

UC Santa Barbara

Econ 196 Honors Thesis

Title

The Cost of Indecision in Coordination Games

Permalink

<https://escholarship.org/uc/item/8454n8d2>

Author

Bjorke, Isaac

Publication Date

2016-10-24

The Cost of Indecision in Coordination Games

Isaac Bjorke

advised by Emanuel Vespa

Abstract

I propose a model for a coordination game that examines the potential effects of costly indecision between actions in a repeated stag hunt game. With small enough costs (which I call frictions) the conventional game theoretic predictions are unchanged, which are well-known to be unreliable for coordination games such as the stag hunt. Depending on certain assumptions about behavioral strategies, conflicting hypotheses for the effect of these frictions on outcomes in this game can arise. To gauge these hypotheses I take this game to a laboratory setting to compare a repeated stag hunt with frictions against a standard repeated stag hunt. Comparing short-term and long-term behavior between these games, I find little difference in the behavior between these games after multiple rounds, but significant strategic differences in the first period of these games, suggesting a difference in the development of prior beliefs in these games.

I. Introduction

The stag hunt game can be interpreted as the crux of many social conflicts and coordination issues, from group-effort problems to global political games. This game (along with the large family of similar coordination games) contains multiple pure-strategy Nash Equilibria, which naturally provides a challenging environment to make predictions about behavioral outcomes. The stag hunt model captures the basic dilemma of multiple and Pareto-ranked Nash Equilibria in coordination struggles, in a minimalistic setting. It may be thought that these simple coordination games would converge to the Pareto dominant equilibria when played repeatedly. In response to this argument, it is important to note the differences in the games shown in Figure 1, where $A > a \geq b > 0$. The first game, which I call game A is a plausible game in which the Pareto dominant outcome would arise. The two equilibria (S, S) and (H, H) are symmetric in every way except their payoffs. The second game, which I will now call game G and is the traditional stag hunt game, is different than A only in the off-equilibrium payoffs.

	<i>S</i>	<i>H</i>
<i>S</i>	<i>A, A</i>	0, 0
<i>H</i>	0, 0	<i>b, b</i>

Game *A*

	<i>S</i>	<i>H</i>
<i>S</i>	<i>A, A</i>	0, <i>a</i>
<i>H</i>	<i>a, 0</i>	<i>b, b</i>

Game *G*

Figure 1

We can clearly see that the significant difference between these games comes from the safeness of playing *H* in Game *G*. Playing *H* in Game *G* guarantees a positive payoff, despite the opposing player's action. Game *A* does not have this safeness in either action, and thus it is logical to believe that due to the symmetry in riskiness between the actions, that players would simply go for the action whose equilibrium payoff is highest. This would also be safest for the opposing player, and thus there are no higher levels of strategic uncertainty in playing *S* in Game *A*. Comparatively, the strategic uncertainty in *G* is relatively high compared to *A*. The riskiness of *S* makes *H* "risk dominant", and thus players have more reason to be uncertain about the actions of the other players. Being uncertain about the opposing player's actions increases the incentives to play *H*, and thus game *G* illustrates the anti-cooperative and uncoordinated behavior caused by the risk involved with opting toward pareto dominance. The stag hunt game brings this risk into perspective in a minimally reduced setting, and reveals the fundamental force of risk in non coordinative outcomes.

Predicting outcomes of the finitely repeated version of game *G* is just as difficult as the one-shot game. Since both *S, S* and *H, H* are rational Nash Equilibria in each round, any strategy profile that designates *S, S* or *H, H* in any round can be categorized as a Subgame Perfect Equilibria. While strategic uncertainty is still prominent, there are signaling factors involved when the game is repeated that can be used to justify the elimination of many of the SPE solutions. If an agent considers the action of their opponent in the last round to be a strong signal of their action in the upcoming round, then the best response is to match their action. Considering the history of the game as a signal of the next round can lead to the proposition that the two SPE that are most likely points of convergence are those with either strictly *S, S* or *H, H* in all rounds. This elimination still leaves two possible outcomes, and the strength of our theoretical predictions for the finitely repeated game match that of the one-shot game, each having two potential outcomes.

Interpreting the stag hunt game as the crux of coordination conflict suggests that coordination is difficult to achieve primarily due to the risk associated with pareto optimal actions. While this model uses this risk as an explanation for many coordination failures, it shrinks coordination into an unrealistically simple environment. This game assumes that players can easily choose between S and H without any difficulty, and thus ignores many possible challenges of coordination that could result from the effort involved with cooperation. It is noteworthy that the stag hunt does not incorporate the potential costs of switching between actions. In a repeated stag hunt, agents are able to adjust their temporary strategies automatically. Under the classic story of hunters in a forest, this simple choice of adjustment in the repeated repeated stag hunt assumes that hunters are able to costlessly and instantaneously travel from the stag-dwelling areas of the forest to the hare-dwelling areas. Transitioning between states in actual coordination applications can require agents to undergo certain costs associated with their change of action: hunters lose time by travelling through the forest, coordinating meetings brings forth transportation costs, political revolvers need to purchase weapons and change wardrobes, and any action change may require the purchasing of different tools or goods. I suggest a finitely repeated stag hunt model that incorporates these potential costs when altering actions between periods without changing subgame perfect equilibria. This model will act as a treatment of the stag hunt game whose behavioral differences from the original stag hunt can be completely accredited to the intertemporal costs, which I now call *friction*.

The purpose of this research is to question how/if friction between action can also be considered an explicit force in coordination conflicts, or if the effects of friction can provide insight into the behavioral strategies involved with equilibria selection in coordination games. While the stag hunt isolates the conflict between risk dominance and payoff dominance as a simplistic model of coordination difficulty, understanding the behavioral differences between the stag hunt with and without friction will bring considerable light onto other fundamental forces behind the dynamics of games with multiple equilibria.

II. Literature Review

Coordination games have been used to model social conflict throughout modern literature on normal form games. Schelling (1980) establishes a large collection of social environments in which these games model the difficulties of cooperation. The existence of multiple plausible Nash Equilibria in the stag hunt exists due to the conflict between risk dominance and payoff dominance, unlike in the “battle of the sexes” game or pure coordination games which have conflict arise strictly from the existence of symmetric multiple equilibria. The conjecture that Pareto dominant behavior will naturally arise from repeated play can be traced back to Harsanyi and Selten (1988), and perhaps even earlier back to theories similar to the invisible hand. Of course as we look into the experimental literature we can easily reject this hypothesis.

Many coordination experiments have studied the fluctuations in behavioral equilibria for stag hunt games and games of similar structure. For one, Cooper et al (1990) show that both pure strategy equilibria as well as the mixed strategy equilibria can be observed in coordination games, rejecting the argument that Pareto dominant outcomes must be a focal point of coordination. While the results from their experiment are not direct stag hunt results, they shed some light on why agents fail to coordinate to the Pareto-efficient outcome aside from strict riskiness. They are able to show that in general, having outside options and complications, even when they are dominated strategies, can interfere with the selection process of agents. Thus, we can see that in coordination games agents often have difficulty reasoning their way towards equilibria. This may seem obvious and can be seen in many settings, such as the experiment analyzed by Nagel (1995), but has deeper importance when arguing against Pareto-dominant outcomes, in that it suggests that agents cannot necessarily reason their way towards pareto-dominance due to the nature of a best response function being just that: a best response to their opponent, which is not necessarily the best response to a coordinating rational agent.

The experimental results of Van Huyck et. al (1990) suggest that the fluctuations in cooperation are due to strategic uncertainty (the uncertainty of the type of each player's opponent). With their experimental analysis of minimum effort games they find that coordination failures are directly related to the number of players. The higher the number of players, the higher the strategic uncertainty, and thus more incentive to play the risk-dominant equilibrium action. As signaling was earlier mentioned to have large behavioral effects in repeated versions of the stag hunt, Crawford (1991) adopts an evolutionary model to interpret the results found in Van Huyck et. al, which he notes easily extends to the simple stag hunt. As recognized by Crawford, the stag hunt is simply the most reduced version of the minimum effort game, and he justifies three equilibria using evolutionary definitions of stability: S,S , H,H , and a mixed strategy dependent on the parameter values of A , a , and b . The mixed strategy is generally unstable due to the discrete time repetition of the game, for after each round a signal will influence each agent to opt away from an interior probability solution. Determining the selection between S,S and H,H depends on the current beliefs of each player and their learning from the signals of the previous rounds.

Equilibrium selection is hypothesized by Battalio et al (2001) to depend on the optimization premium, which formally models the intuition that high differences between Nash Equilibrium payoffs motivate payoff-dominant behavior. Crawford (1991) focuses his analysis on learning dynamics, but this experimental study makes significant steps toward our predictions of prior beliefs in the stag hunt game. The surprising result of this experiment is that they find significant effects on the magnitude of payoffs associated with playing a best response, despite holding the beliefs necessary to play S constant. Among the 3 treatments in their experiment, it is a dominant strategy to play S if and only if they believe that their opponent will play S with .8 probability. However, the magnitude of their payoffs from

playing best responses (which they call the ‘optimization premium’) directly affects the probability of coordination towards payoff dominance, as well as general coordination rates. From this collection of literature we can see that coordination results can be severely influenced by both beliefs as well as equilibria payoff steepness.

Aside from extending the general principles of coordination failure in the stag hunt to a continuous multiplayer setting, the minimum effort game plays another important role in the light of this particular paper by introducing the notion of cost and effort into coordination games. Modelling pareto dominant equilibria as the result of agents undergoing a cost is a framing that can be easily understood to model realistic coordination environments. Experiments by Goeree and Holt (2005) show that the particular size of the cost parameter can directly affect the cooperation rate in this game, which follows the optimization premium results found in Battalio et. al and extends the general cruciality of parameter magnitudes from stag hunt games to minimum effort games. The purpose of the experiment being proposed in this paper is to examine whether or not costs between actions can cause significant differences in coordinative and cooperative outcomes. There are a few subtle but significant differences between these frictions and the cost parameter in the minimum effort game. The costs of the minimum effort game are associated with the momentary choice of high effort levels, while the frictions in the proposed experiment only assign costs to intertemporal indecision. The cost of frictions will not apply to an agent who continually plays with high effort, but will apply to an agent who moves from playing high effort to low effort. The minimum effort game models the difficulty of achieving pareto-optimality through costs associated with the effort. The model in this paper questions how pareto-optimality and coordination rates can be affected by the costs of indecision and the impending punishment of miscoordination.

III. Model

The one-shot stag hunt game can be seen in Figure 1 (Game G), with $A > a \geq b > 0$. The two pure-strategy Nash Equilibria in the one-shot game are at (S, S) and (H, H) . Thus, in a repeated game with a finite n periods (which I now refer to as game G), any Subgame Perfect Nash Equilibrium can be categorized with each period resulting in (S, S) or (H, H) . Any combination of these outcomes is subgame perfect, and all subgame perfect equilibria maintain either (S, S) or (H, H) in every period.

Now let us define the game G' as the same repeated stag hunt with n periods, but with the adjustment that it costs some fixed $\varepsilon > 0$ for either player to play a different action than they had in the previous period. For example if (S, S) is played in period 1, and (S, H) in period 2, then row player earns a net profit of A while column earns $A + a - \varepsilon$ since column player switched their action from period 1 to 2. The set of subgame equilibria in G' is a subset of the equilibria in G , and the difference between the two sets of equilibria depends on the size of ε . For our experimental purposes, we wish to keep the set of subgame

equilibria identical in order to achieve a clear treatment effect from these cost of friction without altering the traditional equilibria concepts associated with the original stag hunt.

Proposition: If $\varepsilon \leq A - a$ and $\varepsilon \leq b$, then the subgame equilibria in G and G' are identical.

Proof of Proposition 1) Suppose play in the $(n - 1)^{th}$ period was (S, S) . Then the payoff for (S, S) in the n^{th} period is A for each player, which is higher than the deviating payoff of $a - \varepsilon$. Thus, (S, S) is a Nash Equilibrium for this subgame. The payoff for (H, H) is $b - \varepsilon$ for each player, which is higher than the deviating payoff of 0. Hence (H, H) is also a Nash Equilibrium in this subgame.

Suppose play in the $(n - 1)^{th}$ period was (H, H) . Then the payoff for (S, S) in the n^{th} period is $A - \varepsilon$ for each player, which is higher than the deviating payoff of a . Thus, (S, S) is a Nash Equilibrium for this subgame. The payoff for (H, H) is b for each player, which is higher than the deviating payoff of $-\varepsilon$. Hence (H, H) is also a Nash Equilibrium in this subgame.

Suppose that the subgame perfect equilibria in the subgame starting from the j^{th} period where $j \in \{2, \dots, n - 1\}$ specifies each period resulting in either (S, S) or (H, H) . We can show that in the subgame starting from the $(j - 1)^{th}$ period, the subgame perfect equilibria specify either (S, S) or (H, H) in period $j - 1$. Let c be the present-value of the payoffs gained from subgame perfect equilibrium at period $j + 1$. Then, given that the subgame perfect equilibrium strategy specifies outcome $X \in \{(S, S), (H, H)\}$ in period j , we notate the present discounted payoff in the $(j - 1)^{th}$ period for player i by the outcome $Y \in \{(S, S), (S, H), (H, S), (H, H)\}$ in the $(j - 1)^{th}$ period as $u_i^{j-1}(Y|X)$.

Case 1) Suppose the subgame perfect equilibria starting from the j^{th} period specifies (S, S) in the j^{th} period.

$$\begin{aligned} u_i^{j-1}(S, S|S, S) &= A + A + c \\ u_i^{j-1}(H, S|S, S) &= a + A - \varepsilon + c \\ u_i^{j-1}(H, H|S, S) &= b + A - \varepsilon + c \\ u_i^{j-1}(S, H|S, S) &= A + c \end{aligned}$$

Clearly $u_i^{j-1}(S, S|S, S) > u_i^{j-1}(H, S|S, S)$, since $A > a$ and $A > A - \varepsilon$. We can also see that $u_i^{j-1}(H, H|S, S) \geq u_i^{j-1}(S, H|S, S)$ since $b \geq \varepsilon$. Due to the symmetry between players, it is clear that the subgame perfect equilibria of the subgame starting from the $(j - 1)^{th}$ period specify either (S, S) or (H, H) in the $(j - 1)^{th}$ period if (S, S) is specified in the j^{th} period.

Case 2) Suppose the subgame perfect equilibria starting from the j^{th} period specifies (H, H) in the j^{th} period.

$$\begin{aligned}
 u_i^{j-1}(S, S|H, H) &= A + b - \varepsilon + c \\
 u_i^{j-1}(H, S|H, H) &= a + b + c \\
 u_i^{j-1}(H, H|H, H) &= b + b + c \\
 u_i^{j-1}(S, H|H, H) &= b - \varepsilon + c
 \end{aligned}$$

Clearly $u_i^{j-1}(H, H|H, H) > u_i^{j-1}(S, H|H, H)$, since $b > b - \varepsilon$ and $b > 0$. Under the condition that $\varepsilon \leq A - a$, then we can also see that $u_i^{j-1}(S, S|H, H) \geq u_i^{j-1}(H, S|H, H)$. Due to the symmetry between players, it is clear that the subgame perfect equilibria of the subgame starting from the $(j - 1)^{th}$ period specify either (S, S) or (H, H) in the $(j - 1)^{th}$ period if (H, H) is specified in the j^{th} period.

Hence, by induction it must be that all subgame perfect equilibria in G' either (S, S) or (H, H) in every period, and that any strategy profile with these outcomes is subgame perfect as long as $\varepsilon \leq b$ and $\varepsilon \leq A - a$.

From this proposition we know that for a small enough value of ε , G and G' do not differ in subgame perfect equilibria. However, we should consider the effects of increased costs brought to light by Goeree and Holt (2005) and strategic uncertainty and belief-formation from Van Huyck et. al (1990) and Crawford (1991) that these frictions could impose. It seems intuitive that these costs, depending on their similarity to the costs in the minimum effort game, could induce higher incentives to opt towards risk dominance despite the unchanged subgame equilibria. However, the belief that an agent may undergo an impending cost in order to achieve pareto dominance may encourage higher rates of playing S in the first round.

We can model these hypotheses using a basins of attraction analysis to narrow down the potential multi-ordered beliefs involved with the play in the initial round of G' . If we assume that agents believe their opponents will play grim trigger strategies, then their choice in the first round is directly dependent on their designated probability of their opponent playing the grim trigger, or simply defecting. Figure 2 shows two potential belief systems that can simplify this complex evolutionary game into a normal form game with an action space of size 2.

	<i>Grim Trigger</i>	<i>Always H</i>
<i>Grim Trigger</i>	nA, nA	$(n-1)b - \varepsilon, a + (n-1)b$
<i>Always H</i>	$a + (n-1)b, (n-1)b - \varepsilon$	nb, nb
	B_1	
	<i>Always S</i>	<i>Reverse Grim Trigger</i>
<i>Always S</i>	nA, nA	$(n-1)A, a + (n-1)A - \varepsilon$
<i>Reverse Grim Trigger</i>	$a + (n-1)A - \varepsilon, (n-1)A$	nb, nb
	B_2	

Figure 2

In game B_1 we are assuming that agents believe that themselves and their opponent is making a choice between playing safe and constantly playing H for all n rounds, or they play a grim trigger strategy that specifies playing S as long as neither player played H any previous rounds. That is, this belief system reflects the idea that playing H is a punishment against your opponent for also playing H . In B_2 the strategies of each player are either constant effort in playing S , or playing H until your opponent shows signs of effort and plays S (this strategy is labeled *Reverse Grim Trigger*). Thus, this belief system represents playing S as a reward to your opponent for undergoing a costly period. The payoffs in each matrix represent the total payoff throughout the n periods without a discount factor.

Consider p_1^* the minimum probability that row player would have to associate with column playing S (Grim Trigger) in order to also play S . Similarly, define p_2^* the minimum probability that row player would have to associate with column playing always S in order to also play S .

$$\begin{aligned}
 p_1^*nA + (1 - p_1^*)[(n-1)b - \varepsilon] &= p_1^*[a + (n-1)b] + (1 - p_1^*)(nb) \\
 \Rightarrow p_1^*[nA - (n-1)b + \varepsilon - a - (n-1)b + nb] &= nb - (n-1)b + \varepsilon \\
 \Rightarrow p_1^* &= \frac{b + \varepsilon}{nA - (n-2)b + \varepsilon - a}
 \end{aligned}$$

To question how the friction ε will affect this required belief, we can take the derivative of p_1^* with respect to ε .

$$\frac{\partial p_1^*}{\partial \varepsilon} = \frac{nA - (n-2)b + \varepsilon - a - b - \varepsilon}{[nA - (n-2)b + \varepsilon - a]^2} = \frac{nA - (n-1)b - a}{[nA - (n-2)b + \varepsilon - a]^2} > 0$$

An increase in the friction will increase the prior belief necessary to make playing S a dominant strategy. Thus, higher values of ε will attribute a low range of player types who will believe that S is a dominant, and thus a high proportion of players whose beliefs designate H as a dominant strategy. We can perform the same calculations on p_2^* .

$$p_2^*nA + (1 - p_2^*)(n - 1)A = p_2^*[a + (n - 1)A - \varepsilon] + (1 - p_2^*)(nb)$$

$$\Rightarrow p_2^*[nA - (n - 1)A - a - (n - 1)A + \varepsilon + nb] = nb - (n - 1)A$$

$$\Rightarrow p_2^* = \frac{nb - (n - 1)A}{nb - a - (n - 2)A + \varepsilon}$$

This probability must be positive, and thus this is only the solution if $nb - (n - 1)A > 0$. Otherwise, $p_2^* = 0$. If $p_2^* = 0$, then there is no local rate of change of p_2^* in ε . For interior solutions,

$$\frac{\partial p_2^*}{\partial \varepsilon} = \frac{(n-1)A - nb}{[nb - a - (n-2)A + \varepsilon]^2} < 0.$$

An increase in the friction will decrease the prior belief necessary to make playing S a dominant strategy. Higher values of ε will attribute a high range of player types who will believe that S is a dominant, and thus a high portion of players whose beliefs designate H as a dominant strategy. Even if $p_2^* = 0$ (which will be true of the experimental parameters chosen in this paper), local increases in the friction parameter will not increase the threshold belief, and will still decrease the payoffs associated with playing H , thus influencing higher play of S .

These two basins of attraction propose two conflicting effects of frictions, each associated with whether or not players interpret the repeated stag hunt as a choice represented by game B_1 (in which playing H is a punishment), or a choice represented by B_2 (in which playing S is a reward). From this analysis we can propose either hypothesis that is easily testable with experimental data. Hypothesis 1 can be supported by the basins of attraction in B_1 , while Hypothesis 2 can be supported by the basins in B_2 .

Hypothesis 1: The higher strategic uncertainty in G' will motivate more risk-dominant behavior in G' than in G . If players expect a lack of coordination and risk averse punishing responses from their opponents, then adding friction will decrease total cooperative behavior.

Hypothesis 2: The impending difficulty of opting toward C in G' if initial play is H will incentivize more payoff-dominant behavior in G' than in G . If players expect a likely choice of future cooperation, then adding friction to the stag hunt will incentivize cooperative play in the initial round.

The inclusion of frictions in G' will generally decrease the incentives of each agent to switch their strategies between rounds. It seems plausible for non-equilibrium outcomes to be achieved less frequently if players are more reluctant to switch, because the classic lack of coordination (S, H in period 1, H, S in period 2, ...) will be less likely if it is more costly to continually switch between S and H . Thus, we can test a hypothesis about the total number of action switches between rounds as well as the total cooperation rates between G and G' .

Hypothesis 3: The costs of switching between actions will cause a lower amount of total switches between actions in G' than G .

If we are to expect less switches between rounds in G' than in G , then it seems logical to expect a general “stability” in G' comparatively to G . There will be less incentive to deviate from either equilibrium with the inclusion of frictions, and thus I expect a higher equilibrium convergence rate in G' than G .

Hypothesis 4: G' will experience higher convergence rates than G within a fixed number of periods, due to the payoff decrease in switching actions.

IV. Experimental Design

Using a between-subject design, all of the subjects in each session played either G or G' . I use a parameterization that closely resembles Battalio et. al (2001), but adjusted with the goal of higher salience in the payoff values. In a treatment in which $A = 45$, $a = 42$, $b = 12$, and $n = 1$, these authors achieved a nearly even distribution of S and H over sessions in which subjects were randomly matched for 75 matches. I adjust these parameters simply to achieve payoffs whose differences are more apparent to the subjects. Specifically, I choose $A = 40$, $a = 36$, $b = 10$, and $n = 4$. The parameters $A = 40$ and $b = 10$ are more easily relatable than those in Battalio et. al, but I ensured to maintain features that are considered essential to the ratio of those who play S . A classic game theoretic prediction about the outcome in coordination games is that the proportion of subjects who play S is negatively correlated to the required belief (basin of attraction)

$$q = P(\text{opposing player plays } S)$$

necessary to make S an optimal strategy for maximizing expected payoff. As shown in the previously mentioned literature, the proportion of subjects playing S is positively correlated with the “optimization premium” δ , which is a measure of the difference in expected payoffs

from playing S or H with a fixed belief q . With these measure is mind, the parameters that I have chosen result in $q = \frac{5}{7}$ and $\delta = 14$ while the parameters in Battalio et. al result in $q = \frac{4}{5}$ and $\delta = 15$. My parameters decrease q which suggest a slight increase in the amount of subjects playing S , but increases δ which suggest a slight decrease in the amount of subjects playing S . I therefore made this change expecting to achieve the previously obtained even proportion between S and H in the control treatment G , allowing for larger and clear differences in the G' treatment.

Between treatments A , a, b , and n will remain constant. The only differences between treatments will be the friction cost ε , which is 0 in the G treatment and 3 in the G' treatment. This parameter is chosen to satisfy the criteria necessary to maintain common strict subgame perfect equilibria across G and G' ($\varepsilon < 40 - 36 = 4$) mentioned in section III, while still providing the largest possible integer cost that is understandable and significant to the subjects.

Subjects will engage in 14 randomly matched anonymous cycles of 4-period repeated games. Their payment will be equal to one tenth of their accumulated payoffs from a single randomly selected match. Paying subjects from a randomly selected match will maintain incentive compatibility across all matches. Providing 4 periods with each pairing of players allows for enough repetition to establish temporary learning of the opponent's strategy, and the possibility of convergence to one of the equilibria within this time frame. For our analysis we will define players in a game to converge if both players in a match play the same action in periods 3 and 4. If each player plays the same action and it is the identical action as their matched player, we consider that an equilibrium convergence. Four periods is minimal enough to achieve convergence without having an excess number or repetitive rounds for players who reach equilibrium convergence early in the match. Fourteen matches will allow agents to learn how to play the game and treat the beginning of each match identically once they have developed a personal strategy.

V. Results

Subjects were recruited from a subject pool of UC Santa Barbara undergraduate students, graduate students, and alumni. Under each treatment G and G' , two sessions were conducted with the intent of having 16 subjects in each session. The G sessions were run with 14 subjects and 10 subjects, while the G' sessions were run with 16 and 12 subjects. The subjects were told that they were to be anonymously matched with 14 people, and with each person they would play 4 rounds of the stag hunt. The random payoff match would be revealed at the end of each experiment.

Analyzing the last 10 matches of each session (as these are the matches in which I consider the subjects to fully understand the experiment), I approximate the following model using a fixed effects linear regression,

$$S_{ij} = \beta_0 + \beta_1 T_j + \gamma_i + \varepsilon_{ij} \quad (1)$$

Where S_{ij} is the probability of an individual i in treatment $j \in \{G, G'\}$ to play S , T_j is an indicator (0 if $j = G$, 1 if $j = G'$), and γ_i represents a fixed effect for each individual i to account for behavioral preferences across subjects.

The errors in all regressions used in this research were clustered across the specific sessions. Experiments with repeated interactions amongst members within a group are subject to session-effects (see Frechette 2012), and thus clustering the errors across sessions accounts for subject specific variation. We run this specified regression on 3 different domains: across all rounds of all matches, across the first round of each match, and across the last match. These three regressions show us a general treatment effect on the proportion of S across the entirety of each match, as well as the beginning and end of each match. Figures 3 and 4 show generic graphical differences in the proportions of subjects playing S between the treatments across matches and rounds respectively.

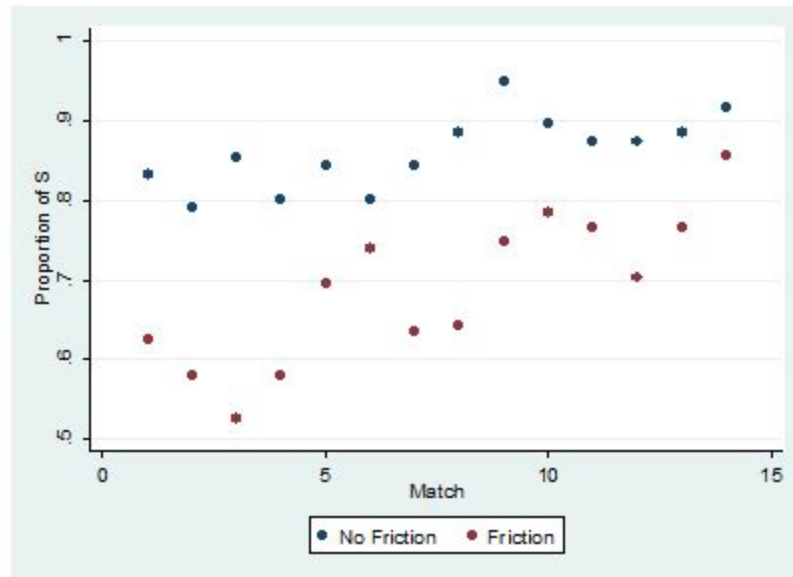


Figure 3

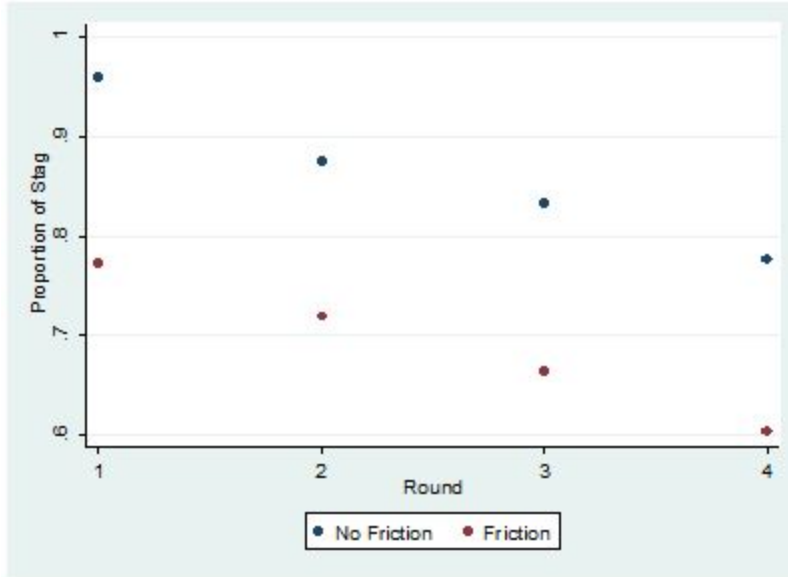


Figure 4

Result 1) Including the friction significantly reduces the proportion of subjects who play *S* in the first round of each match.

The table in Figure 5 reveals the 3 regressions using the main regression model. The coefficient on the Treatment indicator is relatively similar across the 3 subsamples, with a decrease in .142 to .149 probability of *S* in *G'* compared to *G*. However, the only significant coefficient was found in the Round 1 sample. While the decreased cooperation is found across rounds, it is only significant during the first round (at the .01 significance level). In terms of the basin of attraction analysis proposed in section III, this result suggests that agents beliefs may follow in more accordance to the strategies in B_1 rather than B_2 . Thus, it is suggestive that subjects believe that their partners react in a punishing manner rather than a rewarding.

	All Rounds	Round 1	Round 4
	action	action	action
Treatment	-0.142 (1.16)	-0.149 (3.08)**	-0.149 (0.80)
constant	0.877 (10.93)**	0.971 (48.87)**	0.792 (6.17)**
<i>N</i>	2,080	520	520

* $p < 0.05$; ** $p < 0.01$

Figure 5

It is equally important to note the insignificance in the treatment effect of the other two subsamples ($p = .24$ and $p = .42$ respectively). This huge difference in statistical significance between the rounds reveals that while the mean differences across the rounds were relatively similar in our experiment, the standard deviations in the latter rounds diverged tremendously compared to the first rounds. Thus, a larger spread across S and H occurs in the latter rounds of each match, and we cannot note significant differences in the amount of subjects playing S in the latter rounds between treatments.

Result 2) *S proportions were insignificantly different by the end of matches, and coordinating equilibrium convergence rates were insignificantly different between treatments.*

As we saw in the previous result, the overall proportion of subjects playing S was insignificantly different across treatments. We can then determine if the treatments differed in the amount of matches in which the subjects converged to an equilibrium, as well as if the particular equilibriums reached in each treatment differed. First I find no differences in the general convergence when running the regression equation (2).

$$C_{ij} = \beta_0 + \beta_1 T_j + \gamma_i + \varepsilon_{ij} \quad (2)$$

This regression has the same treatment regressor, fixed effects and session-clustered standard errors as regression (1), but the dependent variable C_{ij} is an indicator for convergence in the final round of a match. I run two regressions under this model: one in which the convergence indicator is 1 if the last round has identical action, (S, S) or (H, H), and one in which the convergence indicator is 1 if a subject and their partner matched actions in the last two rounds of each match were identical to each other, and the same in the last two rounds. I refer to these stable 2-period coordination convergence as “mcoordinated” in my regression. These results can be seen in Figure 6, where we can see that with either type of equilibrium convergence, there is no significant treatment effect ($p = .27$ and $p = .28$).

	coordinated	mcoordinated
Treatment	-0.146 (0.62)	-0.138 (0.58)
_cons	0.725 (4.53)**	0.717 (4.32)**
N	520	520

* $p < 0.05$; ** $p < 0.01$

Figure 6

Even with this insignificance in general equilibrium convergence, it could be possible for the achievement of cooperation conditional on convergence to differ across the treatments. Thus, I decided to run regression (1) on the subsample of matches that reached each type of convergence: single round and multi-round. These regression results can be seen in the left and right columns of figure 7 respectively, and once again show very insignificant treatment effects. From the regressions in Figure 6 and Figure 7, it is clear that overall equilibrium convergence between the two treatments is insignificantly different, and thus overall outcomes between G and G' show little variation. There is also an insignificant difference in the number of switches between rounds in each match across treatments. This goes against my expectation for the G' treatments to experience fewer switches due to the friction associated with switching actions.

These results seem inconsistent with the first result that proportions of S are significantly different across treatments in the first round. I attributed result 1 to a difference in the development of prior beliefs about the game. Seeing that the matches in each treatment do not end in different results, it doesn't seem that the prior beliefs would develop differently across treatments. This anomaly leads to some interesting questions about the belief formation in the repeated stag hunt that can be hinted at with a few other results.

Restricted to coordinated matches

	action	action
Treatment	-0.094 (0.95)	-0.081 (0.66)
_cons	0.919 (18.18)**	0.874 (10.92)**
N	2,216	1,656

* $p < 0.05$; ** $p < 0.01$

Figure 7

Result 3) *The cooperation rates in the G treatments were significantly higher than the cooperation rates observed in Battalio et. al (2001).*

The G treatments simply had subjects play randomly matched repetitions of a finitely repeated stag hunt. Battalio et. al (2001) played a randomly matched repetition of a one-shot stag hunt and achieved a nearly even ratio of subject playing S or H . While my treatment parameters were slightly different as discussed in section IV, I justified that the differences in the basin of attraction and optimization premium between my parameters and theirs should support the hypothesis that in a one-shot randomly matched stag hunt my selected parameters would achieve a nearly even proportion. However, as seen in Figure 3, the G treatments had a very large majority of subjects playing S . This cooperative deviation from the one-shot stag

hunt provides a subtle insight into the learning mechanisms of agents in stag hunt games. The higher cooperation rates in the first round of repeated game suggest that subjects learn to expect initial cooperation in the repeated game more than the one-shot game.

Result 4) *S proportions fall between each round of both G and G' treatments.*

This result is perhaps the most surprising and unintuitive result so far. The traditional notion of coordination games (and all games) is that once one of the equilibrium is met by all players, agents will continue to play in accordance to the stability of that equilibrium. However, my results find many subjects deviating from the payoff dominant equilibrium during the progression of each match, despite their initialized state of payoff-dominance. Signalling-based theories of Bayesian updating suggest that reaching a payoff-dominant equilibrium should signal to each player that their opponent will continue to play the payoff-dominant action. Even in the experimental literature, Cooper et. al (1990) find that equilibrium convergence in games similar to the stag hunt (but not actual stag hunt games) occurs according to our theoretical expectations: with subjects stabilizing their actions once a Nash Equilibrium is reached, and most games ending with a Nash Equilibrium outcome. While my results still show a majority of matches ending with a Nash Equilibrium outcome, the significant deviation from the payoff-dominant outcomes in early periods of the repeated game are inconsistent with the conventions assumed by current literature.

The subjects from this experiment are showing high beliefs of cooperation in the first rounds of each match from result 1, but as each match progresses these beliefs are decreasing. For these subjects, previous payoff-dominant convergence is not indicating a strong enough signal to counteract the natural increase in strategic ignorance as the game progresses. One interpretation could be that after each round the number of future rounds per match decreases, which increases the variance of the subject's conditional payoff under the assumption that the partner's action between rounds is independent (similar to a small amount diversification in gambles). This interpretation's assumption of independence across rounds implies that subjects are not updating their beliefs according to the signals provided by the outcomes of previous rounds. Another interpretation could be that subjects do not have self control over their risk aversion during the latter rounds of each match that they had naively discounted in the first round. Regardless, it is clear that subjects are more optimistic about their opponents cooperation in the beginning of each match even after significant learning experience (14 matches), and in some cases the loss of this optimism is strong enough to deviate from the payoff dominant equilibrium. This deviation from payoff dominance is an extremely odd finding relative to the economic status quo of how we think of coordination games.

VI. Hintsights and Future Directions

A large downfall of this experiment was the immense session effects due to the small sample sizes of each session. In particular one of the control G sessions only had 10 subjects, who all played S throughout the first match. The resulting session effect was a near 100% S rate across all 14 matches, aside from a few deviations in the last few matches, which are most likely explained by boredom (see Figure 8). This is the worst possible session effect, which I attribute to the small sample size, and make me skeptical of the overall results. However, all of the previous results are robust when removing this session from the regressions (see Figure 9 for the main regression) with all of the coefficient insignificant at any conventional level, except for the treatment effect during the first round of each match which is significant at the 1% level once again.

To prevent these session effects, it would be helpful to rerun this experiment with larger subject pools. A turnpike design (Frechette 2012) would also help remove session effects, but to guarantee a proper turnpike design would require a session to have twice as many subjects as matches. It would also be useful to have generally more sessions to ensure that these results have greater statistical power and are robust to various session effects.

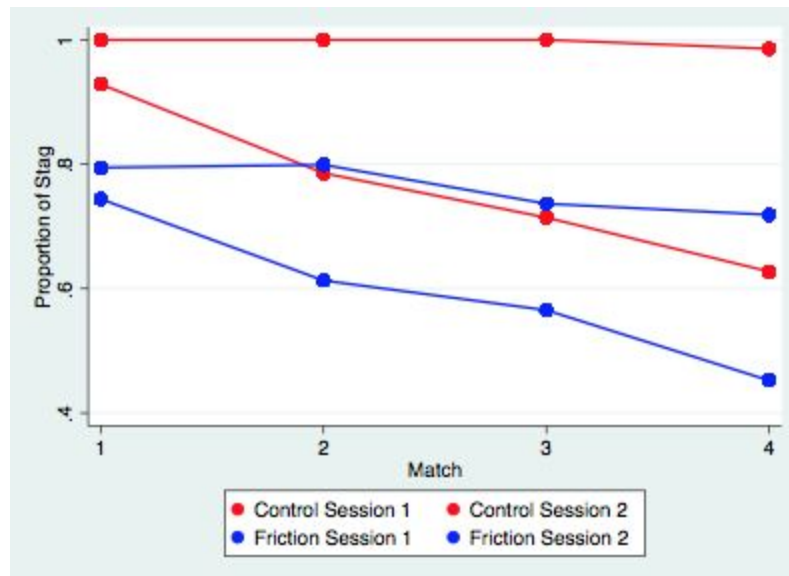


Figure 8

	All Rounds	Round 1	Round 4
	action	action	action
Treatment	-0.058 (0.59)	-0.129 (2.74)**	-0.014 (0.10)
_cons	0.793	0.950	0.657
\bar{N}	1,680	420	420

* $p < 0.05$; ** $p < 0.01$

Figure 9

The parameters chosen for this experiment were chosen with the hopes of achieving even odds of S in the control treatment, but in hindsight we see that this was not the case. Thus, if this experiment were to be replicated it would be important to alter the payoff matrix in order to motivate evenly split behavior. An even split in the control treatment would guarantee the friction treatment has maximal room for deviation from the control.

An unexpected important result of this research is the difference between the observed and expected behavior in a repeated stag hunt. The results in this paper suggest that the current assumptions about basic coordination games may not be consistent with actual behavior, especially in a repeated setting. While many similar settings have been studied, it seems that repeated stag hunts have not been experimentally investigated in depth, and our assumptions in both finite and infinite repetitions seem to be unjustified and may in fact be inaccurate. Further investigation of stag hunts, while simple and seemingly conventional, could perhaps open new doors into the theory behind coordinating behavior.

VII. Conclusion

The results of this experiment reveal important weaknesses in the current economic understanding of repeated coordination games. The difference between these experimental treatments should have no differences in outcomes if we are to categorize each treatment by their Subgame Perfect Equilibria. The difference found between the initial play in each match across treatments leads to interesting questions about the development of subjects beliefs throughout the session and the directions in which they learn to play coordination games. While the treatment effect suggests that subjects expect punishing behavior from their partners, we also see from the control treatment itself new insights that could be gained from studying repetition in coordination games. The amount of cooperative outcomes in the repeated stag hunt differ tremendously from previous literature on the one-shot stag hunt with similar payoffs, and we also find that agents are not maintaining stability in the payoff dominant Nash Equilibrium, which is contradictory to our general notions of equilibrium. The combination of results in this paper suggest that more research in the direction of belief formation and learning over time is essential in our understanding of asymmetric coordination games.

References

Battalio, Raymond, Larry Samuelson, and John Van Huyck. "Optimization incentives and coordination failure in laboratory stag hunt games." *Econometrica* 69, no. 3 (2001): 749-764.

Cooper, Russell W., Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross. "Selection criteria in coordination games: Some experimental results." *The American Economic Review* 80, no. 1 (1990): 218-233.

Crawford, Vincent P. "An "evolutionary" interpretation of Van Huyck, Battalio, and Beil's experimental results on coordination." *Games and Economic behavior* 3, no. 1 (1991): 25-59.

Fréchette, Guillaume R. "Session-effects in the laboratory." *Experimental Economics* 15, no. 3 (2012): 485-498.

Goeree, Jacob K., and Charles A. Holt. "An experimental study of costly coordination." *Games and Economic Behavior* 51, no. 2 (2005): 349-364.

Harsanyi, John C., and Reinhard Selten. "A general theory of equilibrium selection in games." *MIT Press Books* 1 (1988).

Nagel, Rosemarie. "Unraveling in guessing games: An experimental study." *The American Economic Review* 85, no. 5 (1995): 1313-1326.

Schelling, Thomas C. *The strategy of conflict*. Harvard university press, 1980.

Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil. "Tacit coordination games, strategic uncertainty, and coordination failure." *The American Economic Review* 80, no. 1 (1990): 234-248.