Coordination and Transfer

David J. Cooper
Department of Economics
Florida State University
djcooper@fsu.edu

John Van Huyck
Department of Economics
Texas A&M University

February 20, 2017

Abstract: We study the ability of subjects to transfer principles between related coordination games. Subjects play a class of order statistic coordination games closely related to the well-known minimum (or weak-link) and median games (Van Huyck, Battalio, and Beil, 1990 and 1991). When subjects play a random sequence of games with differing order statistics, play is less sensitive to the order statistic than when a fixed order statistic is used throughout. This is consistent with the prediction of a simple learning model with transfer. If subjects play a series of similar stag hunt games, play converges to the payoff dominant equilibrium when a convention emerges, replicating the main result of Rankin, Van Huyck, and Battalio (2000). When these subjects subsequently play a random sequence of order statistic games, play is shifted towards the payoff dominant equilibrium relative to subjects without previous experience. The data is consistent with subjects absorbing a general principle, play of the payoff dominant equilibrium, and applying it in a new related setting.

Author's Note: The research in this paper is the result of a wager. In 1997 – 1998, John Van Huyck visited the University of Pittsburgh where I was an assistant professor. He gave a talk about the experiments subsequently published in Rankin, Van Huyck, and Battalio (2000). This paper is concerned with the emergence of conventions in stag hunt games. Standard wisdom held that repeated play of a fixed stag hunt game leads to emergence of the risk dominant equilibrium as a convention.[1] Rankin *et al.* studied play in a series of closely related stag hunt games. While the games were similar, the location of equilibria and payoffs were randomized with the intent of making the direct relationship between actions less salient. Subjects were instead forced to rely on general features of the games like security, payoff dominance, and risk dominance. The surprising result is that the payoff dominant equilibrium emerges as a convention if any convention emerges.

During John's lecture I asked whether this result would extend to other coordination games. For example, what would happen if a group which had converged to the payoff dominant equilibrium in the stag hunt games was then asked to play a minimum game (Van Huyck, Battalio, and Beil, 1990)? John felt that such groups would successfully coordinate on the payoff dominant equilibrium while I conjectured that play would collapse to the secure equilibrium as is typical for minimum games with large groups. We bet $5 on who was correct.

The bet was settled by running a single experimental session at Texas A&M. Two groups of eight subjects played the same series of closely related stag hunt games used by Rankin *et al.* and then played a series of eight person minimum games. The efficient equilibrium emerged as a convention in the stag hunt games for both groups, but play in the minimum games rapidly converged to the secure equilibrium. Although the bet was settled, John and I agreed that the central intellectual issue was not. The minimum game is a demanding environment. Even if most subjects learn a convention of playing the payoff dominant equilibrium in the stag hunt games *and* apply this convention in the minimum games, a few exceptions who do not learn the convention of playing the payoff dominant equilibrium or do not apply it in the minimum games are sufficient to undo the convention. Thinking more deeply about the nature of conventions in coordination games, additional difficulties emerge. Following a convention of playing the payoff dominant equilibrium requires that you believe others will follow the convention. For a stag hunt game you only need to believe that one person, your current opponent, will follow the convention. In a minimum game, you need to believe that *everyone* in your group will follow the convention. This is a far more challenging standard.[2] John and I agreed that the test we had run for our bet made it unlikely that transfer of conventions would be observed.

---

[1] See Ochs (1995) for a summary of evidence at the time.
[2] This point is a simple matter of arithmetic. Let p be the probability that a subject plays the payoff dominant equilibrium. In choosing whether I should play the payoff dominant equilibrium in a stag hunt game, the relevant number is my belief about p. In an eight person minimum game, what matters is my belief about p^7 which is a much smaller number. For example, suppose p = .8. The probability the one person plays the payoff dominant equilibrium is 80%, but the probability that a group of seven people all play the payoff dominant equilibrium (assuming independence) is only 21%.

We therefore designed a more forgiving test of the central hypothesis that subjects would apply a convention learned in stag hunts to other coordination games where it is relevant. A small number of initial sessions were run at Texas A&M, but life intervened and the project was never completed. When John passed away, I decided that the best way to honor him would be to use the design we had developed and finally complete the project. With an enormous amount of help from Catherine Eckel, Phil Brookins, Laura Magee, and Joe Stinn, I was able to reconstruct the experiments that had been run at Texas A&M and run new sessions at FSU. The paper that follows presents the results of this exercise.

Before getting to the details of the paper, I want to talk about John. He was an exceptional researcher and an exceptional colleague. He never did research as a fishing expedition. John's work was based on a clear understanding of the relevant theory. He wasn't interested in work that was flashy or trendy. John had a deep interest in understanding how people played games, particularly coordination games, and did research with the aim of addressing fundamental issues. John understood the importance of dynamics in a way that many experimenters do not. Equilibrium doesn't spring into existence like Minerva emerging from the head of Jupiter; it develops gradually through a dynamic process. Much of John's work, including this paper, is devoted to understanding that process. John was an incredibly deep thinker. He understood the links between games and economic settings and the relationship between game theory and experiments in a way that very few researchers have achieved. He was brilliant. John also did the little things right. Everybody talks a good game about careful experimental techniques, but John conducted experiments with extraordinary care. He paid attention to little details that the referees will never know about because it is the right way to do science. As a person, John was a fascinating guy. He was frighteningly smart, and had tremendous intensity whenever a topic of conversation caught his interest. He had a wonderful sense of black humor. He cared deeply about taking care of his students and even more deeply about taking care of his family. Academia is a demanding business and it is easy to let the job devour your life. John was passionately interested in experimental economics, but he never forgot that we are more than our jobs. It was an honor to be his friend and colleague.

"If we rely solely on the story of directly relevant experience to justify attention to Nash equilibria, and if by directly relevant experience we mean only experience with precisely the same game in precisely the same situation, then this story will take us very little distance outside the laboratory." Kreps, 1990

1. Introduction: Nash equilibrium (Nash, 1950) is a fundamental concept in game theory. As the name implies, it is an equilibrium concept and does not specify how play arrives at an equilibrium. This is a critical issue in applying the theory to field settings, especially when more than one Nash equilibrium exists. Learning models, which conceptualize equilibrium as the steady state of a dynamic process, offer a good explanation for the emergence of equilibrium. Theorists have enjoyed great success showing conditions under which learning processes will

converge to equilibrium.[3] Experimenters have also studied learning in games extensively. They have shown the existence of strong dynamics in game play, characterized the learning rules being used, and established conditions under which play converges to an equilibrium.[4]

The experimental literature on learning in games typically focuses on repeated play of the same game. This begs the question raised by Kreps in the opening quote. Relatively little work has been done on the issue of transfer, the ability to take information learned in one game and apply it in another related game.[5] The purpose of our experiments is to study transfer between games, particularly in an experimental design that forces subjects to rely on general principles rather structural similarity between games. Examples of structural similarity between games include having the same number of actions, the same location for Nash equilibria, or the same locations of secure and Pareto efficient outcomes.[6]

Our experiments employ two different types of coordination games, stag hunt games and order statistic games. Coordination games are a good environment to study transfer. It is simple to construct games that have similar structures but generate different behavior when played in isolation (e.g. minimum and median games). There are general principles (payoff dominance, risk dominance, and security) that apply to multiple classes of coordination games with distinctly different structures (i.e. different number of actions, payoff structures, and/or locations of equilibria). This makes it possible to study whether subjects are transferring general principles between games rather than merely learning an action and using it in an obviously related game.

Our experiment contains three treatments. The main treatment grew directly out of the bet between John and me. Subjects initially play the series of 75 similar stag hunt games used by Rankin *et al.* in fixed cohorts of eight subjects. They then play a series of 14 "random order statistic" games with their cohort. These are eight player coordination games where subjects simultaneously choose a number between 1 and 7. A subject's payoff is maximized if he matches the OS[th] lowest number chosen by the *other* seven members of his group (OS stands for "order statistic"). The value of OS ranges between 1 and 7 and is randomly drawn for each game. If OS = 1, this is almost the same as the well-known minimum game.[7] Likewise, OS = 4 is (almost) a median game. The structures of the random order statistic games are similar to each other but quite different from the stag hunt games. General principles like payoff dominance and risk dominance apply to both the stag hunt and random order statistic games.

---

[3] See Fudenberg and Levine (2009) for a recent survey.

[4] See chapter 6 of Camerer (2003) and Erev and Haruvy (forthcoming) for surveys.

[5] A related issue is the effect of spillover, the effect on play of participating in related games *simultaneously* (Bednar, Chen, Liu, and Page, 2012; Cason, Savikhin, and Sheremeta, 2012; Savikhin and Sheremeta, 2013; Cason and Gangadharan, 2013). Because play is simultaneous, effects in one game cannot rely on feedback about play or outcomes in the other. Hence, learning does not play a role in spillover effects.

[6] A major difficulty in the literature on transfer is the absence of any clear metric for similarity between games. Our definition of structural similarity is a little bit like Stewart's famous definition of pornography: I know it when I see it. The basic idea is that there exist elements of a game's construction that make the similarity of two games transparent independent of any deeper principles.

[7] The games differ because your actions do not affect the minimum you are trying to match. You therefore cannot not guarantee matching the minimum by choosing 1.

The other two treatments serve as controls. In the first, subjects play a sequence of random order statistic games *without* previous experience in stag hunt games. The second control treatment has subjects playing a sequence of order statistic games where OS is *fixed* rather than randomly changing. In other words, this treatment varies OS between groups but keep it fixed within groups.

A simple model of learning with transfer predicts that the response of subjects' choices to OS will be flatter in the control treatment where the value of OS changes randomly than the control treatment where OS is fixed within groups. Given the previous results of Rankin *et al.*, we expect subjects in the stag hunt phase of the main treatment to converge to the payoff dominant equilibrium (subject to some convention emerging). Assuming that the concept of payoff dominance transfers between games, we predict that choices in the random sequence of order statistic games will be shifted upward for the main treatment where subjects have previous experience with stag hunt games relative to the control treatment where subjects have no previous experience in similar games.

The results are in line with our hypotheses. The most important result is the presence of positive transfer between the stag hunt games and the random order statistic games. Relative to the control treatment, choices in the random sequence of order statistic games are higher when subjects have previous experience in the stag hunt games. Choices are highest in cohorts that had previously converged to the payoff dominant equilibrium in the stag hunt games. A closer look at the data indicates that initial play in the stag hunt games has little predictive power for the random order statistic games. Play in the random order statistic games depends on what subjects learn in the stag hunt games rather than their initial behavior. In other words, the effect does *not* appear to be driven by subjects' types. Subjects' play is consistent with them learning a principle (payoff dominance) and then transferring it to a different game.

Comparing behavior in the two control treatments, the response of subjects' choices to the value of OS is flatter when the value of OS changes randomly between rounds. To our surprise, subjects' use of feedback is *not* sensitive to how close previous values of OS are to the current value. Subjects treat past experience as being more relevant to the current game than is warranted. With experience, play is shifted upwards for all values of OS when the value of OS is fixed relative to when OS changes randomly. This shift is extremely large for high values of OS. There is strong transfer between games when subjects play randomly changing order statistic games, but this transfer does not lead to more efficient outcomes.

Our experimental results make a valuable contribution to the literature on transfer in games. Two strands of the literature are particularly relevant for our work. The first concerns the role of precedence in coordination games. Van Huyck *et al.* (1991) study whether a precedent established in one version of a median game would carry over to another version of the median game. They find little evidence of transfer. Other studies have yielded more positive results. Devetag (2005) studies transfer from critical mass games to minimum games. The critical mass game typically yields efficient coordination while play in the minimum games usually collapses to the secure equilibrium. Devetag finds that a precedent of efficient play in the critical mass

game carries over to the minimum games with play shifting toward the payoff dominant equilibrium. Brandts and Cooper (2006) study a design where subjects initially play a minimum game with payoffs very unfavorable to coordination above the minimum. Players are then switched to a version of the game with favorable payoffs for coordination at higher effort levels. If the payoff table is returned to the original harsh environment, play does not collapse back to the minimum effort level. Instead, there is a persistent positive effect from having experienced coordination at high effort levels. Weber (2006) studies the precedent established in moving from minimum games with a small number of players (where efficient coordination is common) to games with larger numbers of players. He finds that efficient coordination is maintained as the group size grows. Cason, Savikhin, and Sheremeta (2012) study transfer between median and minimum games. They find strong positive transfer from the median games to subsequent play in minimum games. Interestingly, no such transfer is found from minimum games to subsequent median games and there is little spillover when the games are played simultaneously. Overall, the bulk of the evidence points to positive transfer between coordination games with a precedent of efficient play carrying over to environments where efficient play is more difficult.

A second strand of the literature looks at whether subjects can transfer general principles. The evidence on this count is mixed. Ho, Camerer, and Weigelt (1998) find positive transfer, but of an odd variety. Subjects are switched between different p-beauty contest games, with iterated removal of dominated strategies being the relevant strategic concept in all cases. Subjects switched to a new game initially play no differently from inexperienced subjects but learn faster. Cooper and Kagel (2005 and 2008) study the ability of subjects to take principles learned in one signaling game and apply them in another. The action space and general principles are the same across games, but the form of the equilibrium (pooling vs. separating) and the actions used in equilibrium differ. The results are mixed, with negative transfer reported in the 2005 paper and positive transfer reported in the 2008 paper. Subsequent work (Cooper and Kagel, 2009) shows that the differing results can be attributed to the use of abstract versus meaningful context to frame the games. Haruvy and Stahl (2012) study play in a series of dissimilar 4 x 4 normal form games. They find that subjects learn rules and transfer these rules between games.[8]

The papers studying transfer between games share several common features: (1) With the exception of Haruvy and Stahl, the games being played are structurally similar. There are obvious cues that the games are related and obvious mappings of actions from one game into the other. (2) For papers studying transfer of general principles, the principles involved are strategic principles such as not playing dominated strategies, rather than principles of equilibrium selection such as choosing a payoff dominant equilibrium. (3) Little attention is paid to determining what features make subjects treat games as more or less similar to each other.

Our work makes several contributions to the existing literature. We find that precedents transfer between coordination games, as in previous papers, but with two important differences.

---

[8] For additional related papers see Rick and Weber (2010) and Grimm and Mengel (2012). More broadly, there is a large psychology literature on learning and transfer with generally discouraging findings. See Gick and Holyoak (1983) for a classic article on the topic and Alfieri, Nokes-Malach, and Schunn (2013) for a recent meta-study.

First, transfer does *not* necessarily lead to more efficient outcomes. In the control treatment with a random sequence of order statistic games there is strong transfer between games which harms payoffs relative to the treatment where OS is held fixed. Second, the random order statistic games have a different number of actions, a different payoff structure, and different locations for equilibria than the stag hunt games. Like Rankin *et al.*, we have deliberately reduced the ability of subjects to rely on structural similarity rather than general principles. The positive transfer between stag hunt and random order statistic games shows that structural similarity is *not* a necessary condition for precedents to influence play in coordination games.

We also contribute to the literature on the transfer of general principles. Our data is consistent with subjects learning and applying a general principle. Transfer of general principles has been observed before, but our example is unusual because the stag hunt and random order statistic games have little structural similarity. This should increase our optimism about the ability of subjects to learn and transfer general principles. Our experiments are also novel because the principle being transferred is not a principle of strategic play but rather an equilibrium selection concept.[9] Finally, our subjects are surprisingly insensitive to differences between games. This suggests that a central feature of the ABEE concept (Jehiel, 2005), the treatment of broad classes of games as equivalent, may be realistic.

2. Experimental Design and Hypotheses

A. Games: Two types of coordination games were used in our experimental design, random order statistic (ROS) games and stag hunt games. The ROS games are played in eight player groups. An order statistic $OS \in \{1, 2, 3, 4, 5, 6, 7\}$ is announced to all players prior to the start of the game. Players then simultaneously choose a number $X \in \{1, 2, 3, 4, 5, 6, 7\}$. The payoff for Player i is given by Equation 1. The variable X_i is the number chosen by Player i and X_{-i} is the vector of choices by the other seven group members. The function $EOS(X_{-i}, OS)$ gives the OS^{th} order statistic for the choices of the *other* seven people in the group, where $OS = 1$ yields the lowest choice of the other seven group members, $OS = 2$ yields the second lowest choice, and so on. This is an "exclusive" order statistic and does *not* depend on the choice of Player i. ***Throughout the remainder of this paper, if we refer to the "order statistic" we mean OS, the parameter determining what game is being played. If we refer to the "exclusive order statistic" we mean EOS, the function giving the OS^{th} order statistic for the other seven people in the group.*** [10] The values of a, b, and c are constants. For all of our sessions we used the parameters from the Γ payoff table in Van Huyck, Battalio, and Beil (1991) multiplied by a factor of 200: a = 20, b = 10, and c = 120.

---

[9] See Duffy and Fehr (2015) for a case where equilibrium selection principles fail to transfer.

[10] Use of an *exclusive* order statistic has several advantages. Most important, a player's choice does not influence the value of $EOS(X_{-i}, OS)$. Play of the maximum game, for example, would be trivial if the player's own choice is included in the calculation of $EOS(X_{-i}, OS)$. Using the order statistic also makes it possible to use the same payoff table for all seven games. Otherwise the minimum and maximum games require different payoff tables.

$$\textbf{(Eq 1)} \ \pi(X_i, X_{-i}) = aEOS(X_{-i}, OS) - b\big(X_i - EOS(X_{-i}, OS)\big)^2 + c$$

Table 1 shows the resulting payoff table for the ROS game. Taking the value of OS as given, coordinating on any of the seven possible choices is a pure strategy Nash equilibrium. There are no other pure strategy equilibria. The seven equilibria are Pareto ranked with mutual play of 7 being the payoff dominant equilibrium. The maximin (or secure) choice is 3 regardless of the value of OS.

Table 1: Payoff Table for the ROS Game

		EOS (Exclusive Order Statistic, Other Seven Group Members)						
		1	2	3	4	5	6	7
My Choice	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

To get a better idea of how play can be expected to vary as OS changes, suppose a subject believes the other seven group members are choosing uniformly over the seven possible choices. Table 2 reports the expected payoff for each possible choice as a function of OS, the order statistic. The best response has been highlighted in yellow for each value of OS. In each case the best response is equal to the order statistic. In the absence of a clear historical precedent, there are strong incentives to move to higher choices as the value of OS increases.

Table 2: Expected Payoffs vs. Random Play

		Value of OS						
		1	2	3	4	5	6	7
My Choice	1	141.7	138.5	123.7	95.4	53.4	-1.6	-61.8
	2	140.9	153.6	156.2	145.5	121.0	83.4	39.1
	3	120.0	148.7	168.8	175.6	168.6	148.5	120.0
	4	79.2	123.8	161.3	185.6	196.2	193.5	180.8
	5	18.3	78.9	133.9	175.7	203.8	218.6	221.7
	6	-62.5	14.0	86.4	145.8	191.4	223.6	242.6
	7	-163.4	-70.9	19.0	95.9	159.0	208.7	243.5

The stag hunt games used in our experiments were identical to those employed by Rankin *et al*, subject to scaling the payoffs down. The basic payoff table is shown in Table 3. The value of

w always equaled 370, y was drawn from a uniform distribution over [1, 2, …, 369], and ε was drawn from a uniform distribution over [1, 2, …, 50]. Subjects knew the payoff table prior to play. There are two pure strategy Nash equilibria, mutual play of A or mutual play of B. The former is always the secure or maximin equilibrium and the latter is the payoff dominant equilibrium. The identity of the risk dominant equilibrium depends on the value of y. If y > 185, the risk dominant equilibrium corresponds to the secure equilibrium. Otherwise the payoff dominant equilibrium is also risk dominant.

Table 3: Payoff Table for Stag Hunt Games

	A	B
A	$y + \varepsilon, y + \varepsilon$	$y + \varepsilon, \varepsilon$
B	$\varepsilon, y + \varepsilon$	$w + \varepsilon, w + \varepsilon$

The most important way in which the stag hunt games varied was whether risk dominance corresponds to security or payoff dominance, but additional steps were taken to obscure the relationship between the games. The variable ε is a constant added to all payoffs. Making the payoff for the payoff dominant equilibrium vary from round is designed to reduce the transparent relationship between games. We also randomly flipped whether A corresponded to the first or second row/column. This randomly switches the location of the secure and payoff dominant equilibria between the upper left and lower right corners. As Rankin *et al* note, the changing payoffs and scrambled locations of the equilibria are "an effort to blunt the salience of retrospective selection principles."

B. Treatments: The design contains three treatments. Subjects participated in one of the three treatments. In other words, this is a between subjects design.

Random Order Statistic (ROS) Treatment: Subjects were formed into fixed groups of eight subjects ("cohorts") and played 42 rounds of the ROS game. There were six blocks of seven games, with each value of OS ∈ {1, 2, 3, 4, 5, 6, 7} used once in each block. The order of values for OS was randomly scrambled within each block. The order was drawn in advance and the same order was used in all cohorts. The order for the first and sixth blocks was forced to be identical. The instructions state, "The order statistic will change from period to period." No details were given to subjects about how the order statistic would be changing. Subjects knew the number of rounds in the experiment.

Fixed Order Statistic (Fixed) Treatment: We were interested in how playing a sequence of related but different order statistic games affected convergence relative to playing the same order statistic game throughout. The Fixed treatment provides the relevant point of comparison. As in the ROS treatment, subjects were formed into fixed groups of eight subjects ("cohorts") and

played 42 order statistic games. Unlike the ROS treatment, the order statistic is fixed for all 42 games. The instructions tell subjects that, "The order statistic will be chosen by the computer at the beginning of the session for each group of eight participants [cohort] and will remain constant during the session." Cohorts were run using all of the order statistics except OS = 7. Given that play rapidly converges to the payoff dominant equilibrium (mutual choice of 7) for fixed order statistics of 5 or 6, there wasn't much to be learned by running cohorts with OS = 7.

Stag Hunt → Random Order Statistic (SHROS) Treatment: This is the central treatment in our experimental design. If subjects learn a convention of playing the payoff dominant equilibrium in a series of related stag hunt games, can they transfer that convention to the ROS game?

At the beginning of the session, subject are formed into fixed groups of eight ("cohorts"). The instructions tell subjects that they will interact within their cohort throughout the experiment. Subjects initially play 75 rounds of the stag hunt game. Payoffs were perturbed as described above, with the order of games drawn from Rankin *et al*. Subjects are matched in pairs to play the stag hunt games. The match was round robin, meaning that subjects played each of the seven other members of their cohort once in each seven round block.[11] The order over the pairings was randomly scrambled for each block. The instructions for the stag hunt games (Part 1) inform the subjects that there will be a second part of the experiment, but give no details on what Part 2 will be.

After the stag hunt games were completed, subjects played 14 rounds of the ROS game in the same cohorts of eight subjects. The instructions stress that subjects will "…participate in the **same group** of eight people [cohort] as in Part 1 of the experiment." The order of values of OS was the same as the first 14 rounds of the ROS treatment. Thus, the order of values of OS for the first and sixth seven round blocks of the ROS treatment was identical to the order for the first block of the SHROS treatment and the order for the second blocks of the ROS and SHROS treatments were identical.

C. Procedures: All sessions were run at the xs/fs laboratory at Florida State University. Subjects were recruited from the FSU undergraduate population using ORSEE (Greiner, 2015). Sessions were run using z-Tree (Fischbacher, 2007). The instructions and computer interface were designed to mimic as closely as possible those used by John Van Huyck in the original sessions conducted at Texas A&M, but it was not possible to use his original software. Subjects were paid their total earnings from all rounds plus a $7 show-up fee. Payoffs were denominated in points, with a conversion rate of 400 points equal to one dollar. Sessions lasted 60 – 90 minutes and average earnings (including the show-up fee) were $28.54.

All sessions began with detailed instructions (see online appendices for instructions from the SHROS treatment). For the SHROS treatment, subjects were given two sets of instructions, one before the stag hunt games and one before the ROS games. The payoff tables and the feedback tables are relatively complex for the ROS games, so subjects were given detailed examples

---

[11] This was truncated in the final block as there were only five rather than seven rounds.

beyond the instructions using a script and PowerPoint slides displayed at the front of the room. Following the instructions (both sets for the SHROS treatment), subjects took a quiz testing their ability to read the payoff tables.

The payoff table for the current round was displayed on the computer screen when subjects were making choices. After all choices were made, subjects were given detailed feedback on the outcome for the round. This included a summary of all eight choices for the ROS games. Subjects were shown a history table when making choices which gave a summary of the games played and outcomes for previous rounds. Subjects did not know the identity of the other seven members of their cohort and feedback was given in such a way that choices could not be associated with any specific cohort member.

Because subjects in the SHROS treatment had accumulated payoffs in the stag hunt games before beginning play of the ROS games, it is possible that any treatment effects reflect income effects. To control for this, prior to play of the ROS games (and after play of the stag hunt games in the SHROS treatment) all subjects received a randomly determined bonus. Bonuses were independently drawn across subjects and were uniformly distributed over the interval between 0 and 2000 points (0 and 5 dollars). The instructions tell the subjects, "Your bonus is random and has no relationship to your choices in the experiment or any of your personal characteristics." We use the bonuses as an instrument for the cash balance (defined as the sum of the bonus and any accumulated payoffs) entering the ROS games. Bonuses are positively correlated with the cash balance but independent from any individual characteristics or choices of the subjects.

Table 4: Summary of Treatments

Treatment	Order Statistic (OS)	# of Cohorts (Sessions)
Random Order Statistic (ROS)	Random (1 – 7)	14 (5)
Stag Hunt → Random Order Statistic (SHROS)	Random (1 – 7)	12 (5)
Fixed Order Statistic (Fixed)	1	3 (3)
	2	3 (3)
	3	3 (3)
	4	3 (3)
	5	3 (3)
	6	3 (3)

Table 4 summarizes the data that was gathered for this project. There were seventeen sessions with a total of 352 subjects. A "cohort" is a fixed group of eight subjects. Subjects only interacted with other subjects in their cohort, so cohorts are independent observations.[12]

---

[12] For the Fixed treatment, there were seven sessions. Each session in this treatment included cohorts with different values of the order statistic (OS), with no two cohorts from the same session having the same order statistic. This

Two sessions had problems. For one session of the SHROS treatment (2 cohorts), the payoffs for the stag hunt games were accidentally multiplied by a factor of ten. This had no obvious effect on play for either the stag hunt or ROS games, so the data from this session has been included in the dataset. For one session of the ROS treatment (3 cohorts), two subjects had to leave after the first two blocks (14 rounds) had been completed.[13] Data from the two blocks that were successfully completed is included in the dataset. None of our qualitative conclusions are affected if the data from these two problem sessions is dropped.[14]

D. Hypotheses: As noted in the discussion of Table 2, there are strong incentives to choose higher numbers in the ROS game as OS increases. Previous evidence with related order statistic games (Van Huyck, Battalio, and Beil, 1990 & 1991) also suggests that higher order statistics will lead to convergence at higher numbers. These observations lead to a pair of straightforward hypotheses.

H1: Play in the Fixed treatment will converge to an equilibrium. The average number chosen will be an increasing function of the order statistic (OS) in all treatments.

To form hypotheses for the ROS treatment, consider a version of the Roth-Erev learning model (Roth and Erev, 1995) modified to allow for transfer between similar games. Let w_{ijt} be the weight put on action i, i ∈ {1,2,3,4,5,6,7}, by individual j in round t. Let p_{ijt}, the probability of individual j choosing action i in round t, be given by Equation 2.

$$\textbf{(Eq. 2)} \;\; p_{ijt} = \frac{w_{ijt}}{\sum_{k=1}^{7} w_{kjt}}$$

Let w_{ij1} be the initial weight on action i for individual j. The weights in subsequent rounds are given by Equation 3. The term $I(OS(t) \neq OS(s))$ is a dummy for whether the values of OS in rounds s and t are different, $I(a(s) = i)$ is a dummy for whether action i was chosen in round s, and π_{js} is the payoff earned by individual j in round s. The variable σ measures perceived dissimilarity between the current order statistic game and other order statistic games. By assumption, $0 \leq \sigma \leq 1$. If σ = 0, past experience from all order statistic games is treated as equivalent. If σ = 1, only experience with the *same* order statistic game is given any weight. For simplicity, the dissimilarity weight σ does not vary between values of OS ≠ OS(t). Equation 3 says that the weight put on action i equals its initial weight plus the sum of payoffs from playing action i in previous rounds *weighted by whether or not the same order statistic game was being played*.

---

was done to avoid confounding any session effects with effects of the order statistic. Table 4 shows three sessions for each value of OS. This reflects the cohorts for each value of OS being spread across three sessions.

[13] One subject had an exam and one was worried about being late to his job. Although both had been told in advance about the time commitment needed for the experiment, both were allowed to leave and paid for the portion of the experiment they had completed as per IRB regulations.

[14] The average earnings reported above do not include the two problematic sessions.

$$\textbf{(Eq. 3)} \ \ w_{ijt} = w_{ij1} + \sum_{s=1}^{t-1}\left[\left(1 - \sigma I(OS(t) \neq OS(s))\right) * I(a(s) = i) * \pi_{js}\right]$$

For simplicity assume that all players have the same initial weights on the seven actions and assume that initial weights do not depend on the value of OS for either Round 1 or Round t.[15] Let $f(EOS; OS)$ be the distribution over exclusive order statistics (EOS) for Round 1 subject to the value of OS used for Round 1. If $A > B$, $f(EOS; A)$ first order stochastic dominates $f(EOS, B)$. From Equation 1, it can be shown that $\Delta\pi/\Delta EOS$ is an increasing function of X_i, player i's choice. It follows that as OS increases, the expected payoffs for high actions in Round 1 are increased relative to low actions. This is a fancier version of the point made by Table 2.

What does this imply for play in Round 2? Equation 3 implies that the expected increase in weight for Round 2 is proportional to the expected payoff in Round 1. Noting that the values of OS must be different for Rounds 1 and 2 by design, the higher the value of OS drawn in round 1, the more the weights in Round 2 are biased towards higher actions. For low values of OS in round 2 this implies an upward bias in the probability distribution over actions and for high values of OS it implies a downward bias. The same argument follows for later rounds. Relative to play with a fixed order statistic, actions should be higher for low order statistics and lower for high order statistics.

The intuition is simple. Suppose I rely on experience from games with other values of OS. In a game with a low value of OS, that experience probably comes from games with higher order values of OS where high actions tend to do well. I will overestimate how well a high choice is likely to do and bias my choice for the current round upwards. Similar logic implies that I will bias my choice downwards when playing a game with a high value of OS.

H2: The slope of the relationship between OS and subject choices will be flatter in the ROS treatment than in the Fixed treatment.

For the SHROS treatment to be interesting, a necessary condition is that we replicate the main result of Rankin *et al.*

H3: If cohorts in the stag hunt portion of the SHROS treatment converge to a convention, they will converge to play of the payoff dominant equilibrium.

There is no obvious mapping between actions in the stag hunt games and actions in the ROS games. Rankin *et al.* changed payoffs and scrambled the locations of the equilibria to reduce the similarity between games, forcing subjects to rely on general principles. Switching from the stag hunt games and the ROS games is a more extreme version of this. To the extent that experience from the stag hunt games is applicable in the ROS games, it is through the principle of playing the payoff dominant equilibrium. The payoff dominant equilibrium (mutual play of 7) is easily

---

[15] These are simplifying assumptions but not critical for our conclusions.

identified is the ROS game, so applying this principle is not difficult. Ideally subjects would play the payoff dominant equilibrium in the ROS games, but realistically we conjectured that prior play of the stag hunt games would shift play toward the payoff dominant equilibrium in subsequent plays of the ROS game.[16]

H4: Compared with the ROS treatment, play in the SHROS treatment will be shifted toward higher choices.

3. Results

A. Stag Hunt Games: A necessary condition for the SHROS treatment to be interesting is a general tendency to converge to the payoff dominant equilibrium. In other words, we want to confirm that the convergence to the payoff dominant equilibrium is common both unconditionally and conditional on convergence to some convention.

Table 5: Play in the Stag Hunt Games

Cohort ID	Rounds 1 - 15			Rounds 61 - 75		
	Payoff Dominant	Risk Dominant	Secure	Payoff Dominant	Risk Dominant	Secure
101	0.90	0.53	0.10	1.00	0.47	0.00
102	0.79	0.63	0.21	0.73	0.69	0.28
501	0.86	0.61	0.14	0.89	0.57	0.11
502	0.85	0.62	0.15	0.95	0.52	0.05
601	0.93	0.51	0.08	1.00	0.47	0.00
602	0.88	0.57	0.12	1.00	0.47	0.00
603	0.75	0.70	0.25	0.79	0.66	0.21
901	0.92	0.52	0.08	1.00	0.47	0.00
902	0.77	0.63	0.22	0.70	0.70	0.30
1301	0.88	0.58	0.12	0.97	0.50	0.03
1302	0.78	0.65	0.22	0.82	0.65	0.18
1303	0.92	0.52	0.08	0.76	0.59	0.24

---

[16] In terms of the model, we can devise similarity functions that yield H4. The similarity function determines how payoffs from past actions are weighted when determining weights over currently available actions. For example, suppose payoffs from Action A in the stag hunt games reinforce only Action 3 in the order statistic games since both are maximin choices. Likewise suppose payoffs from Action B in the stag hunt games only reinforce Action 7 in the order statistic games since both are played in the payoff dominant equilibrium. If play in the stag hunt games converges to the payoff dominant equilibrium, it follows that play of the payoff dominant equilibrium is reinforced in the order statistic games and prior experience with the stag hunt games shifts play toward the payoff dominant equilibrium in the ROS games. The preceding yields H4 as a prediction but relies entirely on the assumed form of the similarity function, begging the main empirical question: Do subjects use general principals such as payoff dominance to identify choices in different games (i.e. stag hunt and order statistic games) as being similar?

Table 5 summarizes play from the stag hunt games. The data is broken down by cohort. Recall that a total of seventy-five rounds of the stag hunt games were played. To allow comparison of the beginning and end points, data is shown for the first fifteen games (Rounds 1 – 15) and the final fifteen games (Rounds 61 – 75). We display the proportion of subjects whose play is consistent with the payoff dominant, risk dominant, and secure equilibrium. Because the risk dominant equilibrium may overlap either the payoff dominant or secure equilibrium depending on the specific game being played, the proportion of payoff and risk dominant equilibrium generally sums to more than 1.

Define a cohort as having converged to a convention if 80% or more of play is consistent with that convention for the final fifteen rounds. Cases where cohorts have converged to a convention in the last fifteen rounds are highlighted in yellow. Eight of twelve groups have converged to the payoff dominant equilibrium and a ninth group just misses. No groups converge to the risk dominant or secure equilibrium. In eleven groups, the payoff dominant equilibrium is played more frequently than the risk dominant equilibrium over the final fifteen rounds (and there is a tie for the twelfth group). Play of the payoff dominant equilibrium increases with experience for all of the groups that converged to playing the payoff dominant equilibrium. In contrast, play of the risk dominant equilibrium is as likely to decrease (7 of 12 groups) as increase with experience.

Conclusion 1: When cohorts playing the stag hunt games converge to a convention, they converge to play of the payoff dominant equilibrium. The majority of groups converge to this convention. The data is consistent with H3 and replicates the main result of Rankin et al.

B. ROS Games: The main results of the paper can be seen in Figure 1. Data is shown for the first two blocks (14 rounds) of play in the ROS game. Recall that all subjects played the same games in the same order for these fourteen rounds. The data is broken down by treatment and by the order statistic (OS) being used. The vertical axis shows the average choice by subjects. Recall that possible choices range between 1 and 7. The differences between treatments are similar in the two blocks, albeit more extreme in Rounds 8 – 14.

Three clear patterns can be seen in the data. (1) In all three treatments, the average choice is an increasing function of the order statistic (OS). This is consistent with H1. (2) Comparing play in the Fixed and ROS treatments, the slope of the relationship between the value of OS and average choices is flatter in the ROS treatment. This is consistent with H2. Average choices are generally higher in the Fixed treatment than the ROS treatment. (3) Average choices are always higher in the SHROS treatment than the ROS treatment. This is consistent with H4. This shift is stronger for low values of OS.

Figure 1: Choice by Order Statistic (OS), Rounds 1 - 14
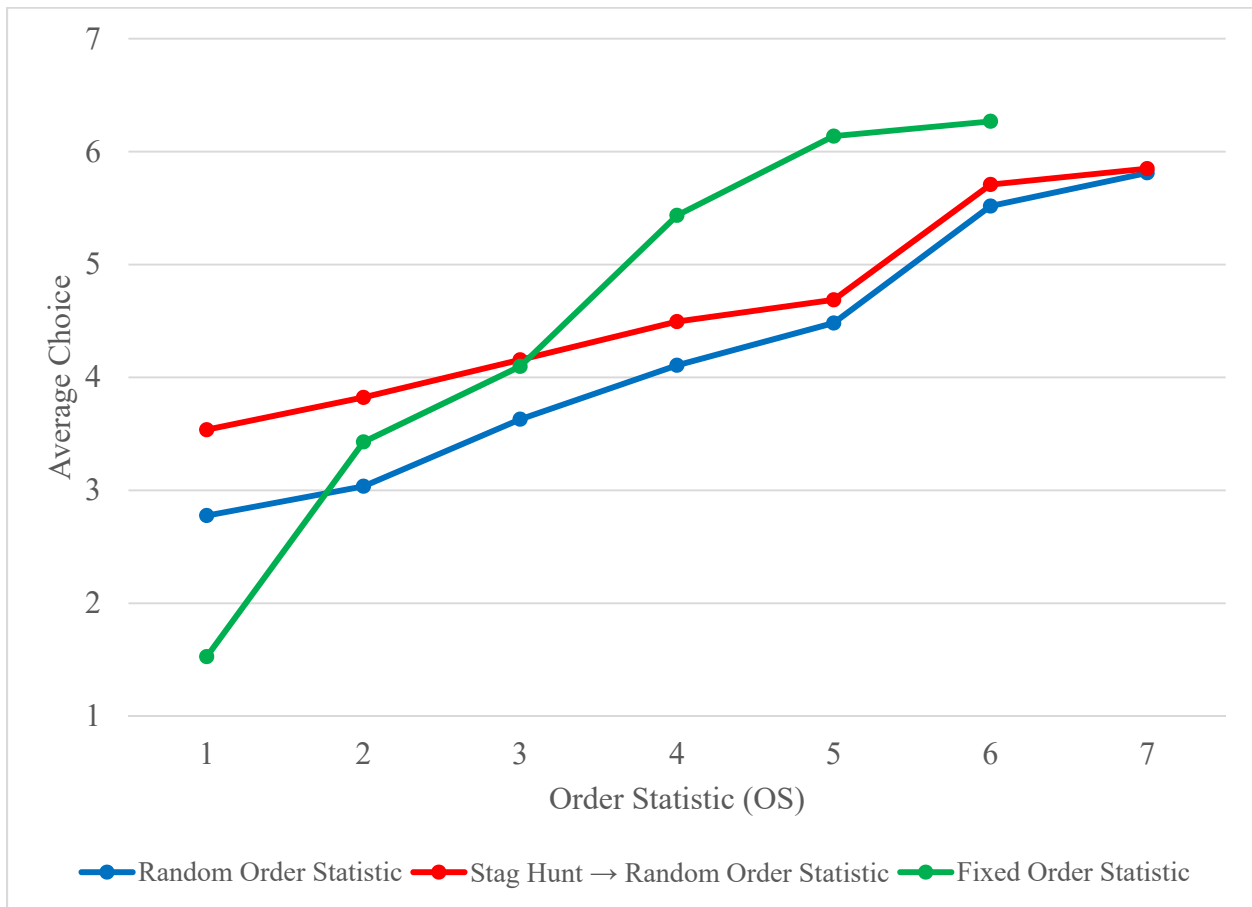
The regressions shown in Table 6 establish the statistical significance of these three patterns. The data is taken from the first fourteen rounds of play with the ROS games. Each observation is a single choice, ranging between 1 and 7, by an individual subject. Given that the choices are categorical in nature and have a natural ordering, an ordered probit model is used. Standard errors are reported in parentheses. Observations from the same cohort are correlated, so standard errors are corrected for clustering at the cohort level. Note that this implies that all observations from the same individual subject are part of the same cluster and are *not* treated as independent.[17]

In Model 1, the independent variables are dummies for the SHROS and Fixed treatments (the ROS treatment serves as the base), the order statistic (OS), a dummy for the second block (Rounds 8 – 14), and an interaction term between the order statistic and the dummy for the second block. Both treatment dummies are positive and statistically significant. The significant estimate for the SHROS dummy is consistent with H4. The estimate for OS is positive and significant. Notice that the dummy for the second block is negative while the interaction term is

---

[17] Both regressions are one observation short. In the ROS treatment session that had to be cut short, there was one subject who never entered her final choice.

positive. Both parameters are statistically significant. The slope of the relationship between subject choices and OS becomes steeper with experience.

Table 6: Ordered Probits, Treatment Effects in the ROS Games

Independent Variable	Model 1	Model 2
Stag hunt → ROS	0.334**	0.740***
	(0.143)	(0.272)
Fixed Order Statistic	0.449***	-0.915***
	(0.165)	(0.334)
Order Statistic (OS)	0.351***	0.273***
	(0.035)	(0.033)
Block 2 (Rounds 8 – 14)	-0.875***	-0.934***
	(0.098)	(0.103)
Order Statistic x Block 2	0.280***	0.302***
	(0.026)	(0.029)
Order Statistic x Stag hunt → ROS		-0.102**
		(0.044)
Order Statistic x Fixed		0.395***
		(0.085)
Log-Likelihood	-7379.59	-7088.35
Number of Observations (Cohorts)	4,927 (44)	4,927 (44)

Notes: Dependent variable is row choice by subject. Standard errors are corrected for clustering at the cohort level. Three (***), two (**), and one (*) stars denote statistical significance at the 1%, 5%, and 10% levels respectively.

Model 2 adds interaction terms between the two treatment dummies and the value of OS. The dummy for the Fixed treatment is now negative while the interaction between this dummy and OS is positive. Both parameters are statistically significant. Together these estimates indicate that the slope of the relationship between subject choices and OS is steeper in the Fixed treatment than the ROS treatment as predicted by H2. The dummy for the SHROS treatment is positive while the interaction term between this and OS is negative. Both parameters are statistically significant. It follows that the slope of the relationship between subject choices and OS is flatter in the SHROS treatment than the ROS treatment. There is only a difference between the ROS and SHROS treatments for the lower values of OS. For the maximum value of the order statistic, OS = 7, the estimated difference is a miniscule 0.028. This difference is not statistically significant. Finally, it is worth noting that the positive relationship between OS and subjects' choices is statistically significant in all three treatments.[18] This is consistent with H1.

When we designed the experiments, we were concerned that the treatment effects of the SHROS treatment could be due to income effects as subjects in this treatment start the ROS

---

[18] To show this formally, we reran Model 2 with the SHROS treatment as the base. The parameter estimate for OS is 0.171 with a standard error of 0.036, significant at the 1% level. See Appendix A for full output (Table A.1, Model 2).

treatment with a positive cash balance due to earnings from the stag hunt games. To control for this possibility we added a random shock to subjects' cash balances prior to beginning play of the ROS games. We then reran Model 1 from Table 6 with two changes. First, we used a tobit rather than an ordered probit to allow us to use an instrumental variable. Second, we added the starting cash balance as an independent variable. Since this is endogenous in the SHROS treatment, we use the income shock as an instrument for the starting cash balance. This has no impact on our conclusions. The estimate for the starting cash balance is small and does not approach statistical significance. The two treatment dummies both remain positive and are both statistically significant.[19]

H1 predicted that play would converge to an equilibrium in the Fixed treatment. The data supports this prediction. A good measure of convergence is the absolute difference between subjects' chosen numbers and the exclusive order statistic (EOS) for their group. This statistic equals zero if play converges to an equilibrium. For the final block (Rounds 37 – 42) of the Fixed treatment, the average absolution difference between subjects' chosen numbers and EOS equals .05. This indicates a high degree of convergence. As a point of comparison, for the final block of the ROS treatment the equivalent figure is .49. Convergence is stronger in the Fixed treatment than the ROS treatment.

Conclusion 2: (a) There is strong convergence to equilibrium in the Fixed treatment. The average number chosen is an increasing function of the order statistic (OS) in all three treatments. The data is consistent with H1. (b) The slope of the relationship between OS and choices is steeper (flatter) in the Fixed (SHROS) treatment than in the ROS treatment. The relationship between the ROS and Fixed treatments is consistent with H2. (c) Compared with the ROS treatment, play in the SHROS treatment is shifted toward higher choices. This finding is consistent with H4. The difference between the ROS and SHROS treatments is due to shifts at low values of OS.

The intuition underlying H4 is that subjects in the SHROS treatment learn to play the payoff dominant equilibrium in the stag hunt game and transfer this convention to the ROS games. This implies that there should be a positive relationship between playing the payoff dominant equilibrium in the stag hunt games and choices in the ROS games. The data is consistent with this prediction. For each cohort in the SHROS treatment, we calculate the proportion of choices consistent with the payoff dominant equilibrium in the stag hunt games and the average choice in the ROS games. Figure 2 shows a scatter plot of these two variables. The positive relationship between the two variables can be seen easily.

---

[19] See Appendix A for full output (Table A.1, Model 1).

Figure 2: The Relationship between Play in Stag Hunt and ROS Games
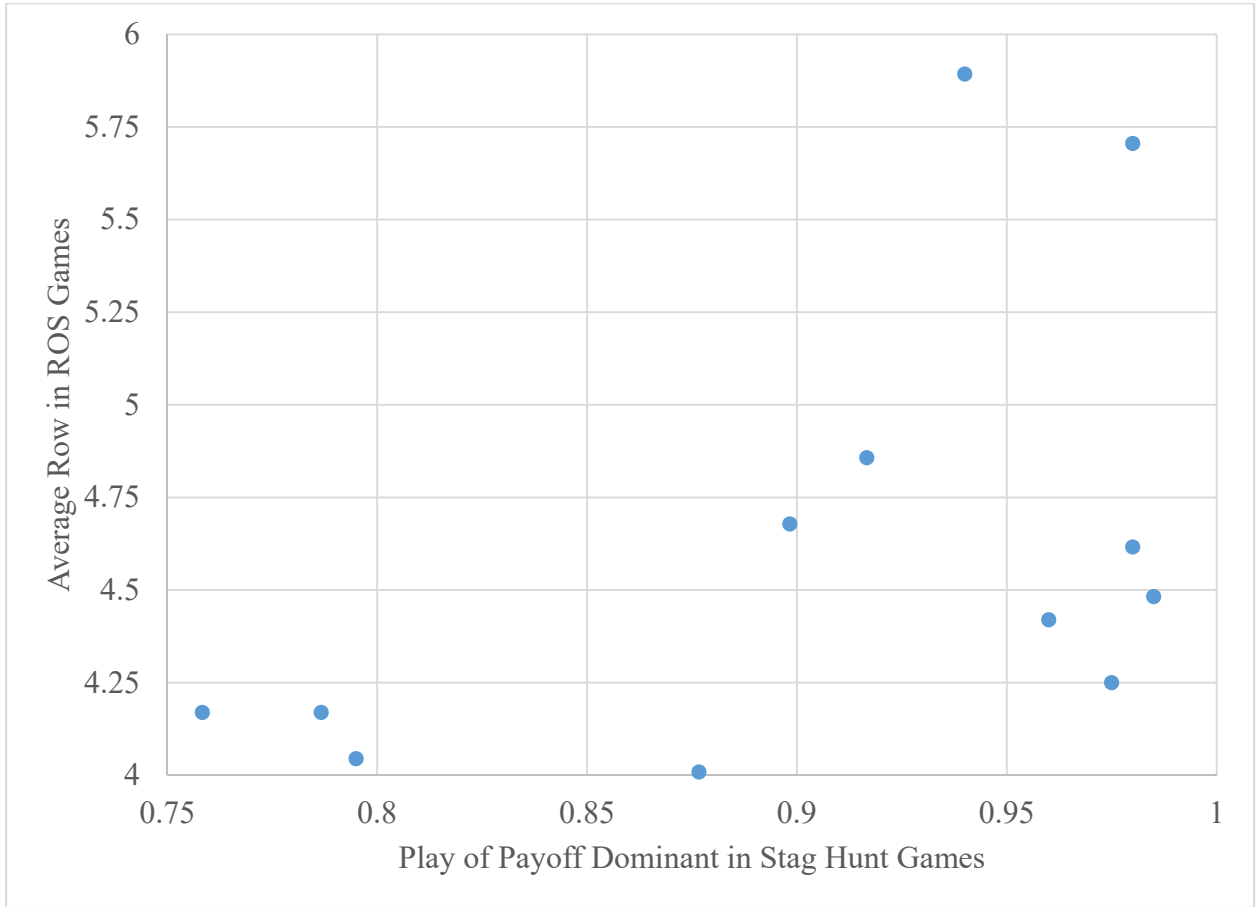
The regressions shown on Table 7 add more detail on this point. The dataset includes all of the ROS games in the SHROS treatment. As previously, these are ordered probits. Standard errors are reported in parentheses and are corrected for clustering at the cohort level. The dependent variable is an individual's choice in a single round of the ROS game. Both regressions include the value of OS and a dummy for the second block (Rounds 8 – 14) as independent variables.

The primary independent variable in Model 1 is the cohort's proportion of choices consistent with the payoff dominant equilibrium in the stag hunt games. In line with Figure 2, the parameter estimate is positive and statistically significant.

Model 1 leaves important questions about causality unanswered. The intuition underlying H4 is based on learning. Subjects learn to play the payoff dominant equilibrium in the stag hunt games and then apply this principle in the ROS games. But, looking at Table 5, play does not change much with experience in the stag hunt games. Perhaps the lagged variable in Model 1 captures an individual characteristic (e.g., an individual fixed effect that inclines subjects to choose the payoff dominant equilibrium) rather than an effect due to learning. Model 2 addresses this issue by adding two new independent variables, a dummy for whether the

subject's play in the first stag hunt game was consistent with the payoff dominant equilibrium and the number of *other* subjects in the cohort whose play in the first stag hunt game was consistent with the payoff dominant equilibrium. These variables cannot reflect any effects due to learning since only data from Round 1 is included, prior to any interactions or feedback. Instead, these variables should capture individual effects. In Round 1, 51% of the choices were consistent with the payoff dominant equilibrium. There is no shortage of variation.

Table 7: Ordered Probits, Effect of Play with Stag Hunt Games

Independent Variable	Model 1	Model 2
% Payoff Dominant Equilibrium (Group)	2.791***	2.789***
	(0.957)	(0.994)
Played Payoff Dominant Equilibrium		-0.018
Own Play, Round 1		(0.104)
Played Payoff Dominant Equilibrium		0.032
Other Group Members, Round 1		(0.567)
Order Statistic (OS)	0.308***	0.308***
	(0.044)	(0.044)
Block 2 (Rounds 8 – 14)	0.142	0.142
	(0.093)	(0.093)
Log-Likelihood	-2057.06	-2056.99
Number of Observations (Cohorts)	1,344 (12)	1,344 (12)

Notes: Dependent variable is row choice by subject. Standard errors are corrected for clustering at the cohort level. Three (***), two (**), and one (*) stars denote statistical significance at the 1%, 5%, and 10% levels respectively.

The estimates for the cohort's proportion of choices consistent with the payoff dominant equilibrium in the stag hunt games are virtually identical in Models 1 and 2. The parameter estimates for the two new variables in Model 2 are small and nowhere close to statistical significance. The inclusion of variables that should do a superior job of capturing individual effects adds no explanatory power. This suggests that the effect of previous play in the stag hunt games reflects learning rather than individual effects.[20]

---

[20] The number of clusters in the dataset for these regressions is low, only 12, biasing us in favor of finding statistical significance (a common rule of thumb is that you should have at least 20 clusters). We have experimented with the "wild bootstrap method" developed by Cameron, Gelbach, and Miller (2008) to generate correct standard errors with a low number of clusters. This isn't perfect, since we have to use a linear probability model, but gives a sense of whether our result is driven by the low number of clusters. The statistical significance of the parameter estimate for "% Payoff Dominant Equilibrium (Group)" is lower (p = .06) as expected, but our qualitative conclusions are unchanged. (Oddly, the change in p-values is more due to use of a linear probability model than the method used to generate standard errors.) Play consistent with the payoff dominant equilibrium in the stag hunt games has significant explanatory power for play in the ROS games, and adding the two first Round 1 variables has a negligible effect.

Conclusion 3: Subject in cohorts that converged to play of the payoff dominant equilibrium in the stag hunt games choose higher actions in the ROS games. This positive correlation cannot be attributed to uncontrolled individual effects.

Table 8: Ordered Probits, Effect of Lagged OS and EOS

Independent Variable	Model 1	Model 2
Order Statistic (OS)	0.428***	0.432***
	(0.041)	(0.040)
Lagged Order Statistic	0.071***	
	(0.007)	
Lagged Exclusive Order Statistic (EOS)		0.115***
		(0.023)
Lagged Exclusive Order Statistic (EOS)		-0.002
x |OS – Lagged OS|		(0.003)
SHROS Treatment	0.345**	0.309**
	(0.146)	(0.127)
Block 2 (Rounds 8 – 14)	0.213***	0.224***
	(0.044)	(0.046)
Log-Likelihood	-4134.80	-4104.58
Number of Observations (Cohorts)	2,703 (26)	2,703 (26)

Notes: Dependent variable is row choice by subject. Standard errors are corrected for clustering at the cohort level. Three (***), two (**), and one (*) stars denote statistical significance at the 1%, 5%, and 10% levels respectively.

The intuition underlying H2 was simple. Suppose I draw on my experience with the other games when playing in the ROS treatment. If OS is low, most of my experience will come from games with higher order statistics. Since higher choices tend to do better in games with higher order statistics, this biases upwards my estimate of how effective higher choices are likely to be and hence makes me more likely to choose a high choice. Similar logic applies when OS is high, biasing choices downward. This led to a prediction that the relationship between choices and OS would be flatter in the ROS treatment than the Fixed treatment.

Not only does the data support this prediction, the regressions in Table 8 show that the data provides evidence in line with the underlying explanation. The dataset includes the first fourteen ROS games of the ROS and SHROS treatments. The first round is dropped to allow use of lagged variables. Once again these are ordered probits, standard errors are reported in parentheses, and standard errors are corrected for clustering at the cohort level. The dependent variable is an individual's choice in a single round of the ROS game.

Model 1 has three independent variables: the order statistic (OS), the lagged value of OS, and a dummy for the second block (Rounds 8 – 14). The variable of interest is the lagged value of OS. If subjects ignore information from similar games or accurately correct for the differences between games, the lagged order statistic should not affect current behavior. Instead, the parameter estimate is positive and strongly significant. This implies that subjects who have seen

a relatively high (low) value of OS in the previous round will be biased upwards (downwards) in their current choice. This is consistent with the intuition underlying H2.[21]

Model 2 looks directly at subjects' responses to feedback. As independent variables we now include the lagged value of the *exclusive* order statistic (EOS) and an interaction between the lagged EOS and the *absolute value* of the difference between the current and lagged order statistics (OS). We expect the response to the lagged EOS to be positive. If subjects put more weight on experience from games with greater similarity to the current game, we also expect the effect to be smaller when the difference between the current and lagged games is increased. In other words, the first parameter should be positive and the second negative. The predicted signs are correct but the parameter estimate for the interaction term is tiny and nowhere close to statistical significance. Subjects do not appear to account for similarity or lack thereof between different order statistic games when responding to recent experience. This helps explain the results of Model 1.

Conclusion 4: Subject are sensitive to the lagged value of the order statistic, consistent with the intuition underlying H2. Subjects' responses to feedback are not sensitive to the similarity between the current and previous order statistic game.

Our analysis thus far has only looked at the first fourteen ROS games, the time period that overlapped in all three treatments. The ROS and Fixed treatments included 42 rounds of play. Figure 3 examines what happens as subjects become very experienced. Data is shown from the first block of ROS games (Rounds 1 – 7) for all three treatments and the final block of the ROS and Fixed treatments (Rounds 35 – 42). Recall that the order of games is the same for the first and last blocks of the ROS treatment.

Two important points can be taken away from Figure 3. (1) With experience in the ROS treatment (blue lines), the relationship between OS and subject choices becomes steeper. There is no generic shift upwards, with the average choice across all values of OS decreasing slightly between the first and final blocks (4.22 vs. 4.12). Instead, choices become lower for low values of OS and higher for high values. The difference between the first and last blocks of the ROS treatment do not mirror the differences between the first blocks of the ROS and SHROS treatments. In the latter case, choices are higher in the SHROS treatment for low values of OS and virtually identical for high values. Previous experience with the stag hunt game has a different effect on choices in the ROS game than previous experience with the ROS game. We conjecture that in the former case subjects are transferring a general principle while in the latter they are learning actions rather than principles. (2) The effect of experience differs for the Fixed and ROS treatments. In the Fixed treatment there is a clear shift upwards over time. The average choice across all values of OS rises substantially between the first and final block (4.27

---

[21] Model 1 pools data from the ROS and SHROS treatments. If the model estimates separate responses to the lagged order statistic, the parameter estimates are virtually identical for the two treatments (0.071 and 0.072 respectively).

vs 5.16) with play converging to the payoff dominant equilibrium for OS = 4, 5, and 6.[22] Even for lower values of OS the general trend is upwards, albeit weakly. The difference between the ROS and Fixed treatments becomes more extreme with experience.

Figure 3: Comparison of Early and Late Play in ROS Games
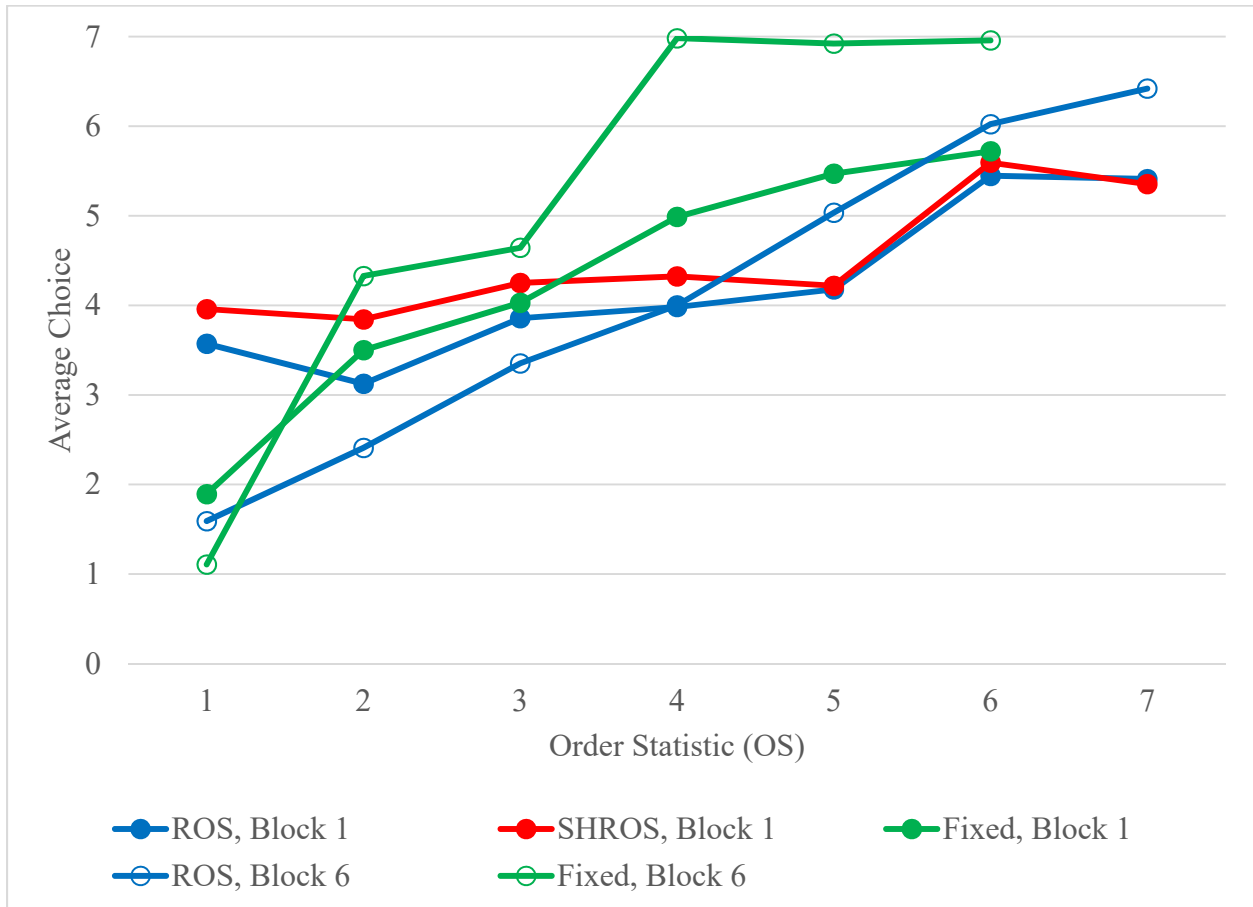
Based on data from the first two blocks (Rounds 1 – 14), Conclusions 2a and 2b state that the data support H1 and H2. The same holds if we consider the data from the final block (rounds 37 – 42) displayed in Figure 3. Consistent with H1, there is a strong positive relationship between average choices and the order statistic in both the Fixed and ROS treatments. The slope of the relationship between OS and subject choices remains steeper in the Fixed treatment than the ROS treatment, subject to the caveat that choices hit a ceiling for the higher values of OS in the Fixed treatment. The data from the final block support H2.

---

[22] Running a Wilcoxon matched pairs signed rank test, the difference between the first and last blocks in the Fixed treatments is significant at the 5% level ($z = 2.42$; $p = 0.02$). Each observation is the average over a cohort for the seven round block. Oddly, the decrease in the ROS treatment is also significant ($z = -2.31$; $p = 0.02$).

4. Conclusion: We study the ability of subjects to transfer precedents learned in one coordination game to other related coordination games. Our subjects exhibit strong transfer between related games. On the positive side, the strong transfer observed in the SHROS treatment indicates that transfer can take place even when games are *not* structurally similar. The ability of subjects to learn a general principle and apply it appropriately in another game is surprisingly strong. That said, comparing play in the ROS and Fixed treatments indicates that transfer of precedents does not always lead to more efficient outcomes in coordination games.

The latter finding contrasts in particular with the results of Cason *et al.* They observe positive transfer from median to minimum games, where transfer is efficiency enhancing, but no transfer from minimum to median games (which presumably would harm efficiency). It remains to be explained why transfer is biased in favor of efficiency in their experiments but not in ours. We conjecture that this is due to the structure of the experiments. In the ROS treatment, subjects see an ever-changing sequence of games. There is never an obvious break as in the SHROS treatment or the Cason *et al.* experiments where the game is changed after a long run of similar games. The strong break may stimulate subjects to think about general principles rather than using a more mechanical rule.

One of the more surprising features of our data is the insensitivity of subjects in the ROS game to differing values of OS when drawing on past experience. The choices induced by OS = 2 and OS = 6, for instance, are quite different, but subjects treat past experience from these two games almost identically when deciding on current actions. This is consistent with the primary assumption underlying analogy-based expectations equilibrium (Jehiel, 2005) – subjects will treat broad classes of games as equivalent when making choices. A high priority for future research is work that better identifies how subjects perceive similarity between games.

All of the games in our experiments are coordination games with a payoff dominant equilibrium and a secure equilibrium. We observe positive transfer across games that are less similar than is typical in experiments studying transfer, but it is reasonable to ask if we could have gone farther and observed positive transfer across games that were even less related. Along similar lines, our subjects remain in the same cohort throughout. They know that their opponents have received similar experience to their own and are likely to have absorbed the same general principles. It is reasonable to ask whether we would observe transfer in the SHROS treatment if the cohorts were not fixed.

Much of our discussion has revolved around the ability of subjects to learn and apply general principles. That said, it is not clear precisely what principle is being learned in our experiments. Are subjects learning to play a payoff dominant equilibrium, or are they learning something more akin to a norm such as choosing an efficient outcome regardless of whether or not this is consistent with an equilibrium.[23] The answer to this question is linked to the process underlying

---

[23] Peysakhovich and Rand (2016) provide an example of cross-game transfer between repeated prisoners' dilemma games and various one shot games. They show that the transfer is driven by a norm of pro-social behavior rather any notion of equilibrium selection, as transfer occurs for non-strategic settings such as the dictator game.

transfer. Is it more deliberative in line with equilibrium selection, or emotive in line with acquisition of a norm? This is an important issue for future research.

It says something negative about the state of the literature that this paper could add significantly to the literature a good fifteen years after it was conceived. Learning in games is a central issue for our understanding of how equilibrium emerges, and our understanding of learning is incomplete without an understanding of transfer. There is an urgent need to answer questions such as the circumstances under which transfer will occur, how similarity is mapped between games, and precisely what is being transferred between games. John Van Huyck was keenly interested in these questions and made pioneering contributions to answering them. Hopefully this paper will play a small role in stimulating further work that builds on his legacy.

Bibliography

Alfieri, Louis, Timothy J. Nokes-Malach, and Christian D. Schunn (2013), "Learning Through Case Comparisons: A Meta-Analytic Review," *Educational Psychologist*, 48, 2, 87-113.

Bednar, Jenna, Yan Chen, Tracy Xiao Liu, and Scott Page (2012), "Behavioral spillovers and cognitive load in multiple games: An experimental study," *Games and Economic Behavior*, 74, 1, 12–31.

Brandts, Jordi and David J. Cooper (2006), "A Change Would Do You Good: An Experimental Study of How to Overcome Coordination Failure in Organzations," *American Economic Review*, 96, 3, 669-693.

Cason, Timothy N, Anya C. Savikhin, and Roman M. Sheremeta (2012), "Behavioral spillovers in coordination games," *European Economic Review*, 56, 233–245.

Cason, Timothy N. and Lata Gangadharan (2013), "Cooperation Spillovers and Price Competition in Experimental Markets," *Economic Inquiry*, 51, 3, 1715–1730.

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

Cameron, Colin, Jonah Gelbach, and Douglas Miller (2008), "Bootstrap-based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 90, 3, 414–427.

Cooper, David J. and Kagel, John H. (2005), "Are Two Heads Better Than One? Team versus Individual Play in Signaling Games," *American Economic Review*, 95, 3, 477-509.

Cooper, David J. and Kagel, John H. (2008), "Learning and transfer in signaling games," *Economic Theory*, 34, 3, 415-439.

Cooper, David J. and Kagel, John H. (2009), "The Role of Context and Team Play in Cross-Game Learning," *Journal of the European Economic Association*, 7, 5, 1101-1139.

Devetag, Giovanna (2005), "Precedent transfer in coordination games: An experiment," *Economics Letters*, 89, 227–232.

Duffy, John and Dietmar Fehr (2015), "Equilibrium Selection in Similar Repeated Games: Experimental Evidence on the Role of Precedents," working paper, University of California, Irvine.

Erev, Ido and Ernan Haruvy (forthcoming), "Learning and the Economics of Small Decisions," in *The Handbook of Experimental Economics, Volume 2,* (John Kagel and Alvin Roth, Eds.), Princeton, NJ: Princeton University Press.

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

Fudenberg, Drew, and David K. Levine (2009), "Learning and equilibrium," *Annual Review of Economics*, 1, 385–420.

Gick, Mary L. and Keith J. Holyoak (1983), "Schema induction and analogical transfer," *Cognitive Psychology*, 15, 1, 1–38.

Greiner, Ben (2015), "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE," *Journal of the Economic Science Association*, 1, 1, 114-125.

Grimm, Veronika and Friederike Mengel (2012), "An experiment on learning in a multiple game environment," *Journal of Economic Theory*, 147, 6, 2220–2259.

Haruvy, Ernan and Dale O. Stahl (2012), "Between-game rule learning in dissimilar symmetric normal-form games," *Games and Economic Behavior*, 74, 208–221.

Ho, Teck-Hua, Colin Camerer, and Keith Weigelt (1998), "Iterated Dominance and Iterated Best Response in Experimental "p-Beauty Contests"," *American Economic Review*, 88, 4, 947-969.

Jehiel, Philippe (2005), "Analogy-based expectation equilibrium," *Journal of Economic Theory*, 123, 81–104.

Kreps, David M. (1990), *Game Theory and Economic Modelling*. Oxford, UK: Clarendon Press.

Nash, John (1950), "Equilibrium Points in n-Person Games," *Proceedings of the National Academy of Sciences of the United States*, 36, 1, 48-49.

Ochs, Jack (1995), "Coordination Problems," in *The Handbook of Experimental Economics,* (John Kagel and Alvin Roth, Eds.), Princeton, NJ: Princeton University Press.

Peysakhovich, Alexander and David Rand (2016), "Creating Norms of Cooperation and Defection in the Laboratory," *Management Science*, 62, 3, 631–647.

Rankin, Frederick W., John B. Van Huyck, and Raymond Battalio (2000), "Strategic Similarity and Emergent Conventions: Evidence from Similar Stag Hunt Games," *Games and Economic Behavior*, 32, 315-337.

Rick, Scott and Roberto A. Weber (2012), "Meaningful Learning and Transfer of Learning in Games Played Repeatedly Without Feedback," *Games and Economic Behavior*, 68, 2, 716-730.

Roth, Alvin E and Ido Erev (1995), "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior,* 8, 164-212.

Savikhin, Anya C. and Roman M. Sheremeta (2013), "Simultaneous Decision-Making in Competitive and Cooperative Environments," *Economic Inquiry*, 51, 2, 1311–1323.

Van Huyck, John B., Raymond Battalio and Richard Beil (1990), "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review*, 80, 234-248.

Van Huyck, John B., Raymond Battalio and Richard Beil (1991), "Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games," *Quarterly Journal of Economics,* 106, 3, 885-910.

Weber, Roberto (2006), "Managing growth to achieve efficient coordination in large groups," *American Economic Review,* 96, 1, 114-126.

Table A.1: Additional Regressions on Treatment Effects in the ROS Games

Independent Variable	Model 1	Model 2
Stag hunt → ROS	.876**	
	(.388)	
Fixed Order Statistic	.794***	-1.655***
	(.288)	(.374)
Random Order Statistic		-0.740***
		(0.272)
Order Statistic (OS)	.587***	0.171***
	(.074)	(0.036)
Block 2 (Rounds 8 – 14)	-1.548***	-0.934***
	(.169)	(0.103)
Order Statistic x Block 2	.512***	0.302***
	(.057)	(0.029)
Order Statistic x ROS		0.102**
		(0.044)
Order Statistic x Fixed		0.497***
		(0.083)
Starting Balance / 100	-.012	
	(.010)	
Log-Likelihood	-38,542.81	-7088.35
Number of Observations (Cohorts)	4,927 (44)	4,927 (44)

Notes: Dependent variable is row choice by subject. Model 1 is a tobit with the bonus used as an instrument for the starting balance. Model 2 is an ordered probit. Standard errors are corrected for clustering at the cohort level. Three (***), two (**), and one (*) stars denote statistical significance at the 1%, 5%, and 10% levels respectively.