

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

Nick Feltovich¹, Atsushi Iwasaki², and Sobei H. Oda³

¹ University of Houston, Department of Economics
Houston, TX 77204-5019, USA
nfelt@bayou.uh.edu

² Kyushu University, Department of Intelligent Systems
Motooka 744, Nishi-ku, Fukuoka, 812-8581, Japan

³ Kyoto Sangyo University, Faculty of Economics
Kamigamo, Kita-ku, Kyoto 603-8555, Japan

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. However, this invariance is an empirical question when “payoffs” are actually money amounts rather than utility amounts. In particular, if individuals treat 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 a game called Stag Hunt, which are identical up to the level of payoffs, under a variety of treatments. We find differences in behavior across the three games; 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 equilibria in such games.

Keywords. experiment, game theory, behavioral economics, stag hunt, learning

1 Motivation

Game-theoretic solution concepts imply that modifying a game by adding the same positive or negative constant to all payoffs—a “change in payoff levels”—should not affect behavior. If 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. If 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 payoff-level changes affect behavior, and at most explicitly rules out such an effect.

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 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 [3], whose “prospect theory” was designed to account for the systematic violations of standard expected-utility theory they observed in a large number of decision-making tasks. Although they did not look specifically at payoff-level changes, prospect theory implies they will affect decision making. More recently, [1] found that decisions in game experiments were 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 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 alternatives that could yield positive payoffs.

We will refer to Cachon and Camerer’s notion of loss avoidance as “certain-loss avoidance”, and draw a distinction between this and “possible-loss avoidance”—a tendency to avoid strategies that give a *possible* negative payoff, in favor of alternatives that give a certain positive payoff. The goal of this paper is to test for these two forms of loss avoidance. We design and run a human-subjects experiment using three versions of Stag Hunt [4], a symmetric two-player game with two strategies: one risky and one safe. The games are shown in Figure 1 (with the strategies labeled R and S, respectively). They differ only in payoff level,

	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)

Fig. 1. The Stag Hunt games used in the experiment

so are game-theoretically equivalent. Nonetheless, our notions of loss avoidance make predictions regarding how individuals’ choices change across these 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), 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, R should then be less attractive in SHM than in SHH, so that R is

less likely to be chosen in SHM than in SHH. Similarly, in the low-payoff game (SHL), choosing S leads to a certain negative payoff, but choosing R could lead to a positive payoff, so that an individual exhibiting certain-loss avoidance will find S less attractive here than it would have been in SHH, so that S is less likely (so that R is *more* likely) to be chosen in SHL than in SHH.

2 Experimental design

In our experiment, subjects played all three versions of Stag Hunt, as well as three additional games. We do not analyze the data from these other games here; they were added simply to hide from subjects the similarity of the three Stag Hunt games.

We used four design treatments. In our O (one-shot) treatment, subjects played each game once and the payoff matrices were publicly announced. 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 deductive reasoning. Our C (complete-information) treatment was similar to our O treatment, except subjects played each game 40 times before moving on to the next game, and 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 due to subjects' experience in playing the game—in particular, whether loss avoidance grows 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 that would enable them to infer the payoff matrices eventually. 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 like the R treatment, except that subjects were matched to the same opponent for all 40 rounds of a game. The R and F treatments allow us to see the effects on loss avoidance of the experience subjects receive in playing each game, without the benefit of any initial introspection. 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.

3 Main results

We first look at initial behavior of subjects playing the Stag Hunt games under complete information, by examining the O treatment (in which each game was played only once) and the first round of each game in the C treatment. Table 1 shows these frequencies of risky-action choices for 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 SHL,

nonnegligible but insignificant ($\chi^2 \approx 1.42$, d.f.=1, $p > .10$) in SHM, significant ($\chi^2 \approx 3.93$, d.f.=1, $p < .05$) in SHH.

Table 1. Initial frequencies of risky action (complete information treatments, all sessions)

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

Next, consistent with certain-loss avoidance, risky-action play is substantially more prevalent in SHL than in SHH. 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 SHH (69.4% of the time in the O treatment, 84.7% in the first round of the C treatment). 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 change test, $p < 0.001$). (See [5] for descriptions of the nonparametric statistical tests used in this paper.) On the other hand, initial behavior shows no evidence in favor of 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 SHH and SHM 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 change test, $p > 0.10$).

We next look at how subject behavior changes over time. Figures 2 and 3 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. (See [2] for a more thorough and rigorous analysis of these data.) Figure 2 suggests that loss avoidance under complete information has only a minor effect. 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, there is some divergence in play across games over time. 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 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,

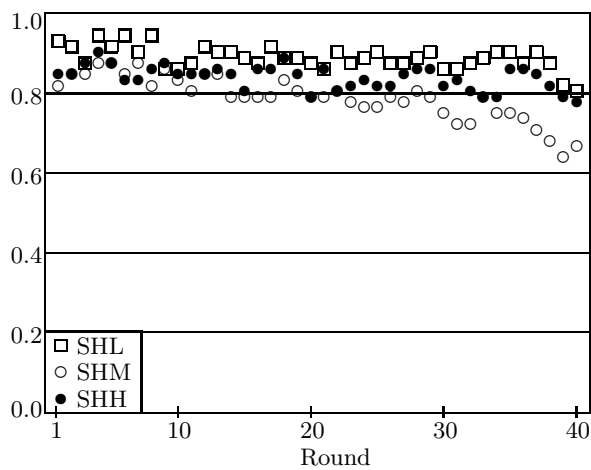


Fig. 2. Frequency of risky action (complete-information treatment, all sessions)

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.

4 Discussion

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

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

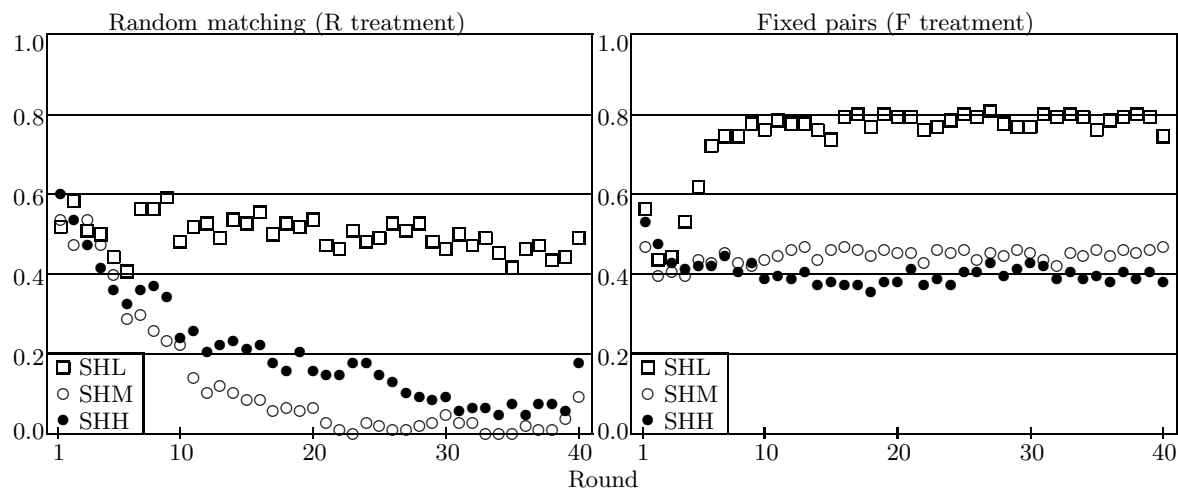


Fig. 3. Frequency of risky action (limited-information treatments, all sessions)

2. Feltovich, N., Iwasaki, A., Oda, S.H.: Payoff levels, loss avoidance, and equilibrium selection in the Stag Hunt: an experimental study. Working paper, University of Houston (2007)
3. Kahneman, D., Tversky, A.: Prospect theory: an analysis of decision making under uncertainty. *Econometrica* **47** (1979) 273–297
4. Rousseau, J.J.: *The Social Contract and Discourses*. J.M. Dent & Sons, London (1973)
5. Siegel, S., Castellan, Jr., N.J.: *Nonparametric Statistics for the Behavioral Sciences*. McGraw-Hill, New York (1988)