conducted in order to check for robustness. We find weaker, but still positive, evidence of possible–loss avoidance in both the original and the new experiment. We therefore conclude that both types of loss avoidance may be real factors in decision making, and though more research in this area is needed, we argue that researchers should take loss avoidance into account when considering how individuals behave in strategic situations when both gains and losses are possible.

References

- 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 (2010), "Experimental evidence of a sunk-cost paradox: a study of pricing behavior in Bertrand–Edgeworth duopoly," forthcoming, *International Economic Review*.
- 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), "Payoff levels, loss avoidance, and equilibrium selection in the Stag Hunt: an experimental study," working paper, University of Aberdeen.
- Fischbacher, U. (2007), "z–Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics* 10, pp. 171–178.

- Harsanyi, 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.
- Holt, C.A. and S.K. Laury (2002), "Risk aversion and incentive effects," *American Economic Review* 92, pp. 1644–1655.
- Holt, C.A. and S.K. Laury (2005), "Risk aversion and incentive effects: new data without order effects," *American Economic Review* 95, pp. 902–912.
- 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.
- 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), "Anomalies: 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.
- 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.