Testing the Efficacy of Stepping Stone Equilibria in Coordination Games*

Ethan Holdahl†

November 15, 2023

Abstract

Games with multiple equilibria introduce the potential for populations to get stuck in inefficient outcomes. In theory, the introduction of additional equilibria, "stepping stones", could pave the way for a smoother and less risky transition. I run a lab experiment to test if the introduction of these "stepping stones", can facilitate transitions from an inefficient but safe equilibrium to a risky, payoff dominant equilibrium. I employ different payoffs for the transition strategy and examine the effects that different degrees of information about the game have on group's play. I find evidence that adding these "stepping stones" does help populations transition to the efficient equilibrium. I also find that when groups have more information about each other's payoffs they are able to transition to the efficient equilibrium faster and are less prone to cyclical behavior.

*I thank Jiabin Wu, Van Kolpin, Anne van den Nouweland, Sherry Li, Stephan Jagau, the University of Oregon 2023 cohort, and ESA 2022 conference members for their helpful comments. I am also indebted to the recruiting efforts of Jiabin Wu and Mike Peixoto. Any remaining errors are my own.

†Department of Economics, University of California, Irvine, eholdahl@uci.edu

1 Introduction

As a social species, coordination games are ever present in our lives. From the language we speak, where we choose to live, to how much effort we put into our work, often the best decision depends upon matching the choices made by those we interact with. Coordination games are characterized by having multiple equilibria and through repeated play, a convention of a group playing the same equilibrium can be established. However, not all equilibria are as desirable as others; Once a convention has been established, the transition from one equilibrium to another, even to one that is a Pareto improvement, is inherently difficult.[1] Consequently, there is potential for groups to get stuck in conventions of playing equilibria which are payoff dominated. This potential is what makes the study of variations of the stag hunt game, first considered by Jean-Jacques Rousseau in his Discourse on Inequality in 1755, of considerable interest to economists spanning the gamut from experimentalists to macro theorists [Cooper and John, 1988, Romer, 1996, Bryant, 1983]. In this paper, I introduce, experiment with, and examine the effectiveness of the addition of a *stepping stone* equilibrium crafted to aid the transition from the payoff dominated equilibrium to the Pareto efficient equilibrium in the classic stag hunt game.

There is some current research on populations being stuck in inefficient conventions. Bursztyn et al. [2023] describes the social media applications of TikTok and Instagram as collective traps. In an experiment they found that college students would prefer that the products don't exist but because of their popularity they primarily participate due to "fear of missing out". Thus, they are unwilling to remove the apps unless others do trapping them in participation while lacking an effective intermediary transition pathway. Another concurrent paper is Gulesci et al. [2023] which examines empirical evidence of female genital cutting practices in Somalia. Coincidentally also defines *stepping stones* as a transitory state that enables transitions in the intermediate run. The primary difference between my definition and theirs is that in their definition, *stepping stones* are strictly transitory where as I consider conventions, which are self-enforcing, as *stepping stones* in stochastic games. Gulesci et al. [2023] examines practice of female genital cutting in Somalia and treats the norm as a discrete choice problem between three options: a high invasive practice called "Pharaonic", a milder practice called "Sunna", or no cutting. They find that over the past 50 years, Sunna has almost complete displace Pharaonic circumcision. Yet Sunna seems to be an absorbing state as the proportion of uncut remains very low. As such, Gulesci et al. [2023] discusses the implications of trying to correct for harmful norms by creating transitions which may end up being absorbing and creating a new, still not ideal,

[1] An equilibrium E is a Pareto improvement over another equilibrium E' if all players weakly prefer E to E' and at least one player strictly prefers E to E'.

norm. With this in mind, my paper is relevant as it establishes a connection between the *stepping stone* and the facilitation of easier stochastic transitions in the long run. Consequently, supporting such transitions leads to monotonically increasing welfare.

A current example of a transition in progress is the question of personal vehicle choice in the United States. Among a wide range of disadvantages, gas cars tend to produce higher emissions, cost more to operate, and require more frequent maintenance than their electric counterparts [Malmgren, 2016, Harto, 2020]. However, making the switch from gas to electric can be an unappealing decision for many due to the dependence on fueling infrastructure and mechanics. In this sense, vehicle choice is a coordination game, as more people switch to electric, more charging stations are built[2]. While gas stations are nearly omnipresent, the relative scarcity of charging stations can make driving certain routes much less efficient,[3] if not impossible in an electric vehicle. This may explain why in 2018 when passenger vehicles contributed 29% of total US greenhouse gas emissions, electric vehicles (EVs) only accounted for 2% of US auto sales [4].

Trying to change equilibrium selection in these group coordination games with so much inertia behind them can be a challenging and expensive task. Continuing with the EV example, in an effort to speed up the transition from gas to electric, the recent Inflation ReductioFiguren Act is estimated to cost over $14 billion in clean vehicle spending over the next ten years, primarily though EV tax credits [5]. Directly incentivizing the desired strategy should help increase the transition speed to that equilibrium, however, with how important and costly transitions like these are it is valuable to be as efficient as possible. Theoretically, there may be other mechanisms that can help a population transition from one strategy to another that are more efficient; If there exists another strategy that provides an easier and faster transition from the initial equilibrium to the desired state, then the creation or investment in that path may be more efficient than directly incentivizing the desired equilibrium. In the case of vehicle choice, the plug-in hybrid vehicle (PHEV) could be viewed as the strategy to facilitate that transition. The PHEV boasts some of the benefits of EVs, as it has a short electric only range, which is sufficient for most daily tasks, without as many cons, since it is able to use gas. This position the makes the transition from gas to plug-in relatively easy. If plug-ins were then widely adopted, that would incentive more charging stations to be built which would

[2]Firms and the government also build charging stations to help stimulate adaptation, which can be though of as equivalent in effect to a proportion of population adopting the electric choice

[3]The charging rate in electric batteries decreases as current charge level increases. For example, a Tesla 3 can charge from 0% to 50% in 15 minutes using a Tesla Supercharger, however, it takes an additional 41 minutes to charge it from 50% to 100% [Hoffman, 2020]. Consequently, the most time efficient strategy when driving a long route is to only partially charge the battery at charging stations, thus reducing total time spent charging. However, this method requires sufficient charging station density, otherwise drivers may have to spend more time charging than what would otherwise be efficient just to make it to the next charging station.

[4]United Auto Workers [2020]

[5]Congressional Budget Office [2022]

make the full transition to EVs easier. Thus, the inclusion of a PHEV as a strategy theoretically acts as a stepping stone, a strategy that makes the transition from gas to electric easier.

To examine if *stepping stones* are effective, I design a study to test if *stepping stones* impact transition dynamics in group coordination games. In the experiments, subjects played 200 rounds of a stag hunt coordination game where the group was initiated with starting at the safe, Pareto dominated, equilibrium. Groups were treated with *complete* or *incomplete information* about other's payoffs and group played games with a *high payoff stepping stone*, *low payoff stepping stone*, and *no stepping stone* (control). Halfway through the 200 rounds, any *stepping stone* strategy was removed and players played the *no stepping stone* game for the remaining 100 rounds with everyone starting at the safe equilibrium again. The idea being to see if any treatment effects in the first 100 rounds would impact play after the treatment is removed, even in the worst case scenario.

The results of the experiment show that groups were able to transition to and play the efficient equilibrium quicker and more often when there was a *stepping stone* present, providing experimental evidence that a third equilibrium can speed up the transition from the initial to the efficient equilibrium. Additionally, the presence of compete information about the other players payoffs impacted the group's play. When *complete information* was present, players more quickly and more briefly utilized the transition strategy to jump to the efficient equilibrium strategy, something they were less willing to do when there was *no stepping stone*. However, when players only knew their own payoffs, they frequently used the transition strategy and through it more slowly arrived at the Pareto efficient equilibrium. In short, more information about the game appeared to make individuals more froward-looking. As such, it seems that the mere presence of a *stepping stone* may be enough to change behavior, even if the strategy isn't frequently chosen.

Related Literature

There is a large literature focused on experiments addressing the challenge of coordination failures [Devetag and Ortmann, 2007]. Games like the minimum effort game [Van Huyck et al., 1990], are among the most studied and illustrate that in coordination games even though a Pareto efficient outcome may be a Nash equilibrium, it may be difficult to achieve if it is risk dominated by another equilibrium. While many studies of the repeated minimum effort game have shown that as group size increases, groups are more likely to exhibit coordination failure [Van Huyck et al., 1990, Knez and Camerer, 1994, Goeree and Holt, 2005, Di Girolamo and Drouvelis, 2015], the game examined in this paper should be less prone to coordination failure due to group size since payoffs depend on the full distribution of other player's strategies, not just the minimum effort level.

There is a small, growing experimental literature studying deviations from myopic best response behavior in laboratory games. Hwang et al. [2018] examines which convention emerges between five strategies in a bargaining game. They find that deviations from myopic best response is payoff dependant as their subject displayed intentional bias. They suggest that this mechanism leads to the egalitarian solution being the most likely bargaining norm to evolve. Lim and Neary [2016] studies the behavior of a large population playing the Language Game of Neary [2012]. They find that deviations depend on the myopic best-response payoff but not on the deviation payoff and that deviations decrease over the rounds played in the experiment. Mäs and Nax [2016] tests the decisions made in a coordination game when players are in a fixed network as described by Ellison [1993]. They also find that deviations are payoff dependant and that there is evidence of individual heterogeneity in those deviations. I contribute to this literature by studying a non-competing coordination game with payoff rankable Nash equilibria.

The idea of intermediate transitions being used to enable a population to move from one state to another (i.e. the process of biological evolution) can perhaps first be attributed to the 19^{th} Century naturalists Wallace [1858] and Darwin [1859]. 140 years later, Ellison [2000] formalized this idea of step-by-step evolution in a game theory setting. Here, I give name, *stepping stones*, to those states that are used to speed up or enable evolution from an initial state to a latter state.

The games studied here are most similar to those found in Cooper et al. [1990]. Like Cooper, I use a stag hunt game and expand it to a 3x3 game and investigate how the added strategy might affect equilibrium selection and transition dynamics. However, the games vary as Cooper added a dominated strategy to the stag hunt game to help solve the coordination failure problem. In addition, in Cooper et al. [1990] players only played 20 rounds, against each other player twice. Since the added strategy was a dominated one, it is reasonable to think that as play evolves the frequency that the dominated strategy is played, and along with it the belief that others would play it, would converge to 0. Thus resulting in the irrelevance of dominated strategies for equilibrium selection [Kohlberg and Mertens, 1986]. I differ from Cooper by adding a third equilibrium to the game instead and examining play in an evolutionary setting against the same players for 100 rounds. This experiment also differs as the game begins with the risk dominant equilibrium as the established choice.

The remainder of the paper is organized as follows: In the next section I introduce the theoretical framework and design of the experiment followed by the hypotheses and the procedure. Following that I report the results of the experiment at the group and individual level. Finally, I provide a summary of the results in my concluding remarks.

2 Experiment

In this section I lay out the design for the experiment followed by my hypotheses and then the procedure. The primary objective of the experiment is to test if and how effective transitory equilibria are in coordination games where players play against the field.

2.1 Design

When players make decisions in a game we assume they are best responding to their beliefs about what the other players in the game will play. Most evolutionary game theory assumes that individuals play the myopic best response the majority of the time [Kandori et al., 1993, Canning, 1992, Young, 1993]. This translates to players believing that their opponents will play the same action in the future as they did in the past. Recent experimental evidence from evolutionary games supports this idea as 90% to 96% of decisions from those experiments could be explained by myopic best response play [Hwang et al., 2018, Mäs and Nax, 2016, Lim and Neary, 2016].

Since myopic best response appears to well explain behavior of players in long repeated games, it is natural to use Young [1993]'s adaptive play model of learning, which is entirely backwards looking, to form the basis of analysis for evolutionary games. As such, I adapt the theoretical framework from Young [1993] and Ellison [2000] to fit a repeated game where everyone in a population plays every period and all plays of that period are observed and recalled by all players in the subsequent period:

2.1.1 Theoretical Framework

Let G be a symmetric w-strategy game. A population of N players repeatedly match with all other players in the population to play G in every period $t > 0$. In this environment each player i has a semi-persistent strategy, $s_i(t)$ they use to play G every period. In each period, players can attempt to change their strategy by selecting an action $a_i(t) \in A$ to play. Since this is a symmetric game, the action sets for all players are the same. [6] After selecting their action, with probability p their strategy is updated such that $s_i(t) = a_i(t)$. Otherwise, their strategy remains unchanged from the previous period: $s_i(t) = s_i(t-1)$ where $s_i(0)$ is given. Each period after players strategies have been determined, G is played $N-1$ times, once against every

[6] Since the stage game has w strategies, it follows that the number of actions in A, $|A|$ is w.

other player, earning a total period payoff of $\Pi(s_i(t), s_{-i}(t)) := \sum_{j \neq i} \pi(s_i(t), s_j(t))$ where $\pi(s_i, s_j)$ is the payoff player i receives playing against player j in the game G.[7]

During each period t, each player i observes each other players' last period strategy, $s_{-i}(t-1)$ and responds by playing the action that maximizes their payoff given the other players strategies remain unchanged: $a_i(t) \in BR_i(s_{-i}(t-1)) := \arg\max \{\Pi(a_i, s_{-i}(t-1)) \mid a_i \in A\}$. Such an action choice is referred to as a myopic best response.

Definition 1. *MYOPIC BEST RESPONSE: A decision that is a best response to the previous period's strategy profile.*

However, instead of playing their best response, occasionally players choose an action at random. For some $\epsilon \in [0, 1/w)$, each player randomly selects an action with probability $w\epsilon$. With probability $1 - w\epsilon$ the player selects an action that is in their set of myopic best responses. Actions that are not myopic best responses are called "mistakes" and are thus played with probability ϵ.

Given that only the strategies from the most recent period of play are considered each period, the probability of advancing to some strategy profile in period $t+1$, $s(t+1)$ depends only upon the strategy profile in period t, $s(t)$. Thus, the set of all possible states is equal to the set of strategy profiles. I refer to both as S which is equal to the set of action profiles: $S = A^N$. As such each strategy profile is a state in a discrete-time homogeneous Markov process as detailed by the decision rules above where $P^{\epsilon}_{ss'}$ describes the probability of moving directly from state s in one period to s' in the subsequent period. Through an unperturbed process where $\epsilon = 0$, a self-enforcing pattern of play, called a convention, has the potential to arise.

Definition 2. *CONVENTION: A convention is a state s such that in the unperturbed process $P^0_{ss} = 1$.*[8]

Thus, to escape a convention requires perturbations. As perturbations are assumed to occur infrequently, the least number of perturbations required to transition from one state to another, known as the resistance, describes how difficult the transition is to make.

Definition 3. *RESISTANCE: For any two strategy profiles s, s' the resistance $r(s, s')$ is the number of perturbations required to make the direct transition from s in one period to s' in the subsequent period.*

Given any two distinct states s, s', consider all the directed paths that begin with s and end with s' and call the collection of those paths $Z_{ss'}$. Among all such paths, let $\zeta^*_{ss'}$ be the path with the least sum total

[7] Note that this is a symmetric game so it doesn't matter if a player is the "row player" or the "column player". This payoff setup is equivalent to one where every player plays once against a randomly drawn opponent and players are risk neutral.

[8] Since memory is restricted to 1, a strategy profile is a strict pure strategy Nash equilibrium if and only if it is a convention.

resistance for each step along each directed path. Let t be the period the directed path $\zeta^*_{ss'}$ starts such that $s(t) = s$ and $T > t$ be the period s' is reached so $s(T) = s'$ and define $r_{ss'}$ as the sum total of the resistance for every step: $r_{ss'} = \sum_{t}^{T-1} r\big(s(t), s(t+1)\big)$. Thus, $r_{ss'}$ measures the least amount of perturbations necessary to transition from s to s'.[9]

Define a recurrent class, $E \subseteq S$, which has the property $r_{ss'} = 0$ and $r_{s's} = 0$ if and only if $s, s' \in E$. In other words, a recurrent class is a closed subset of states that the unperturbed Markov process cannot escape from once it enters the class. As such, a strict pure strategy Nash equilibrium is by itself a recurrent class in this setting.

Note that the resistance of the transition between recurrent classes: $E_1, E_2, \ldots, E_K$ is largely characterized by the difficulty of escaping the basins of attraction $D(\cdot)$ of the initial recurrent class. The following definition is due to Ellison [2000]:

Definition 4. *BASIN OF ATTRACTION: A state s is said to be in the basin of attraction of a recurrent class E if in an unperturbed process:*

$$s \in D(E) := \{s \in S | Prob(\exists T > t \;\; s.t. \;\; \forall t' > T \;\; s(t') \in E | s(t) = s) = 1\}$$

As such, in order for transitions to occur from recurrent classes E to E', play must first escape the basin of attraction of E and then make its way into the basin of attraction of E'. This transition may occur in one step, or play could, for example, move from E to $D(E'')$ to E'' then to $D(E')$ and finally to E'. If the path of least resistance doesn't involve a direct transition from $D(E)$ to $D(E')$ but instead includes transitioning to some other recurrent class E'' then I call E'' a *stepping stone* from E to E'.

Definition 5. *STEPPING STONE: A stepping stone from one recurrent classes E to another, E' is a recurrent class E'' if the path of least resistance $\zeta^*_{ee'}$ from some $e \in E$ to some $e' \in E'$ includes $e'' \in E''$.*[10]

Now construct a directed graph with K vertices, one for each recurrent class. Call the vertex corresponding to the recurrent class E_i vertex i. The weight on the directed edge from vertex i to vertex j is r_{ij}. The tree rooted at vertex i contains $K - 1$ directed edges such that there is a pathway from each vertex $j \neq i$ to vertex i. The resistance of each rooted tree is calculated as the sum of weights of the $K - 1$ directed edges in the tree.

[9] Note that s' does not need to immediately succeed s. It may often be the case where the path of least resistance involves multiple steps. See an example in Figure 5.

[10] Note that this definition of stepping stone differs from that of Gulesci et al. [2023] which examines a stepping stone in an intermediate run dynamic and defines stepping stones strictly as states belonging to a transient class. The definition I use is similar to the step-by-step process used by Ellison [2000] and how using a step-by-step process can affect the modified coradius.

Figure 1: Path of Least Resistance from A to C in two 3x3 Games

No Stepping Stone from A to C

B is a Stepping Stone from A to C

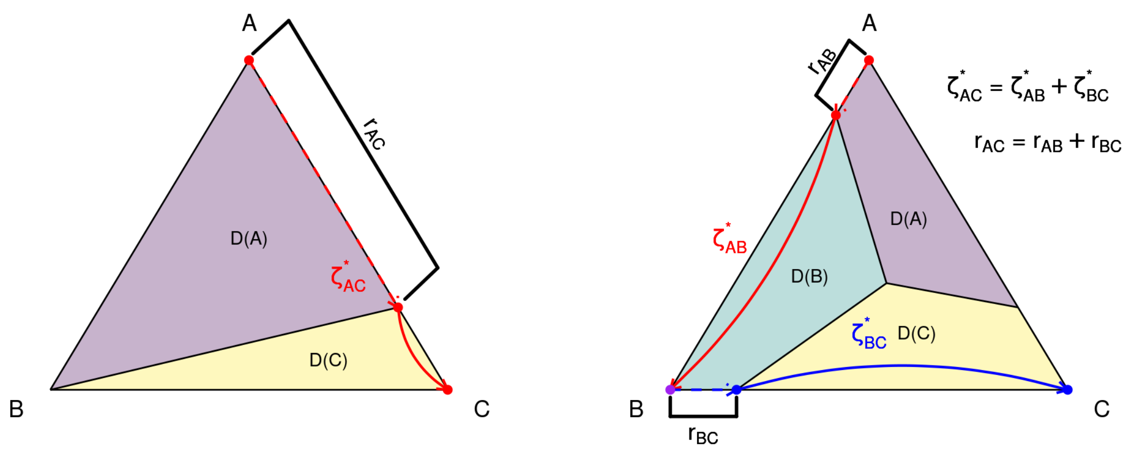

In the game on the left there are only 2 recurrent classes: pure strategy equilibria of all A and all C. The path of least resistance is to in one step move from A to the nearest point in $D(C)$ and then move from $D(C)$ to C with no added resistance. In the game on the right, the path of least resistance is instead from A to $D(B)$ to B to $D(C)$ to C.

Definition 6. *STOCHASTIC POTENTIAL: The stochastic potential, γ_i, of a recurrent class i is the tree rooted at i with the lowest resistance.*

Young [1993] showed that the stochastically stable states are those contained in the recurrent class with the minimum stochastic potential in the game.

Definition 7. *STOCHASTIC STABILITY: A state s is stochastically stable if s has the smallest stochastic potential of all states.*

Given that the inclusion of a stepping stone E'' in a game reduces the resistance from E to E', stepping stones not only reduce the resistance of directed transitions but may also affect the set of stochastically stable states.

The goal of the experiment is to test if stepping stone equilibria are effective as theoretically predicted. That is to say, that when populations play long repeated games they utilize stepping stones to transition between conventions. In the next section, I describe how I choose the payoffs in each game played in the experiment.

Figure 2: Trees rooted at C

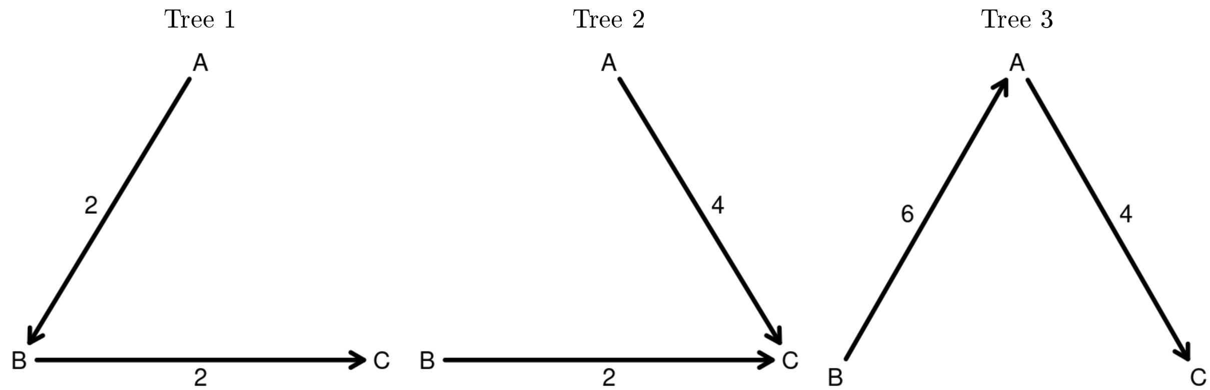

In a game with 3 recurrent classes there are 3 vertices, each with 2 directed edges such that there is a pathway from each vertex to C. In this example, The resistance of each rooted tree is 4, 6, and 10 for Tree 1, Tree 2 and Tree 3 respectively. As such, the stochastic potential of γ_C is 4.

2.1.2 Treatment Design

Figure 3: Parameter Selection

		Player 2		
		A	B	C
Player 1	A	a_G	b_G	e_G
	B	c_G	d_G	g_G
	C	f_G	h_G	i_G

The game is symmetric and the reported payoffs are those of the row player

I design a lab experiment to test if players utilize stepping stones and if so, if there are certain aspects of the game that make a stepping stone more effective. Of particular interest, I study how the strategy profile of the group evolves when playing a coordination game with Pareto-rankable equilibria when the population starts the game at a Pareto dominated equilibrium.[11] To do this, I use a 2x3 treatment design.

One dimension of the treatment design is selecting the games groups played. Groups of size 8 were asked to play one of three augmented stag hunt games all of which had groups starting the game playing E_A. In two of the three games (Game 2 and 3), a stepping stone was added to the classic stag hunt game. In Game 2 a stepping stone that payoff dominates the starting equilibrium was added. I will refer to this as the high payoff stepping stone treatment. In Game 3 I instead add a stepping stone that is payoff dominated by the starting equilibrium. I will refer to Game 3 as the low payoff stepping stone treatment. The motivation for

[11] A game with Pareto-rankable equilibria means that at least one equilibrium is preferred by all players in the game to another equilibrium.

the different stepping stone levels is to see if payoff dominance in transitions makes a difference in group's play. In the control game (Game 1), groups played without a stepping stone and instead with an added strategy that guaranteed the worst payoff possible.

On the other dimension of the treatments, I varied the amount of information subjects were given. Players were given either *complete information* meaning that they knew both their payoffs and the payoffs of the other players in the game, or players were given *incomplete information* meaning they did not know what payoffs the other players received for a given strategy profile. This is important since having common knowledge that players are playing a coordination game and are starting at an inefficient equilibrium may impact player's decisions, presumably by making them more forward looking.

I create a 3×3 symmetric game with action space $\{A, B, C\}$ with the payoffs received by the row player in game $G = \{1, 2, 3\}$ as depicted in Figure 3 with three pure strategy Nash equilibrium: $E_A = (A, A)$, $E_B = (B, B)$, and $E_C = (C, C)$ in each game. With 9 payoff variables in each game G, parameters were chosen to create a large difference in the path of least resistance from E_A to E_C between the game without a stepping stone, Game 1, and the games with a stepping stone, Games 2 and 3, while observing certain restrictions, namely:

1. E_A and E_C must be strict Nash equilibria. Thus, $a_G > c_G, f_G$; $i_G > e_G, g_G$ $\forall G$
2. E_B must be a strict Nash equilibrium in Games 2 and 3. Thus, $d_i > b_i, h_i$ $i \in \{2, 3\}$
3. E_C must be the Pareto Efficient equilibrium. So, $i_G > a_G, d_G$ $\forall G$
4. The variables a_G, e_G, f_G, i_G must remain the same across all games.
5. The resistance from E_A to E_B must be equal to the resistance from E_B to E_C in and between Games 2 and 3.[12]
6. In Game 1 (no stepping stone), b_1, c_1, d_1, g_1, h_1 must be equal to the lowest payoff in the game.
7. The variables a_G and e_G must be equal so that Game 1 is essentially a stag hunt game.
8. $a_G + e_G > f_G + i_G$ so that E_A pairwise risk dominates E_C.
9. In Game 2 (high payoff stepping stone), the payoff at the transition equilibrium must be greater than that at the starting equilibrium $d_2 > a_2$.

[12]By controlling for the resistance between equilibria between and across games I can test for the effect of pairwise and global payoff dominance.

10. In Game 3 (low payoff stepping stone), the payoff at the transition equilibrium must be smaller than that at the starting equilibrium $d_3 < a_3$.

11. To make calculations as simple as possible for subjects, all payoffs must be single digit integers.

12. Payoffs must be greater than 0 to avoid behavioral distortions [Tversky and Kahneman, 1991, Gneezy and Potters, 1997].

In all games E_A pairwise risk dominates E_C which implies that E_A is the stochastically stable equilibrium in the 2x2 game with just A and C which is the result of Theorem 4.2 in Young [1998]. This dynamic lends additional justification for initiating the game with everyone playing A.

2.1.3 Games and Theoretical Analysis

Figure 4: Games Played in the Experiment

Game 1

Player 1	Player 2: A	B	C
A	7	1	7
B	1	1	1
C	1	1	9

Game 2

Player 1	Player 2: A	B	C
A	7	3	7
B	6	8	4
C	1	7	9

Game 3

Player 1	Player 2: A	B	C
A	7	1	7
B	6	6	4
C	1	5	9

The games are all symmetric and the reported payoffs are those of the row player

Figure 4 shows the construction of the three different games played in the experiment. All three games are symmetric with the payoffs reported being those of the row player. In addition, players play against the field meaning they are playing against the full distribution of pure strategies of all other players in the population. All 3 games are coordination games with 3 pure strategy Nash equilibria on the diagonal. I will refer to each equilibrium, where all 8 players play the same strategy, as follows: $E_A = (A, A, \dots, A)$, $E_B = (B, B, \dots, B)$, $E_C = (C, C, \dots, C)$. Game 1 represents the case where there is no stepping stone from E_A to E_C. Note that Game 1 is essentially a stag hunt game. I include B as a strategy to increase the confidence that any change in play between different games is due to the change in payoffs and not due to a change in the strategy space of the game. As a benefit, including B in game 1 guarantees players will receive the worst possible payoff in the game allow us to test if random errors that don't take into account payoffs occur.[13]

[13] This is assuming that players don't utilize B as a way to punish other players, or in games with *incomplete information*, use B to see if that may help other players move from A to C

The idea of trying to solve the coordination failure problem as seen in stag hunt games with the addition of another strategy was examined in Cooper et al. [1990]. However, I differ here by adding a third equilibrium, a stepping stone, to game 2 and 3 instead of a dominated strategy. Note that Game 3 is constructed by taking a payoff transformation of Game 2 that preserves the best reply structure, specifically by subtracting the payoff the row player receives when the column player plays B by 2 [Harsanyi et al., 1988]. This transformation allows me to test what, if any, effect payoff dominance for the stepping stone strategy has. In theory, there shouldn't be any difference between individuals whose play is motivated by myopic payoff differences. However, there is reason to believe that player's decisions are also influenced by payoff dominance [Harsanyi et al., 1988, Jagau, 2022].

Resistance Calculations

In order for the added equilibrium, E_B, in Games 2 and 3 to be considered a stepping stone from one equilibrium, E_A, to another, E_C, the path of least resistance from E_A to E_C must go through an indirect path through E_B compared to the most efficient direct path from E_A to $s \in D(E_A)$ to E_C.[14] Simply stated, E_B is a stepping stone if the resistance from E_A to E_C is smaller in Games 2 and 3 than Game 1.

Below I calculate the resistance from E_A to E_C in Games 1, 2, and 3. Refer to Figure 5 for a depiction of the directed paths of least resistance on a simplex. Note that the graphs depict the mapping of the strategies of the other players in the group from the perspective of a player who always plays their myopic best response. This distinction is made since players best respond to the other players' strategies and this set of strategies will differ across players if their own strategies are not the same as one another. As such, by examining the best reply structure of a player who always plays their myopic best response, the minimum number of "mistakes" necessary to change the myopic best response of some players, and following that, the entire group is revealed. This is possible because of stochastic strategy updating implemented in this game. The path of least resistance does not necessarily require transitioning directly from E_A to $s' \in D(E_C)$. It is possible to transition from E_A to some $s \in D(E_A)$ then with no further mistakes but with selective stochastic strategy updating, make the transition directly from s to E_C. I elaborate below.

First I will describe the path of least resistance in Game 1, which is also the directed path of least resistance that travels directly from E_A to E_C without entering $D(E_B)$ in Games 2 and 3. Starting from E_A (1) the transition into $D(E_C)$ can be accomplished with the least "mistakes" required by transitioning

[14] All strict Nash equilibria are recurrent classes in this game.

Figure 5: Path of Least Resistance from E_A to E_C in Games 1, 2 and 3

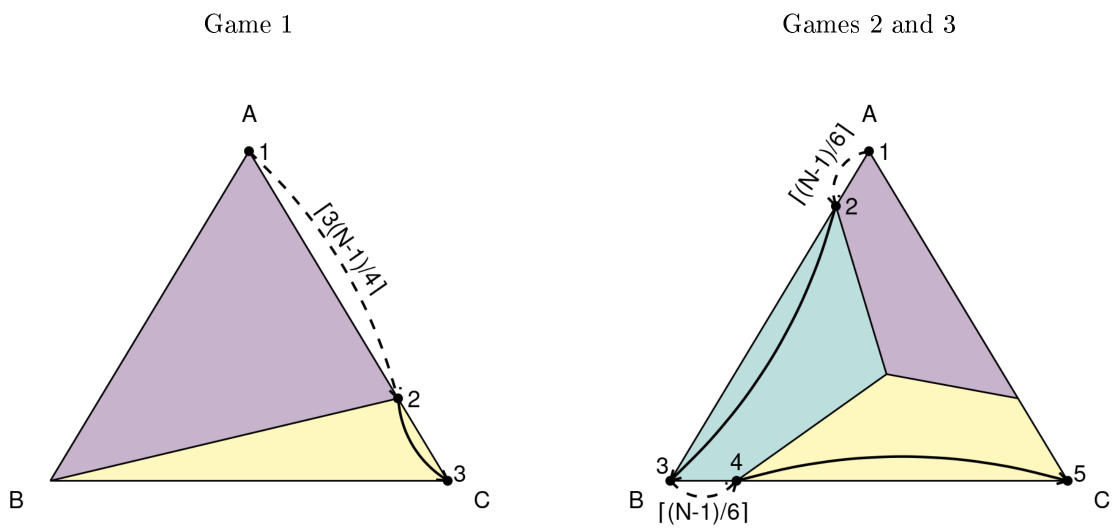

Recall: The best reply structure is the same in Games 2 and 3. Because resistance is a count of only non-myopic best response play (represented by dashed curves), once the game is within the basin of attraction of an equilibrium $D(E)$, it can travel the rest of the way to that equilibrium using only best response play (represented by solid curves). Note that the paths are curved only to make following the steps easier to follow. The directed path accurately mapped onto the simplex is a straight line. The basins of attraction for each recurrent class is color coded in the simplex: purple for E_A, green for E_B, and yellow for E_C.

to a state where $\lceil 3(N-1)/4 \rceil$[15] of the players play C and all other players play A, requiring a minimum of $\lceil 3(N-1)/4 \rceil$ "mistakes". Once in this state, (2), the remaining $N - \lceil 3(N-1)/4 \rceil$ players who have A as their current strategy calculate their expected payoffs for playing against their opponents' last period strategies. If they play A then their expected payoff is $7(N-1)$. If they play C then their expected payoff is $N - 1 - \lceil 3(N-1)/4 \rceil + 9\lceil 3(N-1)/4 \rceil \geq 7(N-1)$. So they can play C as a best response. The other $\lceil 3(N-1)/4 \rceil$ players whose current strategy is C have A as their unique best response. Since $\lceil 3(N-1)/4 \rceil \geq N/2$ it can be shown that the current state is in $D(E_A)$. However, it is possible that all players whose last period strategy was A now play C and their strategies are all stochastically accepted while all the players last period strategy was A now have their actions stochastically rejected. Hence, E_C (3) is reached. The resistance from E_A to E_C in Game 1 is therefore $\lceil 3(N-1)/4 \rceil$.

Now I will describe the path of least resistance in Games 2 and 3. Starting from E_A (1) the transition to a state in $D(E_B)$ can be made by $\lceil (N-1)/6 \rceil$ players picking B as their action by "mistake" and all those actions being stochastically accepted. In the next period (2), the players who did not make a mistake get

[15] The notation $\lceil x \rceil$ means rounding up to the nearest integer that is greater than or equal to x.

and expected payoff of $7(N - \lceil (N-1)/6 \rceil) + 3\lceil (N-1)/6 \rceil$ if they pick A, $6(N - \lceil (N-1)/6 \rceil) + 8\lceil (N-1)/6 \rceil$ if they pick B, and $(N - \lceil (N-1)/6 \rceil) + 7\lceil (N-1)/6 \rceil$ if they pick C. Clearly, the expected payoff for C is less than B. So B is a best response if $7(N - \lceil (N-1)/6 \rceil) + 3\lceil (N-1)/6 \rceil \leq 6(N - \lceil (N-1)/6 \rceil) + 8\lceil (N-1)/6 \rceil$. The expression simplifies to $N \leq 6\lceil (N-1)/6 \rceil$, so B is a best response. So play can transition to E_B (3) without any additional "mistakes". Once at E_B, the transition to a state (4) in $D(E_B)$ can be made by $\lceil (N-1)/6 \rceil$ players picking C as their action by "mistake" and all those actions being stochastically accepted. It is simple to calculate, similar as above, that play can then progress with no further "mistakes" to arrive at E_C (5). Hence, the resistance from E_A to E_C in Games 2 and 3 is $2\lceil (N-1)/6 \rceil$.

As such, the directed paths of least resistance using and not using E_B can now be compared. If $2\lceil (N-1)/6 \rceil < \lceil 3(N-1)/4 \rceil$ then E_B is a stepping stone in Games 2 and 3. It is easy to verify that for all $N > 3$, E_B is a stepping stone in Games 2 and 3.

In the experiment, groups of size $N = 8$ were used. As such, E_B is a stepping stone in Games 2 and 3. In all three games, when the game is at E_A, the amount of simultaneous deviations from myopic best response needed to make C a best response for the remaining players in the next period is then $\lceil (3/4) * (N-1) \rceil = 6$. In Games 2 and 3, the resistance from E_A to E_B as well as the resistance from E_B to E_C is $\lceil (1/6) * (N-1) \rceil = 2$. Examining the resistance in all three games from E_C to E_A is $\lceil (1/4) * (N-1) \rceil = 2$. In Games 2 and 3, the resistance from E_C to E_B and the resistance from E_B to E_A is $\lceil (1/4) * (N-1) \rceil = 6$. This means if perturbations are independent and payoff independent then if conventions change in Games 2 and 3 they should travel almost exclusively in the direction from E_A to E_B to E_C to E_A et cetera spending on average equal time at each.[16] However, if deviations from myopic best response are a function of payoff dominance then populations should spend a greater portion of their time at the payoff efficient equilibrium E_C when cycling or not cycle at all.

In several evolutionary game theory experiments in the literature [Hwang et al., 2018, Mäs and Nax, 2016], subjects were randomly given an opportunity to change their strategy in each round. In the event they weren't given a revision opportunity, their action from the previous round was retained. This is valuable from a data collection standpoint as it slows down the transition from one equilibrium to another, which is where decisions are most important. As discussed in the theory section, in this experiment for all games and in every round, all subjects will be asked which action they want to play. However, with probability p their new action is adopted and with probability $(1-p)$ their strategy from the previous round is retained. In essence, in this experiment I am moving the nature node deciding if they can update their strategy from

[16] Note that by definition just as E_B is a stepping stone from E_A to E_C so too is E_C a stepping stone from E_B to E_A and E_A is a stepping stone from E_C to E_B.

before to after the subject makes their decision. This change in procedure yields the benefit of being able to collect $1/p$ times as much data. This procedure is similar to the strategy method [Selten, 1967] which is often used to boost data collection in extensive form games.

For the experiment trails I use $p = 1/2$. The benefit of using a relatively small p value is that it makes states stickier, thus making equilibria more stable. Additionally, it helps enforce the initial condition of all the games in the experiment: that the game starts with everyone playing A, corresponding to E_A, the safe, payoff dominated equilibrium. This "stickiness" can be demonstrated by considering a player who uses level-K thinking [Nagel, 1995, Stahl and Wilson, 1995]. Consider a level-1 player. They assume every other player plays each strategy with equal probability. So, they expect to face the mixed strategy of $(1-p+p/3, p/3, p/3)$. Thus, since here I use $p = 1/2$, their expected payoffs from playing each strategy is $(36/7, 1, 11/6)$ in Game 1, $(38/7, 36/7, 20/7)$ in Game 2, and $(36/7, 34/7, 18/7)$ in Game 3 for each of strategy (A, B, C) respectively. Thus, A is the unique best response in each game. Level-2 players assume all other players are level-1 players and thus will also play A. It follows that for all players of level-$K > 0$ thinking have A as a best response.[17] Now consider if $p = 1$, the case where players are always able to change their strategy. In this case, level 1 thinkers assume they are facing a mixed strategy of $(1/3, 1/3, 1/3)$ which means their expected payoffs from playing each strategy is $(5, 1, 11/3)$ in Game 1, $(17/3, 6, 17/3)$ in Game 2, and $(5, 16/3, 5)$ in Game 3 for each of strategy (A, B, C) respectively. Thus, in Games 2 and 3, B is their best response when $p = 1$ as in this case enforcing that everyone starts the game playing A is little more than a default option [Thaler and Sunstein, 2008, Samuelson and Zeckhauser, 1988]. This example shows how incorporating stochastic strategy updating can affect decisions and enforce initial conditions.

2.2 Procedure

18 sessions (3 per treatment) of 8 participants each were held in person at the Tattersall Computer Lab at the University of Oregon. A total of 144 subjects were recruited from the University of Oregon student population, each of which made 200 decisions over the course of one hour and were paid, on average, $21.

Each participant was seated at a computer with dividers between the monitors and all participants were seated facing a wall to prevent any in-person interaction or viewing of others' screens. The software used for the experiment was built using oTree [Chen et al., 2016], which is software using Python, HTML, and JavaScript designed for use in laboratory and field experiments in game theory.

[17] Note that in the *incomplete information* treatment players only know their own payoffs so computing the best responses of other players can not be done reliably, especially at the start of the game.

The instructions, quiz questions, and screenshots of the UI during the experiment can be found in the appendix.

Phase I

Upon entering the lab and filling out consent forms, participants were read aloud instructions explaining how the game works and how the experiment is conducted. Typed instructions were also be visible on their computer.[18]

Phase II

Participants were provided with a writing utensil, a basic calculator, and blank paper to make notes and calculations if they desired to use them. After reading the instructions, participants took a short quiz for comprehension to ensure that they understand how the game works and how their payouts would be calculated. Participants had to answer each question correctly before they could proceed to the following question. The number of errors made by each participant was tracked.

Phase III

Participants then played the experiment which comprised of two sets of 100 rounds each. During each round, participants were able to view the payoff they earned in the previous round, the strategies played by the other participants in the previous round, if their action last round was accepted or rejected, the remaining time lest in the round, their payoff table, and depending on the treatment of the study, their opponents payoffs in the payoff table. In each round, participants were be able to change their strategy with probability $= .5$ otherwise, their previous round strategy was retained. Everyone started the experiment coordinating on A. In the first two rounds of each set, participants had 60 seconds to pick a strategy. In rounds 3-5, participants had 45 seconds to pick a strategy, In rounds 6-8, participants had 30 seconds to pick a strategy, in rounds 9-11, participants had 20 seconds to pick a strategy, and in rounds 12-200 participants had 10 seconds to pick a strategy. This shrinking decision time is commonly used in similar evolutionary experiments [Lim and Neary, 2016, Hwang et al., 2018]. Failure to select a strategy in a round resulted in the player's previous round strategy being selected for them. The amount of time it took a participant to select their choice in each round was also recorded.

During the first set of 100 rounds, the group played their treatment game (either Game 1, Game 2, or

[18] I did not read aloud the payoff tables in order to preserve the *imperfect information* treatment.

Game 3). After the conclusion of the first set of 100 rounds the participants were brought to an screen informing them of possible changes that were being made to their payoff table for the second set of rounds. Every group played Game 1 in the second set of rounds. However, the treatment of revealing/concealing the other players' payoffs was maintained for each group across sets.

2.3 Hypotheses

Hypothesis 1: In the sessions where players play a game where E_B is a stepping stone (Games 2 and 3), the groups will be more successful in making a transition from E_A to E_C than the sessions where E_B is not a stepping stone (Game 1).

> This hypothesis is to test if stepping stone equilibria actually work as predicted: to reduce the amount of time it takes to transition from E_A to E_C by creating a path of lower resistance. Under all arms of the study with a stepping stone, both $\max(r_{AB}, r_{BC})$ and $r_{AB} + r_{BC}$ is less than the direct transition r_{AC}, meaning transitions from E_A to E_C are theoretically more probable in Games 2 and 3 than in Game 1.

Hypothesis $\mathbf{2_0}$: In the games with stepping stones (Game 2 and 3), players will spend the same number of periods with each action as their myopic best response.

Hypothesis $\mathbf{2_A}$: In the games with stepping stones (Game 2 and 3), players will spend a plurality of the periods played with the action corresponding to the Pareto efficient equilibrium as their myopic best response.

> As discussed when calculating resistances, since the weight on the directed edges from vertex A to B, B to C, and C to A are then same when $N = 8$, if "mistakes" are uniformly random then each state E_A, E_B, and E_C are stochastically stable. Which means that the adaptive process is expected to spend an equal time at each equilibrium.
>
> However, if deviations from myopic best response are a. payoff dependant or b. a function of equilibrium payoff dominance then populations should spend a greater portion of their time near the Pareto efficient equilibrium, E_C. This is because a. The cost per game played of the first deviation from A to B and B to C is only 1, where as the cost per game played of the first

deviation from C to A is 2.[19] The explanation for preference for playing the payoff dominant equilibrium, b., is self-evident.

Hypothesis $\mathbf{3_0}$: The rate of deviations from the myopic best response of A to choosing action B will be no higher when the stepping stone payoff dominates the starting equilibrium (Game 2 vs Game 3).

Hypothesis $\mathbf{3_A}$: The rate of deviations from the myopic best response of A to choosing action B will be higher when the stepping stone payoff dominates the starting equilibrium (Game 2 vs Game 3).

The theoretical prediction under uniform random errors is that the transition dynamics of these games should be identical since the resistance between equilibria are the same in Games 2 and 3. Myopic payoff dependant deviations also produces the same prediction since when comparing Game 2 to Game 3, the expected payoffs increase by the same amount, depending on how many other players play B, for all strategies a player can choose from.

However, there is reason to believe the transition speed may be faster in Game 2 than Game 3. This is because in Game 2 E_B payoff dominates E_A where E_A payoff dominates E_B in game 3. In essence, this hypothesis tests if deviations from a myopic best response towards a stepping stone are a function of payoff dominance.

Hypothesis $\mathbf{4_0}$: Deviations from myopic best response are payoff independent.

Hypothesis $\mathbf{4_A}$: Deviations from myopic best response are payoff dependant and occur less frequently as the difference between expected payoff of the myopic best response and the next highest expected payoff increases.

Hypothesis $\mathbf{5_0}$: Strategy update success will not influence rate of deviation from myopic best response.

Hypothesis $\mathbf{5_A}$: Strategy update success will increase rate of deviation from myopic best response.

Although update probability is constant and independent, subjects who recently experienced a low update success rate may view deviations as riskier behavior.

Hypothesis $\mathbf{6_0}$: In Game 1, 1/2 of the "mistakes" made are players choosing action B.

Hypothesis $\mathbf{6_A}$: In Game 1, less than 1/2 of the "mistakes" made are players choosing action B.

[19] I remind the reader that each player plays the game against all other players in each round.

This hypothesis is similar to hypothesis 4, but the result is more straightforward since there are few ways to rationalize playing B in Game 1. If B accounts for significantly less than half of the "mistakes" then we have evidence that "mistakes" are not uniformly distributed.

Hypothesis $\mathbf{7_0}$: The rate of deviations from the myopic best response when the myopic best response does not correspond to the Pareto efficient equilibrium will be equal across games with *complete information* vs *incomplete information*.

Hypothesis $\mathbf{7_A}$: The rate of deviations from the myopic best response when the myopic best response does not correspond to the Pareto efficient equilibrium will be higher in games with *complete information* vs *incomplete information*.

Under *incomplete information*, it will take sophisticated subjects time to realize that this is a coordination game, if in fact they do, and they may never realize that their payoffs align in such a way that E_C is Pareto efficient and that E_B is a stepping stone from E_A to E_C.

If players have *complete information*, however, they will immediately know that this is a coordination game and that E_C is an equilibrium and the Pareto efficient outcome. Consequently, a deviation from an equilibrium may be viewed by other players as a costly signal towards a new equilibrium.

If subjects are sufficiently sophisticated, they will realize in Games 2 and 3 that using E_B as a stepping stone is a more efficient path towards E_C than just going directly from E_A to E_C, both in terms of payoff forgone in the transition, and the number of likewise deviations needed to shift the myopic best response. However, some players may instead view the direct jump as a faster method of transitioning. In either case, their rate of deviation from myopic best response when the myopic best response isn't A should be higher than the groups with the *incomplete information* treatment.

If this hypothesis is true, then *complete information* should have an attenuation effect of the amount of time it takes to transition from one equilibrium to the next.

3 Experimental Results

3.1 Group Level Results

Stepping Stone vs. No Stepping Stone

The primary goal of the experiment was to test if the inclusion of a stepping stone equilibrium was effective in facilitating the transition from the initial equilibrium, to the Pareto efficient equilibrium. Here I discuss the first set of each experiment during which groups played one of three games for 100 rounds: Game 1 which had no stepping stone, Game 2 which had a high payoff stepping stone, and Game 3 which had a low payoff stepping stone.

Figure 6 shows the evolution of groups' strategies in the first 100 rounds of 6 different sessions, one from each treatment. If the stepping stones were effective in facilitating transitions from A to C then Groups playing a game with a stepping stone should make it to E_C with higher probability and consequently, spend more rounds playing C.

I find that in all 12 sessions where groups played with a stepping stone they were able to, at least once, make it to E_C. By contrast, in the 6 sessions where players played Game 1 in the first set, only 4 groups were able to make it to E_C. I test Hypothesis 1 using a 1-sided Fisher's Exact Test which yields a p-value $=$ 0.09804. While this is above the .05 threshold traditionally required to reject that groups are just as likely to reach E_C when playing Game 1, it does support the idea that stepping stones are effective as theoretically predicted.

Beyond the binary of "did a group transition to E_C?", the proportion that each strategy was played can be analyzed. Table 1 reports the proportion that each strategy was played in each experiment. Note that in 4/6 of the sessions with no stepping stone A accounted for the majority of strategies in set 1. This is in sharp contrast to the games played with a stepping stone where C made up the plurality of the strategies in every experiment. I use a mixed logistic regression with clustering at the individual and experiment level, the results of which can be found in Table 5, which show that players played C in Game 2 and Game 3 significantly more than in Game 1.

Figure 6: Time Series of Group Strategy in Set 1 (Treatment)

Complete Information Games | Incomplete Information Games

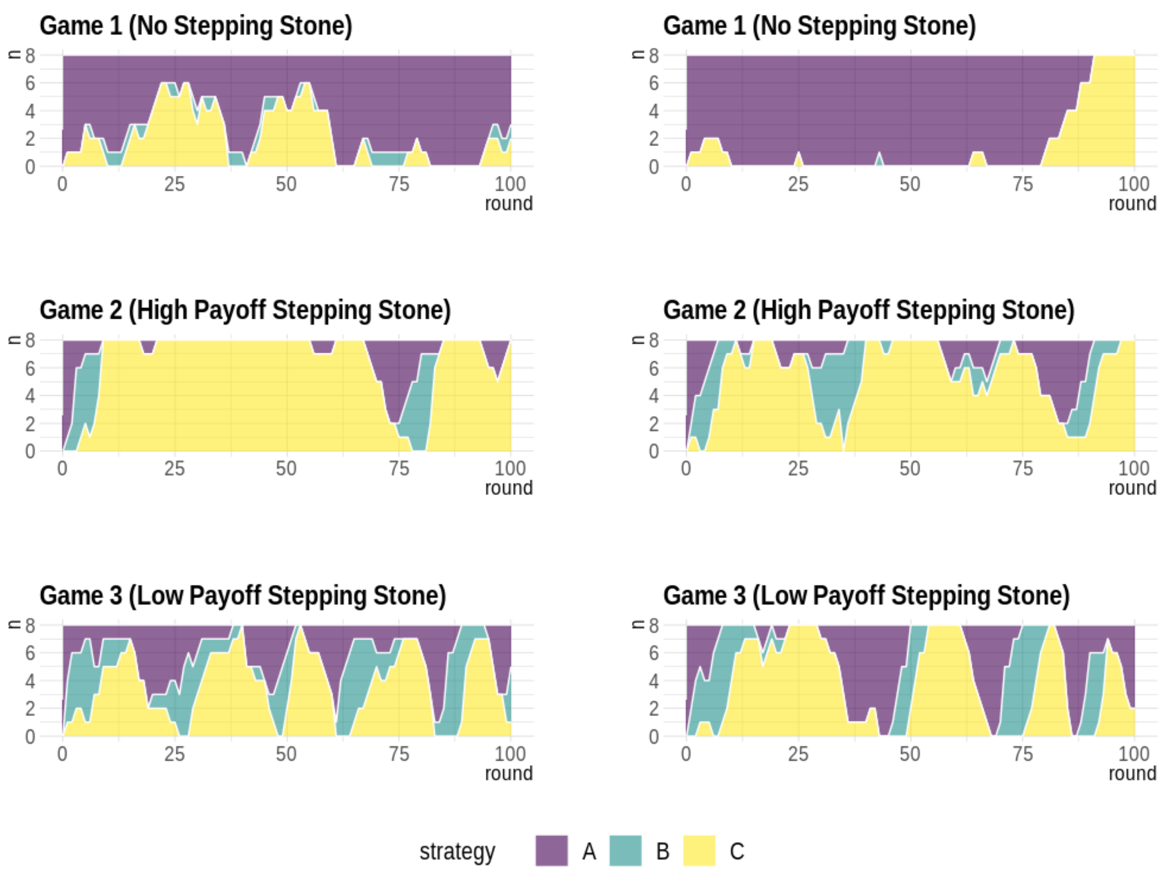

These stacked area plots depict how the proportion of each strategy played changed as the round number increased. The proportion that a strategy was played in a given round is equal to the vertical length with that strategy's color code. The time series of each set of each experiment can be found in the appendix.

High vs Low Payoff Stepping Stones

I've established that stepping stones were effective, here I examine if there was a difference in the effectiveness of the low payoff stepping stone compared to the high payoff stepping stone. Looking again at the regression in Table 5, Game 2 had a point estimate of 1.8157 and Game 3 had a point estimate of .8588. Both with standard errors approximately .31, this is a large and significant difference between the two, indicating that C is significantly more likely to be played in the sessions with a high payoff stepping stone compared to a low payoff stepping stone.

Looking at the frequency of strategy played is informative but doesn't provide much insight into the frequency with which strategy profiles were played. Figure 7 does just that. In Figure 7 and Figure 8, I map the strategy profile faced by each player each round onto a simplex, linearly interpolate between the nearest points to flesh out the graph, then color code by frequency, standardized across plots so they can be compared. I also report the frequency that each action is a myopic best response (mBR) to the strategy faced. As can be seen in the Figure 7, play is much more concentrated around E_C in the games with a high payoff stepping stone compared to the games with a low payoff stepping stone. In fact, in the games with a low payoff stepping stone, A and B were myopic best responses twice as often as they were in games with a high payoff stepping stone.

Figure 7: Frequency of Strategy Profile Faced

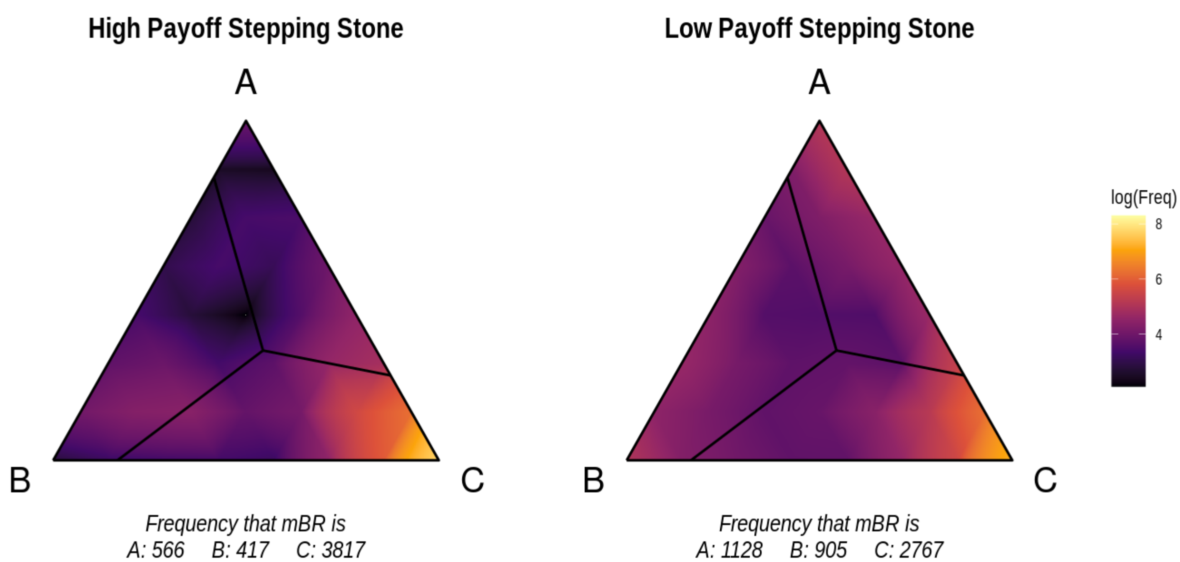

Figure 7 illustrates one of Hypothesis 2, which examines whether players in games with stepping stones (Game 2 and 3) spend the same number of periods with each action as their myopic best response. In the sessions, players were observed to have as their myopic best response action A 1694 times, action B 1322 times, and action C 6584 times.

I conduct a chi-squared test to see if this difference is significant. I get X-squared = 5389.5 with 2 degrees of freedom, which reveals a highly significant p-value of less than $2.2e^{-16}$. This provides strong evidence which indicates that in Games 2 and 3, players tend to spend more time with C, the action corresponding to the Pareto efficient equilibrium as their myopic best response. It is noteworthy that despite the fact that all three states are stochastically stable under uniform perturbations, players exhibit a preference for the

Pareto dominant equilibrium.

Complete vs Incomplete Information

Figure 8: Frequency of Strategy Profile Faced

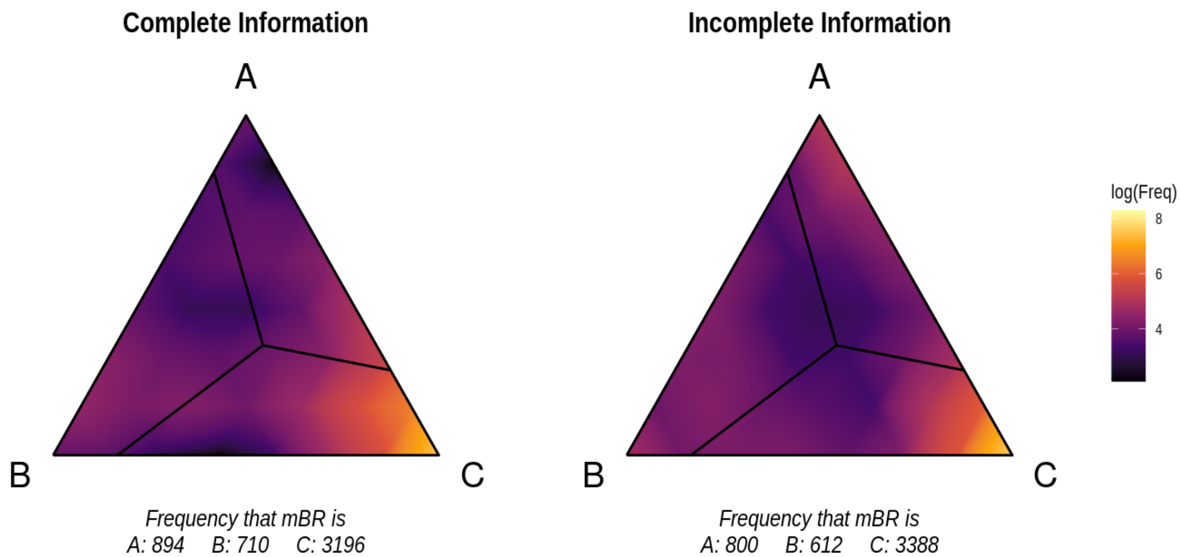

When comparing groups that played with a stepping stone with *incomplete information* to those who played with a stepping stone with *complete information*, in aggregate as in Figure 8 the results appear almost identical. Summary tables 1 and 2 hint towards the biggest difference between games played with *complete information* vs *incomplete information* are when there are no stepping stones present. During the first 100 rounds, players played A with the highest frequency every time (n=3) when they played Game 1 with *incomplete information*. By contrast, in two of the three sessions with *complete information*, groups were able to make the transition and play C for the majority of the set. The explanation for this is essentially that in *complete information* games players know that they are playing a coordination game and that it is in the group's best interest to transition to E_C. As such, it reasons that groups may have a higher propensity to play C even if A is the best response as playing C would likely be viewed as a signal that the player wants to move the group to C and is willing to pay the upfront cost.

It is perhaps because the barrier to transition out of A was so reduced by the stepping stone that there doesn't appear to be much difference between the complete and *incomplete information* treatments with stepping stones. In this sense, one could think that stepping stones are particularly useful for attenuating the difficulty in coordinating inherent to some environments and populations.

Patterns of Play

There were two primary patterns of group play observed once a group made the transition to C. The group would then either stay at C for the remainder of the set, or would fall into cyclical behavior of playing $A \rightarrow B \rightarrow C \rightarrow A \rightarrow \ldots$ until the end of the set. Naturally, those who played Game 1 never fell into the cyclical behavior since playing B guarantees the worst payoff possible. However, more than that, groups who played Game 1 and made it to E_C were the most stable, perhaps recognizing that getting back to C would be difficult if they deviated.

More interestingly, several of the experiment with stepping stones exhibited cyclical behavior. This occurred most frequently in games with *incomplete information* and games with a low payoff stepping stone. There doesn't seem to be a clear reason for why the cyclical process gets initiated, perhaps due to boredom or competitive behavior.[20] However, once the process back to A starts, other subjects are quickly pressured by the payoffs to transition as well. From A they transition back to C through B. This cyclical behavior creates a positive feedback loop through players expectations. Players learn that play moves from A to B to C to A and the stochastic updating probability encourages players to make decisions based on where they think play is heading least they get left behind and punished. This dynamic should be particularly pronounced in games with *incomplete information* as individuals start the game with no basis for strong prior beliefs as to how their opponents will play. Thus, if they see cycling, they may think they are not playing a coordination game.

Set 2 Results

After playing 100 rounds of either Game 1, Game 2, or Game 3, all groups played Game 1 (no stepping stone) for the second set of 100 rounds to see if playing with the stepping stone had any effect. For example, can using stepping stones as a crutch in the short term foster long term coordination in other games between the same population?

Overall, I do not find evidence of correlation between the treatment in the first set and the performance in the second set. Of the eighteen groups, ten of them made it to E_C in the second set with three of the groups having played Game 1 in set 1, four having played Game 2, and three having played Game 3. Four of the ten groups were playing with *complete information* and the remaining six with *incomplete information*.

[20] Sheremeta [2010] has shown that in contests with a prize of zero some subjects are still willing to bid to "win".

What appears to be the biggest predictor of success, meaning making it to E_C and staying there, is if the group ended the previous set at E_C. Although every groups' strategy profile was reset to E_A at the start of set 2, none of the groups who ended set 1 with a strategy profile not E_C were able to make it to and stay at E_C in the second set. By contrast, of the twelve groups who did end set 1 at E_C, eight of them made it back to and stayed at E_C in set 2.

This effect was driven by the groups with *incomplete information* where across all three games, every group except for one (5/6) that ended at E_C in set 1 ended at E_C in set 2. Among the three groups with *incomplete information* that didn't end set 1 at E_C, none of them made it to E_C in set 2. This result is significant under Fisher's Exact Test yielding a p-value of 0.04762.

3.2 Individual Level Results

At the individual level, I am examining the decisions made by each player. In particular, I examine the rate of myopic best response at different positions in the game and test my remaining hypotheses. Table 3 shows the choices made in each experiment by mBR and table 4 shows the aggregated choices by mBR. As expected, most frequently subjects played their myopic best response with a few notable exceptions: in 3 treatments action A was selected as a myopic best response less than half the time. This occurred in both treatments of the high payoff stepping stone game, and in the *incomplete information* treatment of the low payoff stepping stone game.

Figure 9: Choices in High Payoff Stepping Stone Games

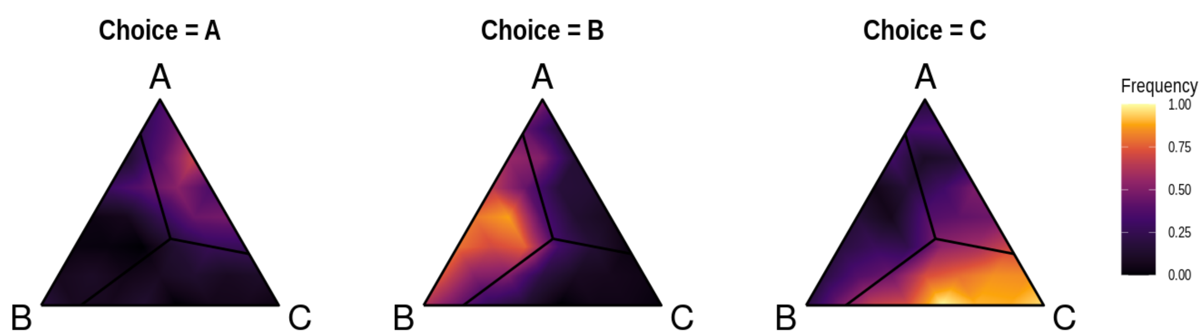

As discussed in the previous section, these are clear difference in play between Games 2 and 3. Recall, the difference between Game 2 and Game 3 is that in Game 2 the row player's payoff increases by 2 when the column player plays B;[21] As a result, in Game 2 E_B payoff dominates E_A and in Game 3 E_A payoff

[21]and because the game is symmetric, the column player's payoff also increases by 2 when the row player plays B

Figure 10: Choices in Low Payoff Stepping Stone Games

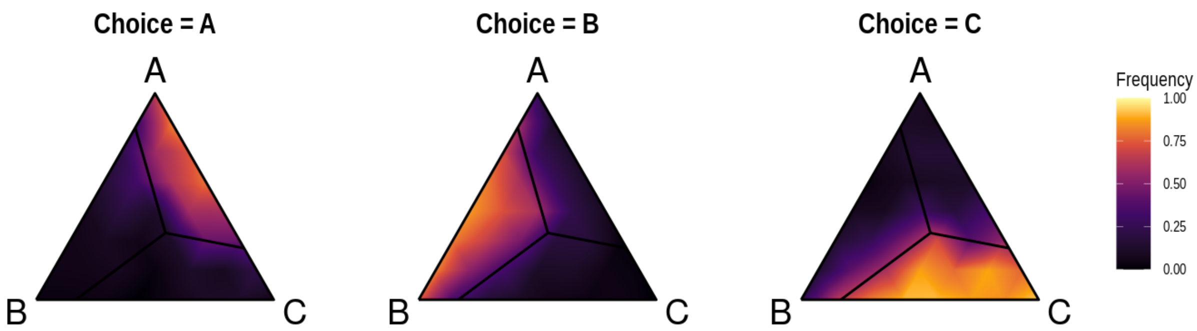

Figure 11: Choices in Complete Information Games with a Stepping Stone

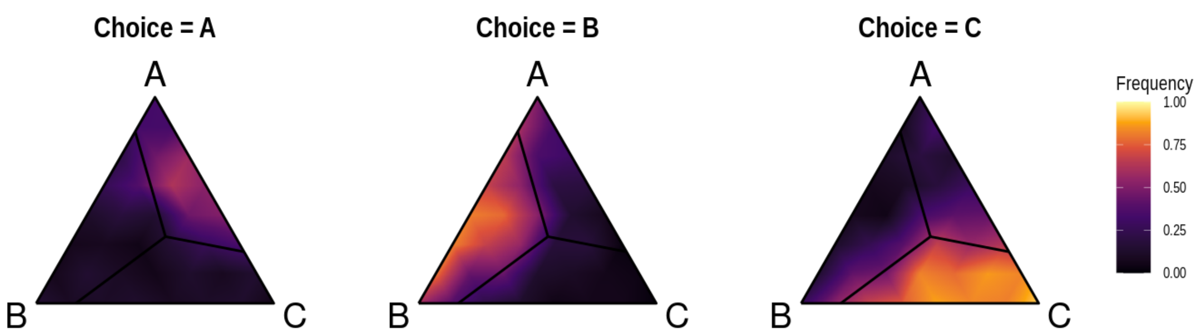

dominates E_B. Since Game 3 is just a payoff transformation of Game 2 that preserves the best reply structure, theoretically, there shouldn't be any difference between individuals whose play is motivated by myopic payoff differences. As such I hypothesised that the biggest difference would be in the transition from A to B as there is evidence that payoff dominance between equilibria plays a role in players' choices [Harsanyi et al., 1988, Jagau, 2022].

I investigate hypothesis 3 by using a generalized logistic mixed model with individual fixed effects to test if the rate of deviation from the myopic best response of A to choosing action B varies significantly across Games 2 and 3. See Table 6 for regression results.

The analysis reveals a difference between the estimates for Game 2 (intercept) and Game 3. The estimated difference is -0.1030. The p-value associated with this difference is 0.697. As result, this test provides no evidence that the rate of deviations towards the stepping stone are payoff dominance dependant. So I can not reject the null hypothesis 3_0.

Figures 9-12 show the proportion of choices made given the strategy profile faced in different treatments.

Figure 12: Choices in Incomplete Information Games with a Stepping Stone

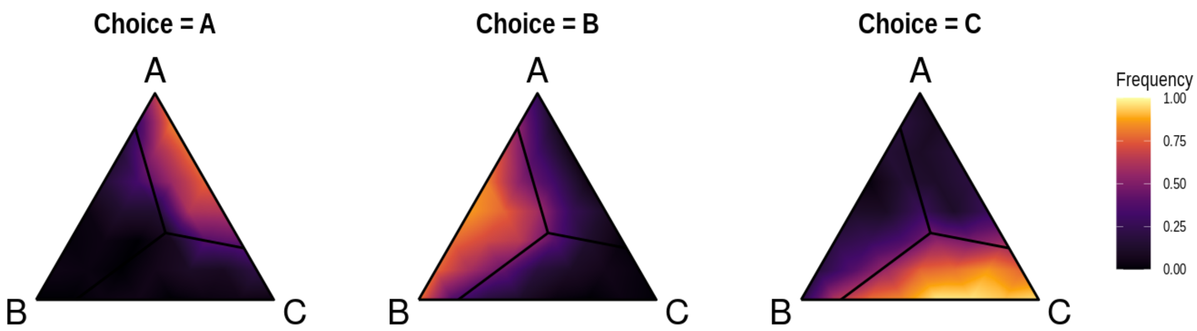

The graphs were constructed in a similar manner to Figures 7 and 8, by linearly interpolating across the simplex. The groups that played with a high payoff stepping stone or with *complete information* played A with much lower frequency than those in *incomplete information* games and low payoff stepping stone games. When comparing the high payoff stepping stone to the low payoff stepping stone this result may be explained by E_A Pareto ranking higher in the low payoff stepping stone game.[22] On the information side, the difference in play can be explained by the uncertainty that a better outcome is stable and not being able to observe an efficient path to achieve it. When there is common knowledge of the game, decisions are more likely to be interpreted as intentional signals and players are likely to believe that others will want to transition to C as fast as possible.

Next, I look at what factors influence players to deviate from playing their myopic best response. Figure 14 shows the difference in myopic best response play between Game 2 and Game 3 and Figure 15 shows the difference in myopic best response play between stepping stone games with *complete information* vs *incomplete information*. Across games and levels of information, when players had C as their best response, the strategy corresponding to the Pareto efficient equilibrium, they played their myopic best response 87-91% of the time, which is inline with the literature examining myopic best response in evolutionary experiments [Hwang et al., 2018, Mäs and Nax, 2016, Lim and Neary, 2016]. However, when another strategy was a best response this rate fell dramatically.

This significant drop off appears to be largely driven by players picking C which partially explains the difference in myopic best response play between when A was the myopic best response and when B was the myopic best response. This is because, by formulation, the basin of attraction for A can contain many more players with the strategy C than the basin of attraction for B can.

[22] E_A payoff dominates E_B in Game 3.

I employ a generalized logistic mixed model to examine which factors affect the rate of myopic best response play. I control for the effect that different myopic best responses have on the propensity to play those best responses, as there is a clear difference demonstrated in Figure 14. I also control for the game played and I incorporate clustering at the subject ID level and within the nested experiment number to address potential correlations within the data. With these controls, I test if there is any significant interaction between *incomplete information* and which strategy is the myopic best response, if the difference in payoff between the myopic best response and the next highest myopic payoff (denoted $\Delta\Pi_{mBR}$), and if past stochastic rejection of strategy updating has any effect on myopic best response play.

Figure 13: Residual Effect of Round and Lagged Failed Stochastic Update on Choice = mBR Probability

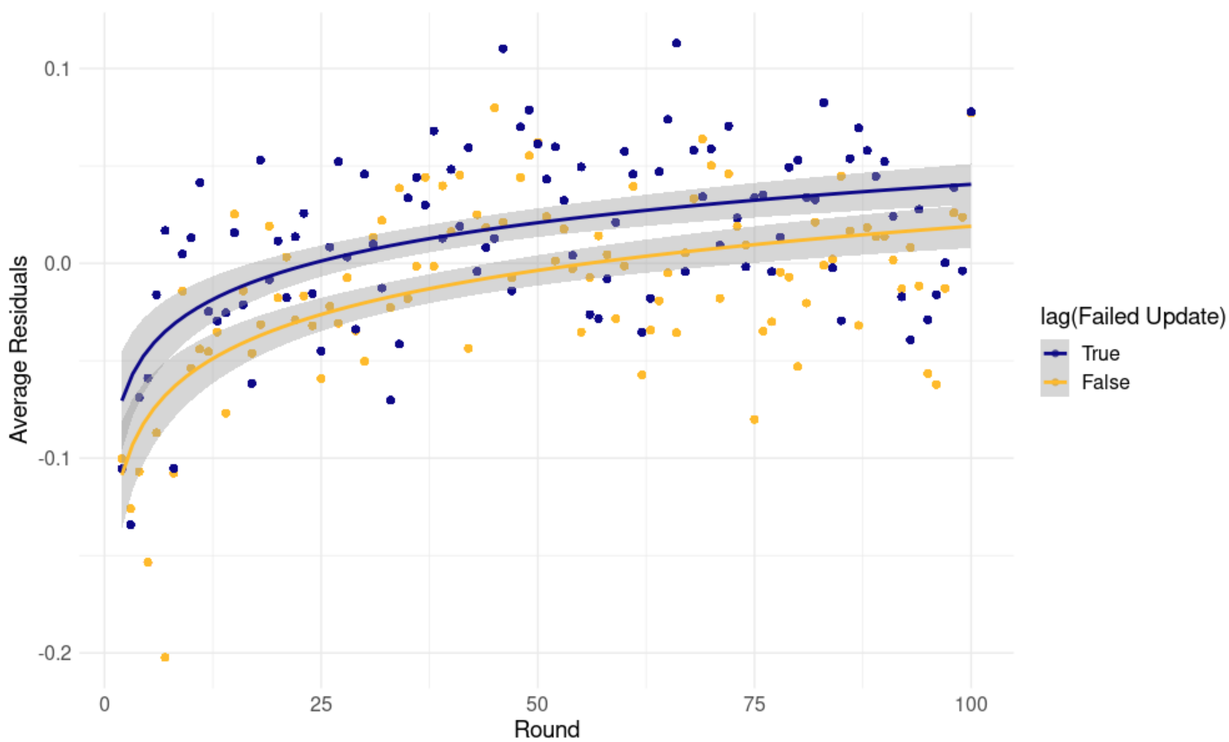

To account for previous findings of Lim and Neary [2016] who found a positive influence of the round number on the rate of myopic best response play, I included it as a control variable. To test the most appropriate way to model the effect of round number on subjects choosing to play their myopic best response I regress upon the above specified model excluding the independent variables round and lagged stochastic rejection. I then use the results of the restricted model to predict the probability that subjects will play a myopic best response and graph the residuals in Figure 13. In Figure 13, the blue points are the average residual for each round when the players choice was rejected in the previous round, the average residual when players' previous choices were accepted are colored goldenrod. There seems to be a non-linear relationship between the round number and residual value. As such i specify a logarithmic relationship between round number and propensity to play myopic best response. With this specification, I regress $log(Round)$ and $lag(FailedUpdate)$ on the residuals of the restricted model to demonstrate the impact of a player's action being stochastically rejected in the previous round on the rate of myopic best response play. The area around the best fit line indicate 95% confidence intervals. Figure 13 clearly shows that players are more likely to choose to play their myopic best response after failed stochastic update in the previous round.

Given the results of the residual test, I include $log(Round)$ in the full logistic model. The results of the regression can be found in table 7.

As discussed, The analysis revealed that stochastically rejection had a significant effect on the odds that the myopic best response was selected in the next round with a point estimate of 0.2602 and a p-value $= 1.24 \times 10^{-6}$. Thus, I rejected hypothesis 5_0 in favor of hypothesis 5_A, indicating that when a player's action was rejected in the previous round, they are more likely to play the myopic best response in the next period.

Figure 14: Rate of Myopic Best Response Play by Strategy Profile Faced

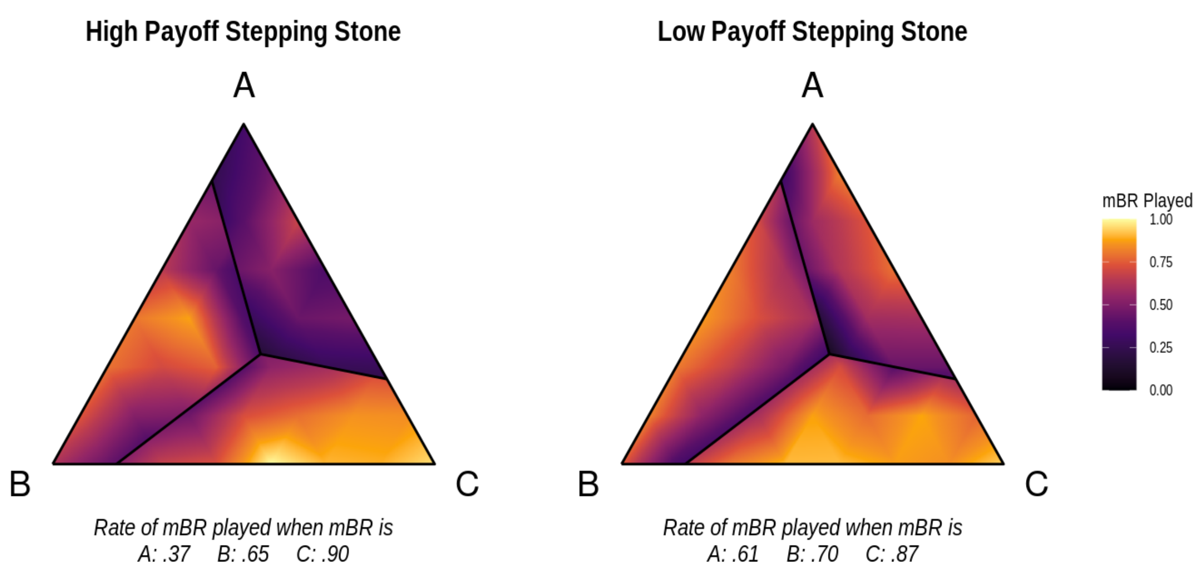

I also examined whether the likelihood of a subject playing the myopic best response is influenced by the difference in expected payoff between the myopic best response and the next highest option. I found a significant positive relationship, with a logistic parameter estimate of 0.1287 and a $p-$value $< 2 \times 10^{-16}$. Therefore, I rejected hypothesis 4_0 in favor of hypothesis 4_A, suggesting that a larger difference in expected payoff leads to a higher probability of myopic best response play. This result in part explains why the mBR plots in Figures 14 and 15 are notably dark around the mBR boundaries.

Figure 15: Rate of Myopic Best Response Play by Strategy Profile Faced

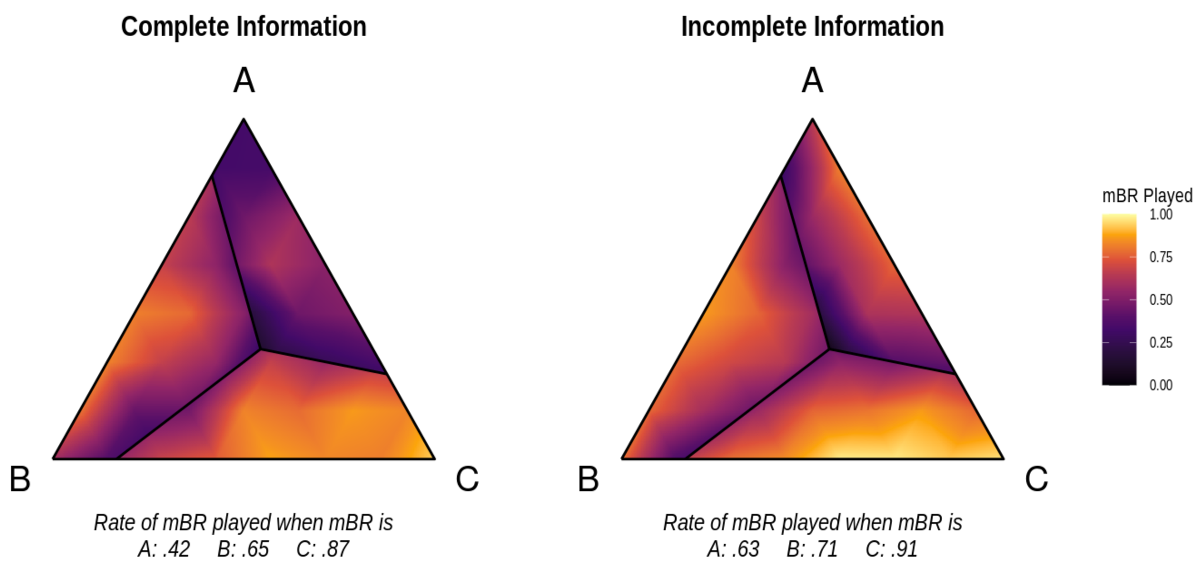

I also examine Hypothesis 7, which explores whether the rate of deviations from the myopic best response, specifically when the myopic best response does not align with the Pareto efficient equilibrium, differs between games of *complete information* and *incomplete information*. To investigate this, I examine the interaction terms of myopic best response and *incomplete information*. The point estimate of 0.7417 with a p-value = 0.000438 indicates that when information is incomplete and the myopic best response is A, subject are much more likely to play A in games with *incomplete information*. However, when the myopic best response is B, the combined point estimate is just 0.199872 and a resulting insignificant 1-sided p-value of 0.232. Although I can not claim that the subjects were statistically more likely to play B when B was the myopic best response when playing with *incomplete information*, the combined results do provide significant support to hypothesis 7_A. As such, I claim that the likelihood of deviating from the myopic best response is significantly higher in games with *incomplete information* when the myopic best response does not correspond to the Pareto efficient equilibrium.

Related to my test of Hypothesis 4, I investigate whether the proportion of "mistakes" made by players choosing action B in sessions playing Game 1 is significantly lower than $1/2$, indicating non-uniform distribution of mistakes. To test this, I employ a simple binomial test.

In the first 100 rounds, there were a total of 606 "mistakes" in sessions playing Game 1. Out of these mistakes, only 67 were players choosing action B. The binomial test yields compelling results. It allows me to reject Hypothesis 6_0 in favor of 6_A, as indicated by a p-value of less than $2.2e^{-16}$.

4 Concluding Remarks

In this paper I use the theoretical foundations of Young [1993] and Ellison [2000] to define *stepping stones*, a recurrent class which reduces the resistance from one recurrent class to another. I then design an experiment to test if injecting a stepping stone into a stag hunt game helps the group transition to the Pareto efficient equilibrium as theoretically predicted. I use a 3×2 treatment design varying the amount of information players receive about the game as well as the game they play and conducted 18 sessions in total, three for each treatment.

The main results are as follows: First, I find that groups that played games with stepping stones were always able to make the transition to the risky, high payoff equilibrium and ended up playing the strategy associated with that equilibrium with the highest frequency. By contrast, groups without a stepping stone occasionally failed to make the transition. I also find that in games where the stepping stone payoff dominated the starting equilibrium, groups were more stable at and ended up playing the Pareto efficient equilibrium significantly more than when the starting equilibrium payoff dominated the stepping stone.

Second, in examining the effect that information about other players' payoffs had on the game, I found that the groups who played with *complete information* were more successful than groups with *incomplete information* when playing the stag hunt game with no stepping stone. I attribute this to the common knowledge that a Pareto improvement exists. However, I find that this effect disappeared when a stepping stone was added to the game, presumably because the stepping stone offered easy to accomplish transitions at a low deviation cost.

Finally, I examine how players made decisions in relation to their myopic best response. examining myopic best response, the decision making mechanism behind adaptive play models. Recent experiments have found that subjects play their myopic best response 90-96% of the time which provides good support for using the adaptive learning model in analyzing evolutionary games. In this experiment I find that subjects played their myopic best response 87-91% of the time when their myopic best response corresponded with the Pareto efficient equilibrium which is in line with what's been observed in the literature. I found several factors that influenced subjects propensity to deviate. Like Lim and Neary [2016], I find that players are more likely to deviate from myopic best response in the initial stages of the game and that players were sensitive to the difference in myopic payoffs. Specifically, players were less likely to play their myopic best response when the difference in myopic payoffs between that and an alternative strategy was lower. I also found that the largest factor in determining if a player would play their myopic best response is if their myopic best response corresponds to the payoff dominant equilibrium, adding support to the theory of Harsanyi et al. [1988] and results of Jagau [2022]. As such, combining the current data with a payoff-dependent mistakes model would give more powerful analysis.

I also adapted the stochastic strategy update probability used in experiments like Hwang et al. [2018] by first soliciting players decisions before the stochastic determination. This pseudo-strategy method allowed me to boost data collection and provided an interesting result. When subjects choices were not accepted in the previous round they were slightly but significantly less likely to deviate from myopic best response in the subsequent round. I attribute this to the increased salience that they could get stuck in an inefficient choice for multiple rounds. Further testing is required to see if this same effect impacts play when the stochastic determination is made prior to the player's decision.

References

J. Bryant. A simple rational expectations keynes-type model. *The Quarterly Journal of Economics*, 98(3): 525–528, 1983.

L. Bursztyn, B. R. Handel, R. Jimenez, and C. Roth. When product markets become collective traps: The case of social media. Technical report, National Bureau of Economic Research, 2023.

D. Canning. Rationality, computability, and nash equilibrium. *Econometrica*, 60(4):877–888, 1992. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/2951570.

D. L. Chen, M. Schonger, and C. Wickens. otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97, 2016.

Congressional Budget Office. Estimated budgetary effects of H.R. 5376, the inflation reduction act of 2022, as amended in the nature of a substitute. Technical Report ERN22335, Congressional Budget Office, 2022. URL https://www.cbo.gov/system/files/2022-08/hr5376_IR_Act_8-3-22.pdf.

R. Cooper and A. John. Coordinating coordination failures in keynesian models. *The Quarterly Journal of Economics*, 103(3):441–463, 1988.

R. W. Cooper, D. V. DeJong, R. Forsythe, and T. W. Ross. Selection criteria in coordination games: Some experimental results. *The American Economic Review*, 80(1):218–233, 1990.

C. Darwin. *On the Origin of Species by Means of Natural Selection*. Murray, 1859.

G. Devetag and A. Ortmann. When and why? a critical survey on coordination failure in the laboratory. *Experimental economics*, 10(3):331–344, 2007.

A. Di Girolamo and M. Drouvelis. The role of gender composition and size of the group in a minimum effort game. *Economics Letters*, 137:168–170, 2015. ISSN 0165-1765. doi: https://doi.org/10.1016/j.econlet.2015.10.040. URL https://www.sciencedirect.com/science/article/pii/S0165176515004553.

G. Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/2951493.

G. Ellison. Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *The Review of Economic Studies*, 67(1):17–45, 2000.

U. Gneezy and J. Potters. An experiment on risk taking and evaluation periods. *The quarterly journal of economics*, 112(2):631–645, 1997.

J. K. Goeree and C. A. Holt. An experimental study of costly coordination. *Games and Economic Behavior*, 51(2):349–364, 2005.

S. Gulesci, S. Jindani, E. La Ferrara, D. Smerdon, M. Sulaiman, and H. Young. A stepping stone approach to norm transitions. Working Paper RWP23-013, Harvard Kennedy School, 2023.

J. C. Harsanyi, R. Selten, et al. A general theory of equilibrium selection in games. *MIT Press Books*, 1, 1988.

C. Harto. Electric vehicle ownership costs: Today's electric vehicles offer big savings for consumers. *Consumer Reports*, 2020.

C. Hoffman. Tesla's 250-kw supercharger only saved us 2 minutes vs. a 150-kw charger, 2020. URL https://www.caranddriver.com/news/a32132062/tesla-250-kw-vs-150-kw-supercharger-tested/.

S.-H. Hwang, W. Lim, P. Neary, and J. Newton. Conventional contracts, intentional behavior and logit choice: Equality without symmetry. *Games and Economic Behavior*, 110:273–294, 2018.

S. Jagau. To catch a stag: Payoff- and risk-dominance in a one-shot coordination experiment. *Working Paper*, 2022.

M. Kandori, G. J. Mailath, and R. Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56, 1993. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/2951777.

M. Knez and C. Camerer. Creating expectational assets in the laboratory: coordination in 'weakest-link'games. *Strategic Management Journal*, 15(S1):101–119, 1994.

E. Kohlberg and J.-F. Mertens. On the strategic stability of equilibria. *Econometrica: Journal of the Econometric Society*, pages 1003–1037, 1986.

W. Lim and P. R. Neary. An experimental investigation of stochastic adjustment dynamics. *Games and Economic Behavior*, 100:208–219, 2016.

I. Malmgren. Quantifying the societal benefits of electric vehicles. *World Electric Vehicle Journal*, 8(4): 996–1007, 2016.

M. Mäs and H. H. Nax. A behavioral study of “noise” in coordination games. *Journal of Economic Theory*, 162:195–208, 2016.

R. Nagel. Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5): 1313–1326, 1995. ISSN 00028282. URL http://www.jstor.org/stable/2950991.

P. R. Neary. Competing conventions. *Games and Economic Behavior*, 76(1):301–328, 2012. ISSN 0899-8256. doi: https://doi.org/10.1016/j.geb.2012.06.003. URL https://www.sciencedirect.com/science/article/pii/S0899825612000875.

D. Romer. Advanced macroeconomics. mcgraw, 1996.

W. Samuelson and R. Zeckhauser. Status quo bias in decision making. *Journal of risk and uncertainty*, 1: 7–59, 1988.

R. Selten. Die strategiemethode zur erforschung des eingeschränkt rationalen verhaltens im rahmen eines oligopolexperimentes. Seminar für Mathemat. Wirtschaftsforschung u. Ökonometrie, 1967.

R. M. Sheremeta. Experimental comparison of multi-stage and one-stage contests. *Games and Economic Behavior*, 68(2):731–747, 2010.

D. O. Stahl and P. W. Wilson. On players’ models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254, 1995.

R. H. Thaler and C. R. Sunstein. *Nudge*. Yale University Press, New Haven, CT and London, 2008. ISBN 978-0-300-12223-7.

A. Tversky and D. Kahneman. Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics*, 106(4):1039–1061, 1991.

United Auto Workers. Taking the high road: Strategies for a fair ev future. Technical report, United Auto Workers, 2020. URL https://region1d.uaw.org/system/files/ev-white-paper-revised-january-2020-final.pdf.

J. B. Van Huyck, R. C. Battalio, and R. O. Beil. Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, 80(1):234–248, 1990.

A. R. Wallace. On the tendency of varieties to depart indefinitely from the original type. *Journal of the Proceedings of the Linnean Society of London. Zoology*, 3(9):53–62, 1858.

H. P. Young. The evolution of conventions. *Econometrica: Journal of the Econometric Society*, pages 57–84, 1993.

H. P. Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, 1998. ISBN 9780691086873. URL http://www.jstor.org/stable/j.ctv10h9d35.

5 Appendix

Figure 16: No Stepping Stone & Complete Information Time Series of Group Strategy

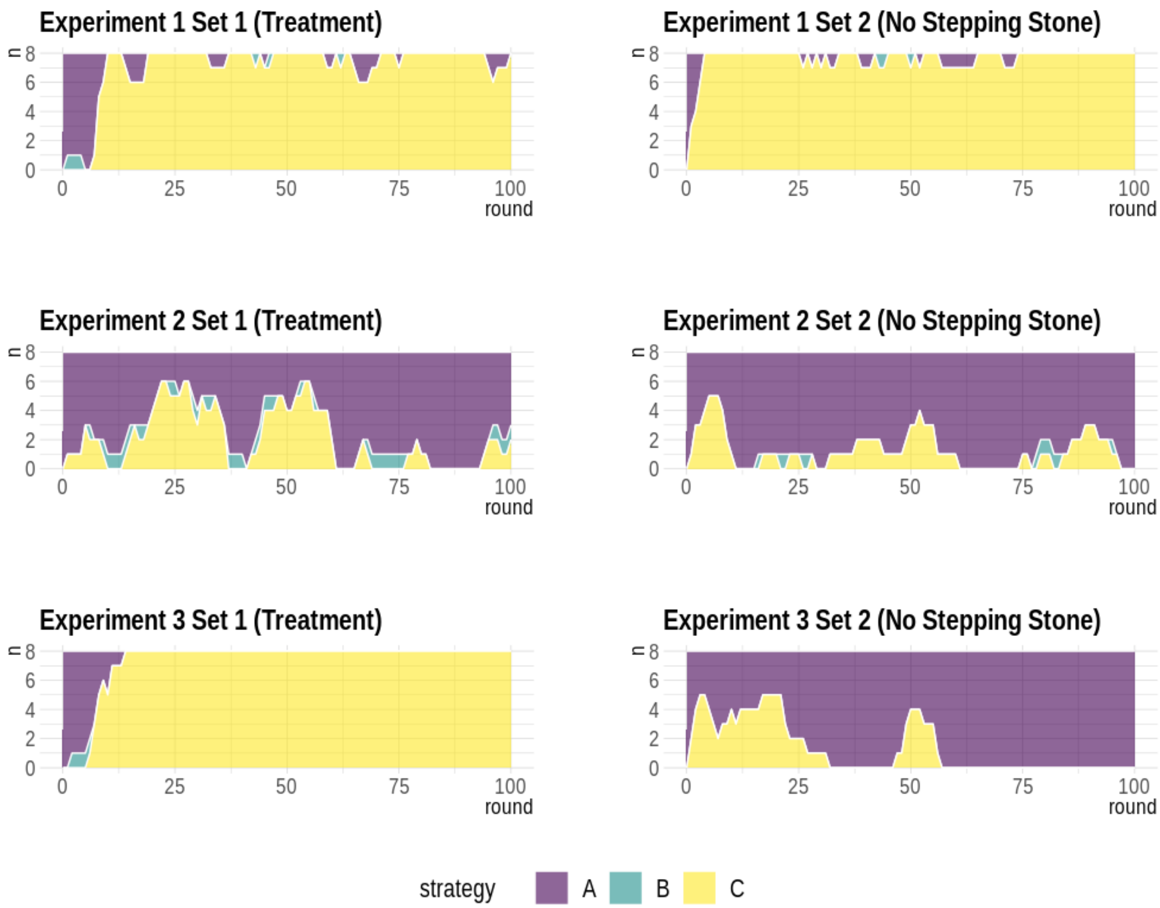

Figure 17: No Stepping Stone & Incomplete Information Time Series of Group Strategy

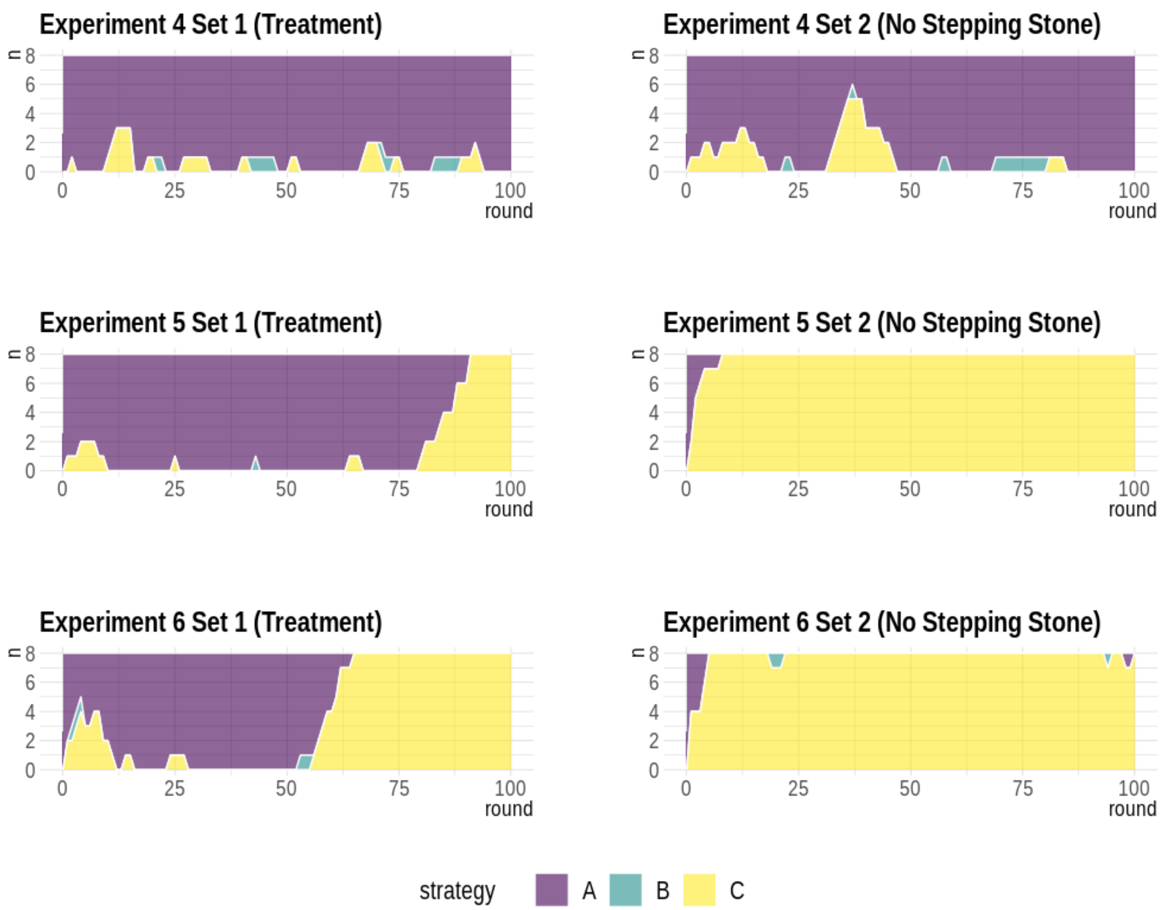

Figure 18: High Payoff Stepping Stone & Complete Information Time Series of Group Strategy

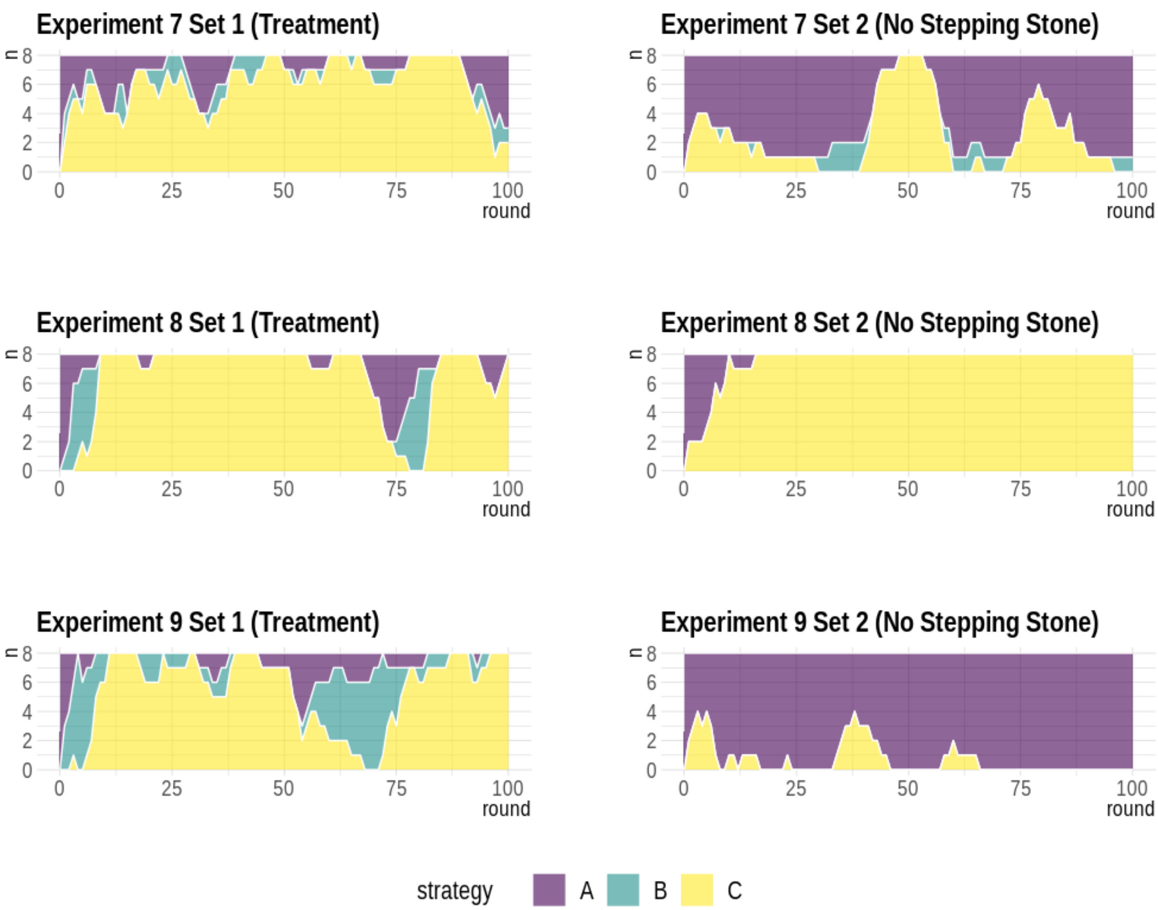

Figure 19: High Payoff Stepping Stone & Incomplete Information Time Series of Group Strategy

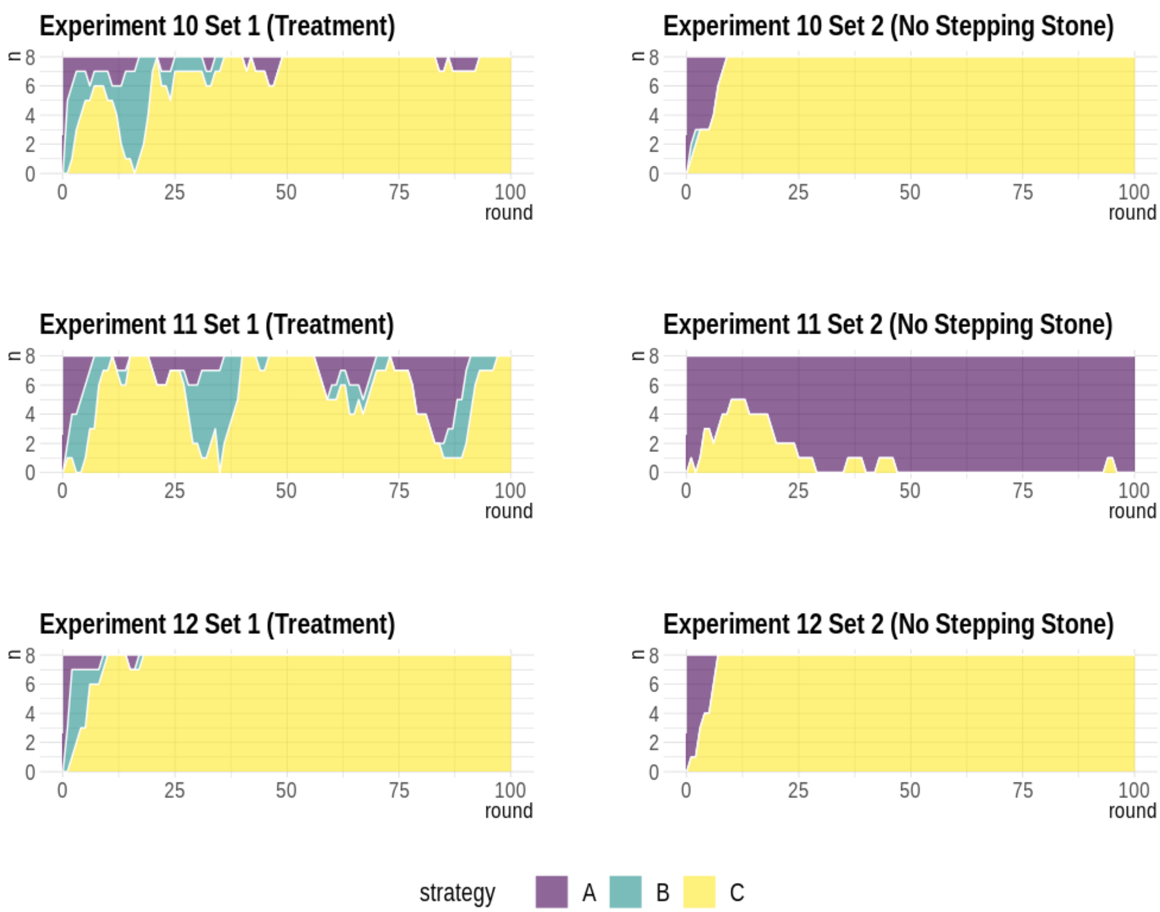

Figure 20: Low Payoff Stepping Stone & Complete Information Time Series of Group Strategy

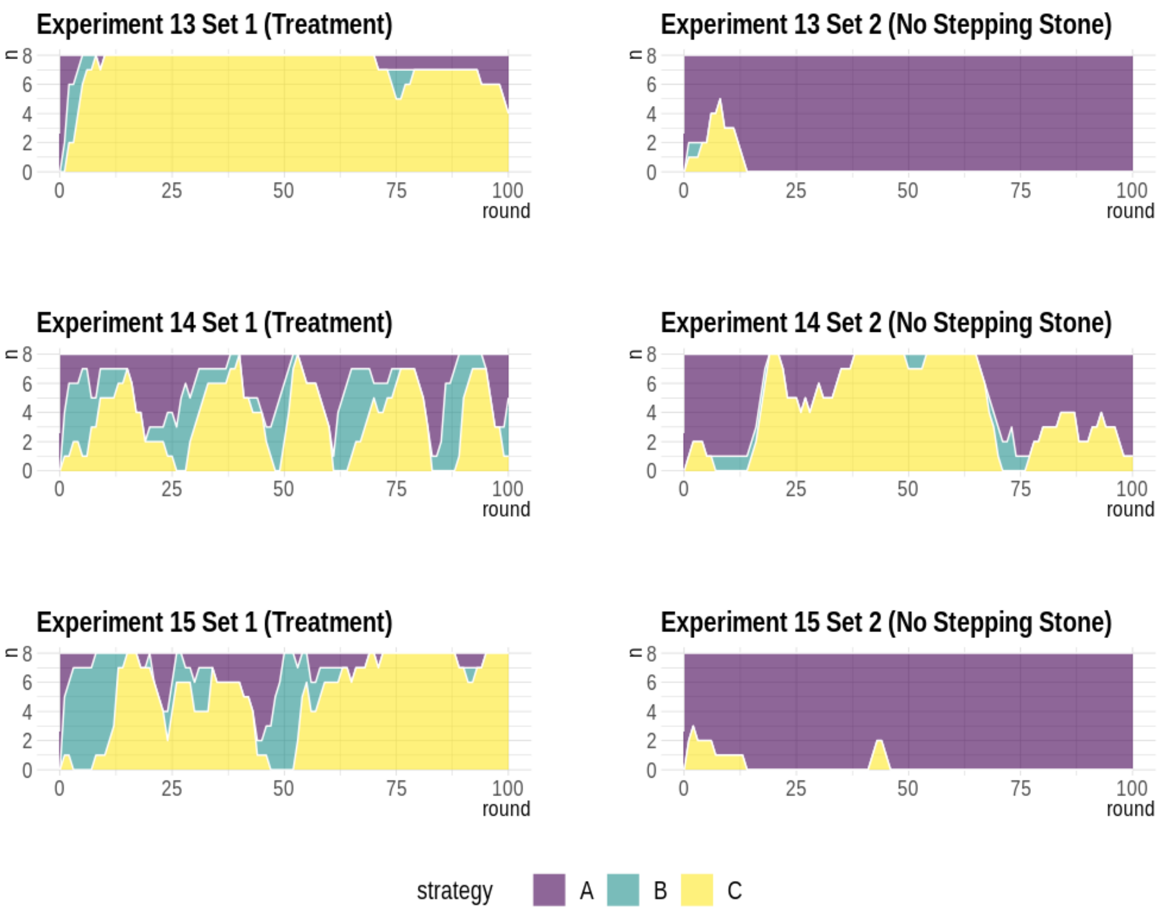

Figure 21: Low Payoff Stepping Stone & Incomplete Information Time Series of Group Strategy

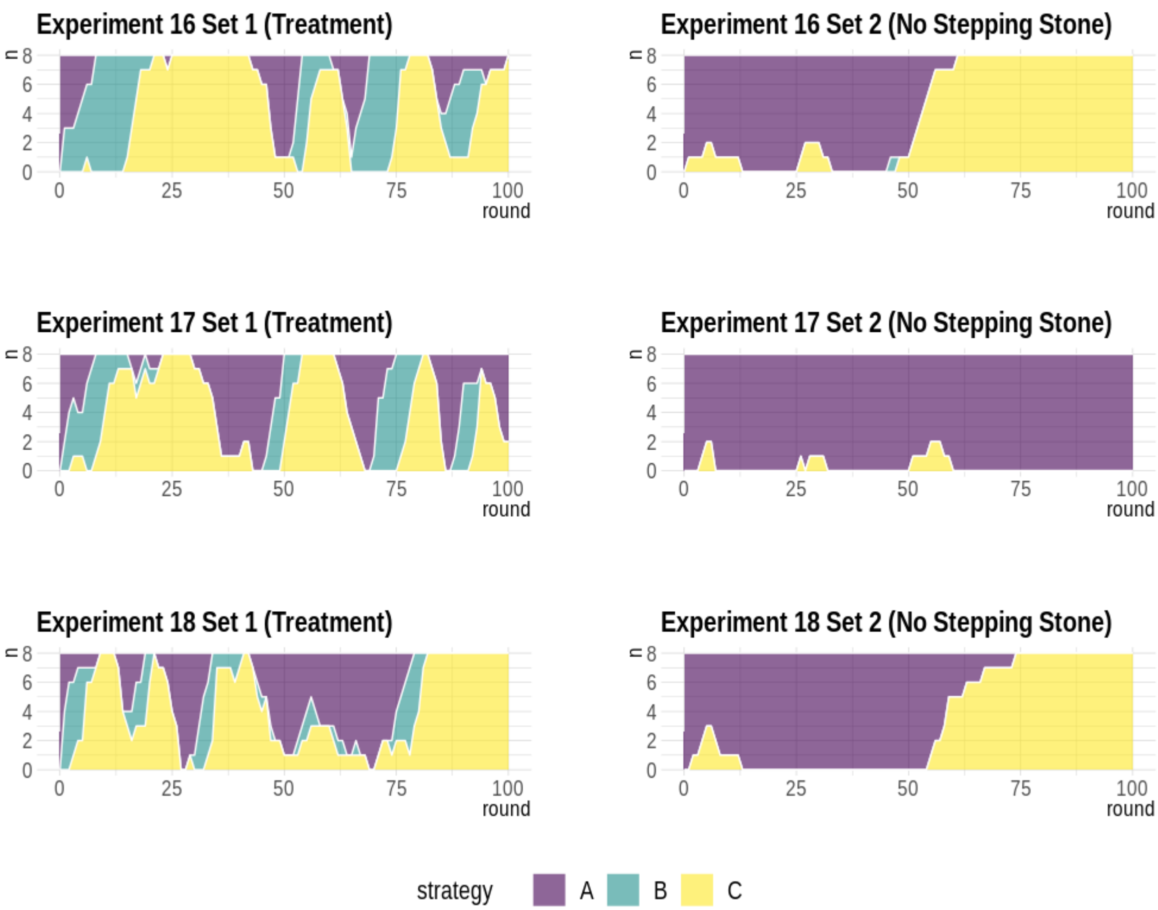

Table 1: Proportion of Strategy A, B, and C Played in Each Session

Game	Information	#	Set 1: Treatment			Set 2: Effects		
			A	B	C	A	B	C
No Stepping Stone	I	1	0.924	0.021	0.055	0.879	0.021	0.1
		2	0.828	0.001	0.171	0.019	0	0.981
		3	0.537	0.007	0.455	0.02	0.005	0.975
	C	1	0.11	0.009	0.881	0.039	0.004	0.958
		2	0.696	0.051	0.252	0.835	0.015	0.15
		3	0.072	0.006	0.921	0.843	0	0.158
High Payoff Stepping Stone	I	1	0.056	0.105	0.839	0.041	0.003	0.956
		2	0.159	0.169	0.672	0.89	0	0.11
		3	0.018	0.034	0.949	0.036	0	0.964
	C	1	0.17	0.075	0.755	0.649	0.051	0.3
		2	0.115	0.086	0.799	0.056	0	0.944
		3	0.106	0.219	0.675	0.92	0	0.08
Low Payoff Stepping Stone	I	1	0.166	0.295	0.539	0.494	0.003	0.504
		2	0.325	0.204	0.471	0.972	0	0.028
		3	0.336	0.138	0.526	0.583	0	0.418
	C	1	0.065	0.03	0.905	0.956	0.004	0.04
		2	0.298	0.248	0.455	0.415	0.038	0.548
		3	0.135	0.205	0.66	0.968	0	0.032

Table 2: Mean Proportion of Strategy A, B, and C Played in Each Treatment

Game	Information	Set 1: Treatment			Set 2: Effects		
		A	B	C	A	B	C
No Stepping	I	0.763	0.010	0.227	0.306	0.009	0.685
Stone	C	0.293	0.22	0.685	0.572	0.006	0.422
High Payoff	I	0.078	0.103	0.82	0.322	0.001	0.677
Stepping Stone	C	0.13	0.127	0.743	0.542	0.017	0.441
Low Payoff	I	0.276	0.212	0.512	0.683	0.001	0.316
Stepping Stone	C	0.166	0.161	0.673	0.78	0.014	0.207

Table 3: Choices in each Session by MBR

Game	Information	#	mBR	Set1: Treatment n	A	B	C	Set 2: Effects n	A	B	C
No Stepping Stone	I	1	A	800	.93	.02	.05	798	.9	.01	.09
			C	0	-	-	-	2	.5	0	.5
		2	A	722	.89	0	.11	30	.13	0	.87
			C	78	.03	0	0	770	0	0	1
		3	A	496	.86	.01	.13	38	.18	0	.82
			C	304	0	0	1	762	0	.01	.99
	C	1	A	126	.4	.01	.59	30	.23	0	.77
			C	674	.05	0	.95	770	.02	.01	.97
		2	A	782	.7	.05	.25	800	.84	.02	.14
			C	18	.67	0	.33	0	-	-	-
		3	A	593	.59	.03	.37	800	.82	0	.18
			C	714	0	0	1	0	-	-	-
High Payoff Stepping Stone	I	1	A	34	.18	.35	.47	62	.4	.02	.58
			B	56	.02	.66	.32	-	-	-	-
			C	710	.04	.03	.93	738	0	0	1
		2	A	152	.38	.11	.51	800	.91	0	.09
			B	102	.08	.67	.25	-	-	-	-
			C	546	.08	.07	.85	0	-	-	-
		3	A	8	.25	.62	.12	54	.43	0	.57
			B	18	.11	.56	.33	-	-	-	-
			C	774	.01	.01	.98	746	0	0	1
	C	1	A	210	.29	.16	.55	698	.73	.06	.21
			B	21	.29	.19	.52	-	-	-	-
			C	569	.11	.07	.82	102	.11	0	.89
		2	A	100	.52	.16	.32	76	.42	0	.58
			B	82	.07	.68	.24	-	-	-	-
			C	674	.04	0	.96	724	0	0	1
		3	A	62	.5	.19	.31	800	.92	0	.07
			B	138	.14	.69	.17	-	-	-	-
			C	600	.06	.1	.84	0	-	-	-
Low Payoff Stepping Stone	I	1	A	116	.63	.28	.09	446	.87	0	.13
			B	256	.04	.79	.17	-	-	-	-
			C	428	.1	.02	.88	354	.01	0	.99
		2	A	268	.76	.13	.11	800	.98	0	.02
			B	175	.06	.74	.19	-	-	-	-
			C	357	.1	.02	.87	0	-	-	-
		3	A	316	.69	.12	.18	528	.84	0	.16
			B	103	.08	.54	.38	-	-	-	-
			C	381	.06	.04	.9	272	.03	0	.97
	C	1	A	48	.21	.1	.69	800	.95	0	.05
			B	18	.17	.56	.28	-	-	-	-
			C	734	.05	.01	.93	0	-	-	-
		2	A	256	.54	.21	.26	498	.64	.04	.32
			B	209	.15	.62	.23	-	-	-	-
			C	335	.21	.07	.72	302	.04	.02	.94
		3	A	124	.38	.16	.46	800	.97	0	.03
			B	144	.1	0	.18	-	-	-	-
			C	532	.09	.04	.87	0	-	-	-

Table 4: Grouped Choices by Treatment

Game	Information	mBR	n	A	B	C
	C	A	994	0.65	0.04	0.3
No Stepping	C	C	1406	0.03	0	0.96
Stone	I	A	2018	0.9	0.01	0.09
	I	C	382	0.01	0	0.99
	C	A	372	0.39	0.16	0.45
	C	B	241	0.13	0.64	0.22
High Payoff	C	C	1787	0.07	0.06	0.87
Stepping Stone	I	A	194	0.34	0.18	0.48
	I	B	176	0.06	0.65	0.28
	I	C	2030	0.04	0.03	0.93
	C	A	428	0.45	0.18	0.36
	C	B	371	0.13	0.65	0.22
Low Payoff	C	C	1601	0.1	0.03	0.87
Stepping Stone	I	A	700	0.71	0.15	0.14
	I	B	534	0.05	0.73	0.22
	I	C	1166	0.09	0.03	0.88

Table 5: Generalized Logistic Mixed Model: Is C Played More in Game 2 and 3 than Game 1?

Fixed Effects	**Estimate**	**Std. Error**	**z value**	**Pr($>|z|$)**
(Intercept)	-0.3045	0.2230	-1.366	0.17202
Game 2	1.8157	0.3147	5.770	7.94×10^{-9}
Game 3	0.8588	0.3130	2.744	0.00608
Model Information				
AIC: 14738.0				
BIC: 14775.8				
Log Likelihood: -7364.0				
Deviance: 14728.0				
Residual degrees of freedom: 14395				

Table 6: Hypothesis 3 Test: Generalized Logistic Mixed Model Results

Data filtered to only include observations where mBR $= A$

Fixed Effects	**Estimate**	**Std. Error**	**z value**	**Pr($>|z|$)**
(Intercept)	-1.7481	0.2022	-8.645	$< 2 \times 10^{-16}$
Game 3	-0.1030	0.2645	-0.389	0.697
Model Information				
AIC: 1448.1				
BIC: 1464.4				
Log Likelihood: -721.1				
Deviance: 1442.1				
Residual degrees of freedom: 1691				

Table 7: Hypotheses 4, 5 and 7 Test: Generalized Logistic Mixed Model

mBR is the Dependant Variable				
Fixed Effects	**Estimate**	**Std. Error**	**z value**	**Pr(> $\|z\|$)**
(Intercept)	-2.359739	0.244356	-9.657	$< 2 \times 10^{-16}$
mBR = B	1.200690	0.127741	9.399	$< 2 \times 10^{-16}$
mBR = C	2.239175	0.097942	22.862	$< 2 \times 10^{-16}$
Incomplete Info	0.741663	0.210947	3.516	0.000438
Game 2	0.097101	0.247757	0.392	0.695116
Game 3	0.168979	0.245713	0.688	0.491637
lag(stochastic rejection)	0.260193	0.053658	4.849	1.24×10^{-6}
$\Delta\Pi_{mBR}$	0.110395	0.004179	26.416	$< 2 \times 10^{-16}$
log(round)	0.290226	0.029797	9.740	$< 2 \times 10^{-16}$
mBR = B:Incomplete Info	-0.541791	0.173655	-3.120	0.001809
mBR = C:Incomplete Info	-0.582463	0.141303	-4.122	3.75×10^{-5}
Model Information				
AIC: 9536.0				
BIC: 9634.3				
Log Likelihood: -4755.0				
Deviance: 9510.0				
Residual degrees of freedom: 14243				

Figure 22: Experiment Instructions (Complete Information)

Introduction

Instructions

In this study, you are playing in a group of size 8 and you will play 2 games, each of which will consist of you making a series of 100 decisions: choosing one of **A**, **B**, or **C**. Each decision will yield a different payoff for you which will depend on the decisions of the other players in the game. After playing both games, the experiment will end and the computer will randomly select 2 rounds from each game. Your **payoffs** in those 4 randomly selected rounds will be added together and divided by 2 * 7. That amount, rounded to the nearest dollar, will be added to your show up fee of $5 and together will be your total payment for the study. Because **payoff** rounds are randomly selected, every decision you make has a chance of directly impacting your payment.

How **Strategies** are Selected

During each round you will be able to make or update your decision **choice** by clicking one of the three purple buttons corresponding to the three decisions available to select from (see table below). Your decision **choice** and its row in your payoff table will be highlighted yellow to help you keep track of what decision you have currently selected. At the end of every round, the computer will randomly (with probability = 0.5) select to update your **strategy** from last round with your selected **choice** from this round. If your **strategy** fails to update, the computer will ignore your **choice** from this round and keep your **strategy** from your last round of play. If you do not make a **choice** during a given round, your **strategy** last round will be selected as your **choice** for this round. Your **strategy** for this round will be highlighted green at the start of the next round.

Note: **Everyone** will start the game with **A** being their "last round **strategy**", and everyone randomly and independently faces the same 0.5 update probability.

Game 1:

*In each cell, the amount to the left in red is the **payoff** for you and to the right in blue is the **payoff** for the other participants.*

	The Other Participants		
	A	**B**	**C**
I will pick A	7 points, 7 points	1 point, 1 point	7 points, 1 point
I will pick B	1 point, 1 point	1 point, 1 point	1 point, 1 point
I will pick C	1 point, 7 points	1 point, 1 point	9 points, 9 points

How **Payoffs** are Calculated

Payoffs for each round will be calculated based on your **strategy** (not **choice**, see above) and the **strategies** (not **choices**, see above) of all the other players in the game. Each round you get payoffs by pairing up with every player in the game. Your **payoff** in a pairing is determined by your **strategy** which determines the row of the payoff table, and the other player's **strategy** which determines the column of the payoff table, together selecting an individual cell in the payoff table. The red value in each cell is your **payoff** if that cell is selected, and the blue value is the **payoff** the other player gets. The sum of the **payoffs** you get from every player in the game is your total **payoff** for a round.

For example, say there are 3 other players in the game.

If for a round your **strategy** is **A** and the three other players **strategies** are **C** then your **payoff** for that round will be: **0*7 points** + **0*1 point** + **3*7 points** = **21 points**.

If your **strategy** is **B** and the **strategies** of the other players are: **1 A**, **0 B**, **2 C**. Then your **payoff** for that round will be: **1*1 point** + **0*1 point** + **2*1 point** = **3 points**.

If your **strategy** was instead **C** and the **strategies** of the other players are: **1 A**, **0 B**, **2 C**. Then your **payoff** for that round would be: **1*1 point** + **0*1 point** + **2*9 points** = **19 points**.

How Rounds Progress

Each round lasts a fixed amount of time. Starting with 60 seconds for the first two rounds, 45 seconds for rounds 3-5, 30 seconds for rounds 6-8, 20 seconds for rounds 9-11, and 10 seconds for rounds 12-100. After each round, **strategies** and **payoffs** for all players are calculated. The information about your decision **choice** and **strategy** selected last round, the **strategies** of the other players, and your **payoff** from last round are automatically reported at the start of every new round. New rounds will automatically start as soon as the previous round finishes.

Next

Figure 23: Quiz Question 1

Quiz to Check Understanding

Consider the following payoff table:

	The Other Participants		
	A	**B**	**C**
I will pick A	7 points, 7 points	1 point, 1 point	7 points, 1 point
I will pick B	1 point, 1 point	1 point, 1 point	1 point, 1 point
I will pick C	1 point, 7 points	1 point, 1 point	9 points, 9 points

If your strategy in a round is A and the strategy of another player is C what payoff do you get for playing against them that round?

Only integers are accepted.

Submit Answer

Figure 24: Quiz Question 2 (Complete Information Only)

Quiz to Check Understanding

Consider the following payoff table:

	The Other Participants		
	A	**B**	**C**
I will pick A	7 points, 7 points	1 point, 1 point	7 points, 1 point
I will pick B	1 point, 1 point	1 point, 1 point	1 point, 1 point
I will pick C	1 point, 7 points	1 point, 1 point	9 points, 9 points

If your strategy in a round is C and the strategy of another player is B what payoff do they get for playing against you that round?

Only integers are accepted.

Submit Answer

Figure 25: Quiz Question 3

Quiz to Check Understanding

Consider the following payoff table:

	The Other Participants		
	A	**B**	**C**
I will pick A	7 points, 7 points	1 point, 1 point	7 points, 1 point
I will pick B	1 point, 1 point	1 point, 1 point	1 point, 1 point
I will pick C	1 point, 7 points	1 point, 1 point	9 points, 9 points

Assume your **strategy** last round was **B**, your **choice** this round is **A**, and the **strategy** of another player this round is **A**. What is the smallest **payoff** you could get for playing againt them this round?

Only integers are accepted.

Submit Answer

Figure 26: Quiz Question 4

Quiz to Check Understanding

Consider the following payoff table:

	The Other Participants		
	A	**B**	**C**
I will pick A	7 points, 7 points	1 point, 1 point	7 points, 1 point
I will pick B	1 point, 1 point	1 point, 1 point	1 point, 1 point
I will pick C	1 point, 7 points	1 point, 1 point	9 points, 9 points

Assume your **strategy** last round was **B**, your **choice** this round is **A**, the **strategy** of another player last round is **B**, and their **choice** this round is **C**. What is the largest **payoff** you could get for playing againt them this round?

Only integers are accepted.

Submit Answer

Figure 27: Quiz Question 5

Quiz to Check Understanding

Consider the following payoff table:

	The Other Participants		
	A	B	C
I will pick A	7 points, 7 points	1 point, 1 point	7 points, 1 point
I will pick B	1 point, 1 point	1 point, 1 point	1 point, 1 point
I will pick C	1 point, 7 points	1 point, 1 point	9 points, 9 points

Assume there are 7 players other than you in the game, your **strategy** in a round is **A** and the distribution of opponents **strategies** are as follows: **2 A**, **4 B**, **1 C**. What would your **payoff** for this round be?

Only integers are accepted.

Submit Answer

Figure 28: Experiment UI

Your selected choice will be submitted in 7 seconds

Your choice last round, A, was randomly **accepted** and your strategy was updated. Consequently, your strategy for last round is **A**.

Your payoff from the last round: **3*7 points + 0*1 point + 4*7 points = 49 points**

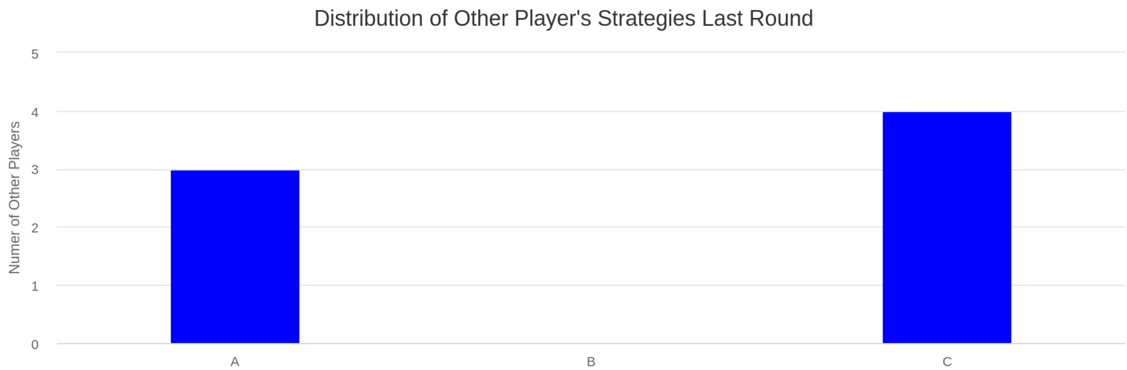

	The Other Participants		
	A (3)	**B (0)**	**C (4)**
I will pick A	7 points, 7 points	1 point, 1 point	7 points, 1 point
I will pick B	1 point, 1 point	1 point, 1 point	1 point, 1 point
I will pick C	1 point, 7 points	1 point, 1 point	9 points, 9 points

Figure 29: Experiment UI

Your selected **choice** will be submitted in 12 seconds

Your **choice** last round, **C**, was randomly **rejected** and your **strategy** failed to update. Consequently, your **strategy** for last round is **A**.

Your **payoff** from the last round: **3*****7 points** + **0*****1 point** + **4*****7 points** = **49 points**

	The Other Participants		
	A (3)	**B (0)**	**C (4)**
I will pick A	7 points, 7 points	1 point, 1 point	7 points, 1 point
I will pick B	1 point, 1 point	1 point, 1 point	1 point, 1 point
I will pick C	1 point, 7 points	1 point, 1 point	9 points, 9 points