When Less Information Is Good Enough: Experiments with Global Stag Hunt Games*

John Van Huyck,[‡] Ajalavat Viriyavipart,[§] and Alexander L. Brown[¶]

Abstract

There is mixed evidence on whether subjects coordinate on the efficient equilibrium in experimental stag hunt games under complete information. A design that generates an anomalously high level of coordination, Rankin et al. (2000), varies payoffs each period in repeated play rather than holding them constant. These payoff "perturbations" are eerily similar to those used to motivate the theory of global games, except the theory operates under incomplete information. Interestingly, that equilibrium selection concept is known to coincide with risk dominance, rather than payoff dominance. Thus, in theory, a small change in experimental design should produce a different equilibrium outcome. We examine this prediction in two treatments. In one, we use public signals to match Rankin et al.'s design; in the other, we use private signals to match the canonical example of global games theory. We find little difference between treatments, in both cases, subject play approaches payoff dominance. Our literature review reveals this result may have more to do with the idiosyncrasies of our complete information framework than the superiority of payoff dominance as an equilibrium selection principle.

JEL Classification: C72, C92, C73, D82

Keywords: Stag Hunt, Global Games, Efficiency, Equilibrium Selection, Threshold Strategies, Risk Dominance, Payoff Dominance, Experiments.

---

*Financial support was provided by the Texas A&M Humanities and Social Science Enhancement of Research Capacity Program. The research was conducted under TAMU IRB approval IRB2012-0664. We thank Ravi Hanumara for his help on z-Tree programming and the Economic Research Laboratory group at Texas A&M for testing the program. We also thank Catherine C. Eckel, Daniel Fragiadakis, Nikos Nikiforakis, and the seminar participants at NYU-Abu Dhabi, American University of Sharjah, Thammasat University, Texas A&M University, the workshop in honor of John Van Huyck, the 2015 North American Economic Science Association meetings, the 2014 Texas Experimental Association Symposium, the 2013 North American Economic Science Association meetings, and the 2013 Southern Economic Association Conference for valuable comments and discussion.

[†] Department of Economics, Texas A&M University, College Station, TX

[‡] John Van Huyck is deceased.

[§] Department of Economics, American University of Sharjah, Sharjah, UAE

[¶] Department of Economics, Texas A&M University, College Station, TX

1 Introduction

The stag hunt game provides a tradeoff between two pure-strategy equilibria: an efficient, risky equilibrium and an inefficient, safe equilibrium. The risks involved are strategic. The strategy associated with the safe equilibrium ensures a constant payoff. The risky strategy yields reduced payoffs if others do not play the same strategy. Thus, the choice between equilibrium depends on players' beliefs (and higher order beliefs) about how individuals naturally coordinate. While philosophical implications of this choice are quite profound, frustratingly, methods in basic game theory cannot select which of these two pure-strategy equilibria is inherently correct.

A great deal of work in game theory and experimental economics has addressed the issue of equilibrium selection in this specific game under complete information. Early experiments find repeated, randomly-rematched, play converges to the inefficient equilibrium (Cooper et al., 1990, 1992). Later research has less absolute findings; structural differences across experiments appear to determine which equilibrium is dominant (see Devetag and Ortmann, 2007). The most glaring exception, Rankin et al. (2000), achieves near universal coordination on the efficient equilibrium by slightly varying (i.e., perturbing) payoffs each round, in contrast to the convention of keeping payoffs constant for several rounds.

Perturbed payoffs are also utilized in theoretical models, but generally to justify coordination *away from* the efficient equilibrium. Carlsson and van Damme (1993a,b) introduce private signals about payoffs in the stag hunt game. Agents best respond by playing the safe strategy up until a sufficiently high signal after which they play a risky strategy. Their "global games" approach determines a unique threshold that is the sole equilibrium that survives removal of iterated dominated strategies. With two players, this threshold is equivalent to the risk dominance solution (Harsanyi and Selten, 1988).

How can perturbed payoffs have such starkly different implications in theory and experiments? The distinction between identical signals that are public or private—while structurally very similar—may determine radically different equilibrium outcomes. In such case, the design of Rankin et al. is perfectly suited to test the predictions of global games in the stag hunt. We can greatly alter game theoretic predictions without changing much else. Further, the Rankin et al. design provides a case known to produce coordination on the efficient equilibrium, so any movement to the global-games prediction under incomplete information is easily observable. To that end, we conduct an experiment where subjects play a sequence of perturbed stag hunt games either with complete or

Tab. 1: A Class of Stag Hunt Games

	A	B
A	$1, 1$	$0, Q$
B	$Q, 0$	Q, Q

incomplete information. Under the complete information treatment—a setting nearly identical to Rankin et al.—subjects coordinate on the efficient equilibrium. Under the incomplete information—a setting designed to match the motivating example of the theory of global games —subjects achieve the efficient equilibrium almost as often, inconsistent with the global games prediction.

Our paper, Rankin et al., and Carlsson and van Damme (1993a,b) all rely on the notion of a general class of stag hunt game forms, depicted in Table 1.[1] When $Q \in [0, 1]$, the game has two strict equilibria: either both players choose A or both players choose B. Both equilibrium selection criteria and human subject behavior can be generally classified by threshold strategies, that is, choosing A if $Q < Q^*$ and B if $Q \geq Q^*$ for some $Q^* \in [0, 1]$. Any strategy profile where both players use threshold strategies with the same Q^* will constitute an equilibrium because they will always end up playing (A, A) when $Q < Q^*$ and (B, B) when $Q \geq Q^*$. In fact, risk dominance (Harsanyi and Selten, 1988) selects (A, A) when $Q < 0.5$ and (B, B) when $Q > 0.5$.

Our complete information treatment replicates the findings of Rankin et al. (2000). Subjects utilize thresholds only on the top quartile of signals (i.e., $Q^* \in (0.75, 1]$), far higher than the predictions of risk dominance. This results in subjects playing choice A with high probability.

Our incomplete information treatment matches the environment of the motivating example of Carlsson and van Damme (1993a,b). They consider an environment where each agent observes a private and noisy signal about Q, the global games parameter, with the true value of Q remaining unknown. In this global stag hunt game, their theory uses iterated elimination of dominated strategies to pin down a unique dominance solvable equilibrium. The unique threshold equilibrium is also $Q^* = 0.5$. In contrast, we find most subjects utilize thresholds that are far higher than the predicted global games threshold. The cohort-level differences between treatments are not statistically meaningful.

In sum, there is little evidence in our experiment to support the predictions of global games or risk dominance as criteria for equilibrium selection. Subjects appear to coordinate on the efficient

[1] It should be noted that our experiment utilizes a linear projection of this game where payoffs are mapped from 0 and 1 to 100 and 500 while the experiments of Rankin et al. use w and $w + 370$ for $w \in [0, 50]$ where the value of w is changed in each period. Carlsson and van Damme (1993a) use this exact game but Carlsson and van Damme (1993b) scale this game by a factor of 4. The scales are theoretically equivalent and produce the same results. We normalize and use the [0,1] scale for ease of reading and interpretation.

equilibrium immediately and do not deviate from it. Further, the effect of introducing private signals—in the spirit of Carlsson and van Damme—to the experimental setup of Rankin et al. appears to be nil.

We caution how far we can extend this conclusion. The level of coordination in stag hunt games under complete information is sensitive to structural changes (Devetag and Ortmann, 2007). The results under Rankin et al.'s design (our complete information treatment) are quite exceptional for their level of efficient outcomes achieved. Our results may only show that changes between incomplete and complete information treatments in this context are minimal, a result consistent with many global games experiments (e.g., Heinemann et al., 2004). That being said, our findings are still quite relevant. Another experiment that compares stag hunt games under complete vs. incomplete information, Cabrales et al. (2007) concludes the criterion of risk-dominance is predictive in global stag hunt games. That paper uses an alternative study for their complete information baseline (i.e., Battalio et al., 2001) that does not provide the high efficiency results of Rankin et al.. We provide a necessary counterpoint.

The remainder of the paper is organized as follows. The next subsection provides an extensive literature review of stag-hunt experiments, global games models and experiments. Section 2 describes our specific analytical framework. Section 3 lays out the experimental design. Section 4 presents results from the experiments, including debriefing questionnaires. We discuss our results in section 5.

1.1 Related literature

There is a long and storied history of experimental stag hunt games under complete information. Beginning in the early 1990s, experiments featured a large number of periods to examine which of the pure-strategy equilibria would survive many rounds of repeated play. Under matching protocols designed to ensure no subject ever encountered the same subject again, Cooper et al. (1990, 1992) find subject play converged to the inefficient pure-strategy equilibrium. Combined with the results of similar coordination games with a greater number of equilibria (Van Huyck et al., 1990, 1991), the idea that strategic uncertainty doomed players to coordination only on low-payoff equilibria became prevalent. Robustness tests show that this result, while common, was not absolute (see Devetag and Ortmann, 2007, for a survey). Altering payoffs to make the risky option relatively more attractive (Battalio et al., 2001; Clark et al., 2001; Dubois et al., 2012; Schmidt et al., 2003;

Stahl and Van Huyck, 2002), utilizing fixed rather than stranger matching (Clark and Sefton, 2001), including preplay communication (Duffy and Feltovich, 2002, 2006) and informing players about their opponent's relatively high risk tolerance (Büyükboyacı, 2014), all appear to increase coordination on the efficient equilibrium. Pairs who survey as being more patient are more likely to coordinate on the efficient equilibrium (Al-Ubaydli et al., 2013).

The strongest evidence of subjects coordinating on the efficient equilibrium comes from Rankin et al. (2000), the basis for our complete information treatment. They report an experiment in which subjects play a sequence of 75 perturbed stag hunt games under complete information where payoffs, action labels, and game forms are changed in each period. They show that payoff dominance emerges as an equilibrium selection principle. Almost all subjects in the last 15 periods use threshold strategies that are close to 1, a payoff-dominant threshold; even when the value of Q is as large as 0.97, slightly more than half of the subjects select a risky choice. Our findings replicate these results.

The 2×2 stag hunt game is also the basis for the initial theory of global games (Carlsson and van Damme, 1993a). The same authors generalize this framework to a generic class of 2×2 games (Carlsson and van Damme, 1993b). Building upon this work, Morris and Shin (1998, 2003) develop a more universal theory and the most popular application of global games theory, the currency attack game. In this game, usually modeled with more than two players, an individual has two choices: "attack" and "not attack." Not attacking produces a sure payoff of 0. Attacking involves a sure cost of T, but may produce a benefit of Y, the global game parameter, provided enough other players also attack. The hurdle for a successful attack is determined by the non-increasing function $h(Y)$. It is important to realize this hurdle function is the model's major departure from global stag hunt games, as reaching the efficient or inefficient equilibrium—holding player choices constant—is independent from the global games parameter in the stag hunt game. Nonetheless, there are many similarities to the global stag hunt game. Both models feature two pure-strategy equilibria in complete information and a single global games threshold under incomplete information.[2]

Beginning with Heinemann et al. (2004), nearly all experiments on global games focus on some form of the currency attack model, rather than the stag hunt game. Their framework features a game of 15 subjects under two information conditions: complete and incomplete information about

[2] Another minor but important difference is that increasing noise in the currency attack game changes the odds of success as well as the expected payoff. However, in global stag hunt games, the unique equilibrium, coinciding with the risk-dominant threshold, is the same for any noise levels that are arbitrarily small, because the expected payoff does not change with different noise levels.

the value of Y. Subjects in their study play 10 independent situations in each of the 10 periods for a total of 100 situations. They present three main findings. First, estimated mean thresholds follow the comparative statics of the global games solution. Second, subjects select "attack" more often in sessions with complete information than in sessions with incomplete information and they coordinate more often than the theoretical predictions for both information conditions. The qualitative results of our data follow these two main findings even though we use different games with different settings. Third, the variations in estimated thresholds among cohorts are large but are about the same for both information conditions despite the different theoretical predictions: equilibrium multiplicity for complete information and uniqueness for incomplete information. Our results are different from both Heinemann et al. (2004) and the theoretical predictions; we observe a larger standard deviation across cohorts with incomplete information than across cohorts with complete information.

One interesting extension to the currency attack game is Szkup and Trevino (2015a) who experimentally allow subjects choose the precision level at a cost.[3] In their experiment, subjects can select from six precision levels, where more precise signals are more costly. The results show that more than one-third of subjects select the equilibrium precision level, which yields the highest expected payoffs. For coordination, subjects who select signals with more precision coordinate more often than those who select less precise signals. We might view this costly information acquisition as a signal of coordination. In addition, when the precision level is exogenously given to them at no cost, subjects in their study behave similarly to subjects in our study. That is, subjects coordinate more often when observing more precise information. Schotter and Trevino (2017) examine subject response time in this framework.

Other papers model and experimentally test the effects of receiving an additional public signal to a currency attack framework (Cornand, 2006). Several make the game dynamic (Shurchkov, 2013), including endogenous entry (Duffy and Ochs, 2012), and exogenous public signals (Shurchkov, 2016). Kawagoe and Ui (2010) examines how ambiguity differentially affects payoff and risk dominance predictions. It should be noted that these extensions to the currency attack game are generally done to understand and predict empirical macroeconomic phenomena rather than more basic questions on whether individuals naturally coordinate.[4] Nonetheless, several find little difference between incomplete and complete information treatments as we do. Some find coordination on the equilibrium predicted by payoff dominance (rather than risk dominance) as we do. However, given that very

[3] Szkup and Trevino (2015b) study this similar game but with a continuum of players.

[4] An exception, Heinemann et al. (2009) compare the results of the currency attack game with known lotteries to measure the level of strategic uncertainty perceived by subjects within the game.

minor differences in the complete information stag hunt game can produce disparate results with regards to coordination on equilibria, it is unlikely one would be able to confidently infer results about the global stag hunt game from any of these studies using the currency attack game.

To our knowledge, only two experiments examine the global stag hunt game. Brindisi et al. (2014) test differences between static and dynamic global stag hunt games, finding significant differences as their theory predicts. Their static treatment uses a variation of the global stag hunt where the parameter gives information about the risky payoff, rather than the safe as in ours. While it was by no means the point of their paper, similar to our study, they find subjects in sessions with incomplete information choose to coordinate more often than theory would predict.

In a design, in many ways, very similar to ours, Cabrales et al. (2007) compare complete information stag hunt games with incomplete information stag hunt games. In sharp contrast to our results, they find that in both treatments subject play converges to the inefficient equilibrium, which under their parameters is the global games solution and risk-dominant threshold solution.[5] Their complete information treatment does not perturb payoffs like Rankin et al.. This creates two issues which make testing global games predictions more difficult: (i) the change between treatments involves both a change of payoff perturbations and information structures; (ii) the game under complete information was already likely to produce coordination on the inefficient equilibrium. Contrary to our results, the authors find support of the risk-dominant and global games solutions. However, the authors conclude that learning through payoff differences, rather than deducing iterated removal of dominated strategies is responsible for the movement to the global games and risk dominant solution. We find the differences between our results fascinating, and speculate on them in our final section.

2 Analytical Framework

We choose the motivating game of Carlsson and van Damme (1993a,b) to test their global games prediction. We examine the perturbation of payoffs in both complete and incomplete information and calculate threshold equilibria under both. That is, all equilibria are of the form choosing A if $Q < Q^*$ and B if $Q \geq Q^*$ for some $Q^* \in [0, 1]$ (see Table 1). Under complete information, any strategy profile where both players use the same threshold strategy will constitute an equilibrium.

[5] Rather than provide subjects with a continuum of possible global games parameters, their design only uses five discrete states. In all five, under complete information the risk-dominant criterion would select the secure strategy. So the risk-dominant threshold and the global games solution is a boundary condition.

However, under incomplete information, the iterated elimination of dominated strategies forces players to conform to the unique threshold equilibrium, which is not efficient. Given this theoretical prediction, it is unfavorable for subjects in the incomplete information treatment to coordinate on the most efficient threshold.

2.1 Stag Hunt Games under Compete Information

Let us generalize the complete information stag hunt game as shown in Table 1 to $n \geq 2$ players. Each player i simultaneously chooses between A and B. Let k denote the number of players, including i, that choose A. Player i's payoff to A is $\pi(k,n) = \frac{k-1}{n-1} \cdot 1 + \left(1 - \frac{k-1}{n-1}\right) \cdot 0 = \frac{k-1}{n-1}$ for $k \geq 1$ and to B is Q.[6]

Suppose $Q \in (0,1)$, consider the strategy profile in which all n players choose A. Since $k = n$, the payoff to A is 1. Deviating from the strategy profile yields Q, which is less than 1 by assumption. Hence, playing A is a best response to the other $n-1$ players choosing A and by symmetry a strict Nash equilibrium. Consider the strategy profile in which all n players choose B. The payoff to B is always Q. Deviating from the strategy profile yields 0, which is less than Q by assumption. Hence, playing B is a best response to the other $n-1$ players choosing B and by symmetry a strict Nash equilibrium.[7]

2.2 Equilibrium Selection

Between the two pure Nash equilibria when $Q \in (0,1)$, all of the players prefer all-A, which yields them 1, over all-B, which yields them less than 1. The presence of multiple Pareto-ranked equilibria confronts the player with a strategy coordination problem. While the high payoff favors A, strategic uncertainty, which is inherent in the strategy coordination problem, may lead players to choose B instead. Intuitively, if Q is high, then it is more likely that players will choose B. Two equilibrium selection principles payoff and risk dominance select different thresholds.

Given the fact that it is more attractive to choose B when Q is high, we expect a player to play a "threshold" strategy where the player chooses A if $Q < Q^*$ and B if $Q \geq Q^*$ for some $Q^* \in [0,1]$. Any strategy profile where all players use threshold strategies with the same Q^* will constitute an equilibrium because the group will always end up playing all A when $Q < Q^*$ and all B when $Q \geq Q^*$.

[6] For example, suppose $n = 8$ and $k = 3$, $\pi(3,8) = \frac{2}{7}$.

[7] There are also mixed strategy Nash equilibria but no other pure Nash equilibria.

Tab. 2: A Class of Global Stag Hunt Games Observed by Player i

	A	B
A	$1, 1$	$0, Q_i$
B	$Q_i, 0$	Q_i, Q_i

Harsanyi and Selten (1988) introduce payoff dominance which compares the efficiency of equilibria and selects the equilibrium that all players prefer. In the class of stag hunt games, this principle selects the all-A equilibrium regardless of the value of Q, which does not capture the intuitive notion that the likelihood of all A should depend on Q.[8]

Harsanyi and Selten (1988) develop risk dominance which compares the riskiness of equilibria. For $n = 2$, risk dominance is equivalent to choosing the equilibrium with the larger basin of attraction under best response dynamics. That is, in this game, we compare between $(1 - Q)^2$, a product of deviation losses from an equilibrium (A, A), and $(Q - 0)^2$, a product of deviation losses from an equilibrium (B, B). The former is larger when $Q < 0.5$, so the (A, A) equilibrium has the larger basin of attraction, in which case both payoff dominance and risk dominance agree on the selection of A by both players. However, the latter is larger when $Q > 0.5$, so the (B, B) equilibrium has the larger basin of attraction, in which case payoff dominance and risk dominance conflict.

For $n > 2$, Harsanyi and Selten (1988, pp.207-209) use the tracing procedure to select the risk-dominant equilibrium. We follow Proposition 3.1 of Carlsson and van Damme (1993a) to show that Q^* is equal to 0.5, which is the same as the case where $n = 2$, see Appendix for details.

Proposition 1. *For all $n \geq 2$, in a stag hunt game with complete information, there are multiple equilibria as symmetric and monotone strategies. Any thresholds between 0 and 1 are equilibria. Payoff dominance has a threshold of 1 and risk dominance has a threshold of 0.5.*

Proof. See Appendix for the proof of the condition regarding the risk-dominance threshold, all other statements are trivial. □

2.3 Global Stag Hunt Games

Carlsson and van Damme (1993b) develop an equilibrium selection theory based on the idea that the payoff parameters of a game cannot be observed with certainty. In this global stag hunt game, each player observes the payoff parameters with errors as in Table 2, but his/her payoff is determined by the actual payoff parameters in Table 1. With this assumption, there exists a unique equilibrium

[8] Selecting A is not always the payoff-dominant equilibrium if Q can be larger than 1; in this case, B strictly dominates A. Therefore, payoff dominance is a threshold strategy with $Q^* = 1$.

threshold which equals the risk-dominant threshold of the game with complete information. In fact, Carlsson and van Damme (1993b) show that this is true for any 2×2 games in which the initial subclass of games is large enough and contains games with different equilibrium structures.

Following Carlsson and van Damme (1993a, section 4), we assume that everything about the stag hunt game is common knowledge except the payoff to the safe choice Q. Each player receives a signal $Q_i = Q + \epsilon_i$ that provides an unbiased estimate of Q. The signals are noisy so Q is not common knowledge amongst the players. Let Q denote a random variable that is distributed on the interval $[a, b]$ where $a < 0$ and $b > 1$. So it is possible that $Q > 1$, in which case B strictly dominates A and it is possible that $Q < 0$, in which case A strictly dominates B. Let $(\epsilon_1, \epsilon_2, \ldots, \epsilon_n)$ denote an n-tuple of zero mean independently and identically distributed random variables. The ϵ_i are assumed to be independent of Q and to be distributed within $[-E, E]$ where $E \leq -\frac{a}{2}$ and $E \leq \frac{b-1}{2}$. The incomplete information model is described by the following rules:

1. $(Q, \epsilon_1, \epsilon_2, \ldots, \epsilon_n)$ are randomly generated to obtain $(Q, Q_1, Q_2, \ldots, Q_n)$.
2. Player i observes $Q_i \equiv Q + \epsilon_i$ and chooses between A and B.
3. Each player i receives payoffs which are determined by choices made in step 2 and the actual value of Q in the Table 1 with the mean matching protocol.

Carlsson and van Damme (1993b) use the iterated elimination of strictly dominated strategies to show that there exists a unique dominance solvable equilibrium with a threshold $Q^* = 0.5$. That is, player i would select A if $Q_i < Q^*$ and B if $Q_i > Q^*$. Remarkably this is true for any $E > 0$ that is arbitrarily small. Carlsson and van Damme's argument thus gives another reason to expect the risk-dominant equilibrium if the players have arbitrarily small uncertainty about Q. We provide a detail of the proof by starting with a global stag hunt game with $n = 2$ and extend it to a general case of $n > 2$ in the Appendix.

Proposition 2. *For all $n \geq 2$, in a stag hunt game with incomplete information (i.e., a global stag hunt game), there exists a unique equilibrium threshold of 0.5. This threshold is consistent with the risk dominant equilibrium of the underlying complete information game.*

Proof. See Appendix. □

Tab. 3: Version of Global Stag Hunt Game Form Used in the Experiment

	A	B
A	500,500	$100, Q$
B	$Q, 100$	Q, Q

	A	B
A	500,500	$100, Q_i$
B	$Q_i, 100$	Q_i, Q_i

(a, left) Actual payoff table; (b, right) Subject i's estimated payoff table.

3 Design, Procedures, and Predictions

3.1 Experimental Design

The experiments implemented the stage game form given in Table 3: each player observed the right payoff matrix but his/her payoff was determined by the left payoff matrix. The stage game was played 100 times to give adequate experience for the iterated elimination of strictly dominated strategies to converge to equilibrium. The values of Q used in the experiment were integers in the interval 0 to 600, that is, $Q \in \{0, 1, 2, \ldots, 600\}$. The sequences of a hundred values of Q were generated by a computer using a uniform distribution. The units denote twentieths of a cent or 2,000 points = \$1.

Two treatments were conducted. The baseline treatment was designed to be intentionally similar to Rankin et al. (2000).[9] With complete information about Q, $\epsilon_i = 0$, that is $Q_i = Q$. The incomplete information treatment was designed to fit the motivating example of Carlsson and van Damme (1993a,b). The global games parameter $Q_i = Q + \epsilon_i$ where ϵ_i is uniformly distributed between -50 and 50, that is, $\epsilon_i \in \{-50, -49, \ldots, 49, 50\}$. The sequences were generated in the same way as the Q sequences in the complete information treatment. That is in each period, eight ϵ's were generated for eight participants in a cohort. The same sequence of Q_i was used in all cohorts of a treatment, but different sequences were used between the two treatments.

As is common in the stag hunt literature (Battalio et al., 2001; Rankin et al., 2000, etc.) subjects were grouped in cohorts of 8. Similar to Stahl and Van Huyck (2002) but unlike Rankin et al. we follow a "mean matching" protocol to accelerate convergence. That is, subjects are matched against all seven other subjects in the cohort each period and receive a payoff equal to the mean of the

[9] Our complete information treatment is different from Rankin et al. in three important ways. First, subjects are matched against everyone in the cohort each period and receive a payoff equal to the mean of the matches. Second, Q is allowed to be smaller than 0 and larger than 1 as required by global games theory to get a unique equilibrium under incomplete information. In order to apply the iterated dominance argument, the initial subclass of games must be large enough and contain games with different equilibrium structures. Lastly, action labels are fixed (a risky choice is always labeled A and a safe choice is always labeled B) and subjects play the games for 100 periods. After each period, each subject receives feedback on the the actual value of Q, the number of subjects in the cohort who chose A and B, and his/her payoff.

matches.[10]

3.2 Experimental Procedures

Each treatment consisted of 72 subjects divided over three session of three cohorts each. The participants were Texas A&M University undergraduates recruited campus wide using ORSEE, see Greiner (2015).

The experiment was programmed and conducted with the software z-Tree, see Fischbacher (2007). The experiment was conducted in the Economic Research Laboratory at Texas A&M University in February and March of 2013. A five dollar show up fee plus their earnings in the session were paid to the participants in private and in cash. The average earnings were about $29.19 for a session that lasted between 70 and 90 minutes.

After the instructions were read aloud, the participants filled out a questionnaire to establish that they knew how to calculate their earnings. After each period, each subject received feedback on the actual value of Q, the number of subjects in the cohort who chose A and B, and his/her payoff.

After the decision-making portion of the session was completed, participants filled out a questionnaire that asked them to explain their behavior in the session.

3.3 Theoretical Predictions

In our experiment, we transform the game in Tables 1 and 2 to the games in Table 3: $G_2 = 400 \times \; G_1 + 0.25\}$, where G_2 is the game used in the experiment and G_1 is the game in Table 1. This transformation matches previous work of Rankin et al. (2000) and Stahl and Van Huyck (2002) and avoids the possibility of decimal points and negative earnings in any periods.

We use a group size $n = 8$. That is, each player is matched with each of the other seven players in a two-player game and earns the average payoff from these matches (a mean matching protocol). Whether one considers this a game with 2 players (as is often convention) or a game with 8 players, Propositions 1 and 2 hold (i.e., they are defined for $n \geq 2$ with payoffs that apply to both cases).

Under complete information, there is no uncertainty about value of Q, that is, $Q_i = Q$. Proposition 1 shows payoff dominance threshold is 500 and the risk dominance threshold is 300 and all

[10] One may consider comparing behavior under incomplete information between two matching protocols: random matching and mean matching. Under complete information, the results from Rankin et al. (2000) (random matching) and Stahl and Van Huyck (2002) (mean matching) show no difference.

other thresholds are equilibria. This result produces a very ambiguous prediction.

Prediction 1. *Each cohorts will coordinate on a single threshold on* $[100, 500]$ *which is an equilibrium for the game.*

Under incomplete information, ϵ_i is uniformly distributed between -50 and 50. Applying Proposition 2, a unique threshold equilibrium Q^* satisfies equation (1) with $n = 8$ or

$$Q^* = \sum_{k=1}^{8} \frac{[100 + \{400 \times (\frac{k-1}{n-1})\}] \times \frac{7!}{(k-1)!(8-k)!}}{2^7}. \tag{1}$$

That is $Q^* = 300$, the threshold that is consistent with the risk dominant equilibrium of the underlying complete information game.

Prediction 2. *Cohorts will coordinate on the global games solution and risk dominant threshold 300, the sole equilibrium to survive removal of iterated dominated strategies in this global stag hunt game.*

4 Experimental Results

4.1 Aggregate Results

Figure 1 displays the overall frequency of choice A by value of Q_i by treatment, for the first 50 periods. Figure 2 provides the same information for the last 50 periods. The solid line on the left denotes the risk-dominant threshold, which is our predicted equilibrium for treatments with incomplete information. That is, subjects will select A when $Q_i < 300$ and B when $Q_i > 300$. All thresholds between 100 and 500 are equilibria for complete information treatments. The dotted line on the far right is a threshold of 500, the payoff-dominant threshold.

In the first 50 periods, the average threshold of both treatments were between risk dominance and payoff dominance. For Q_i below 300, in which payoff dominance and risk dominance agree on selecting A, almost all choices under both treatments were A, as expected. Subjects selected B more often than A only when Q_i was above 450. When Q_i was between 450 and 500, around 35 percent (42 percent) of choices under incomplete (complete) information were A.

The results of the last 50 periods were similar to the first 50 periods. The main difference was that in the last 50 periods, A was selected more often for Q_i below 350 and less often for Q_i

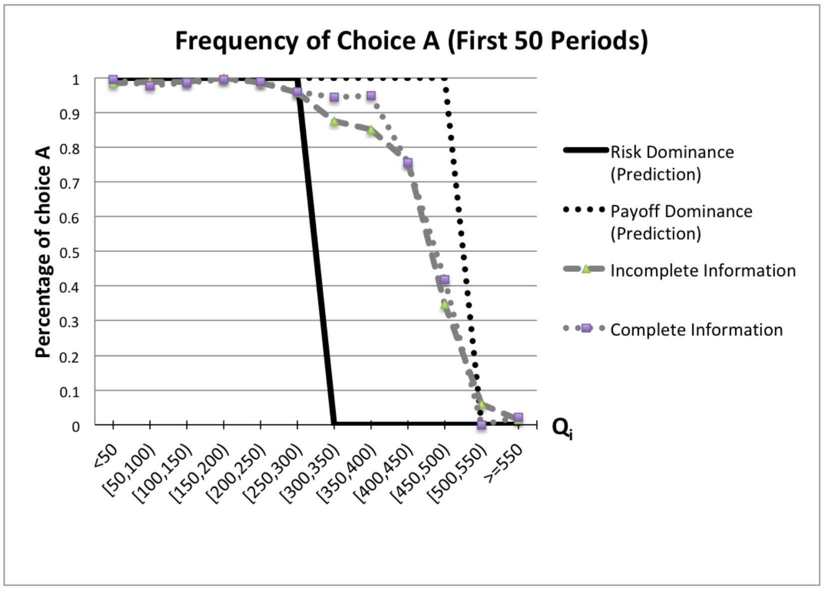

Fig. 1: **Frequency of choice A given the value of Q_i by treatment for the first 50 periods.** Each treatment line represents 3600 subject decisions averaged over the 12 signal realization categories.

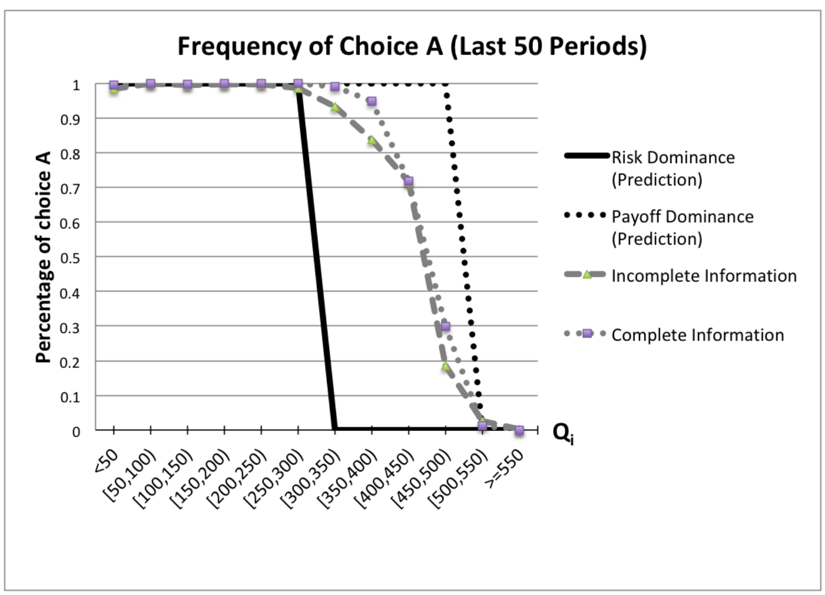

Fig. 2: **Frequency of choice A given the value of Q_i by treatment for the last 50 periods.** Each treatment line represents 3600 subject decisions averaged over the 12 signal realization categories.

above 400. Although A was selected more often under complete information than under incomplete information, the results under the two information conditions were similar.

4.2 Cohort Thresholds

Our predictions are made at the cohort level because if there is convergence to a single threshold, all subjects in the cohort must be playing the same strategy. This idea is common in both complete information stag hunt games and global currency attack games, despite the differences in procedures. As a consequence, we follow Heinemann et al. (2004) in estimating the following logit model for each cohort j,

$$\text{logit}\left(\Pr(A_{ijt}|Q_{it}, j)\right) = b_{0j} - b_{1j} \times Q_{it}, \tag{2}$$

where the probability of an individual, i, in cohort j, choosing to play A in period t is dependent on a fixed term determined individual's cohort (b_{0j}) and an interaction term of that cohort and observed private signal (b_{1j}). As Heinemann et al. note, the threshold can be estimated by $b_{0j}/b_1 j$ and the level of cohort dispersion by $\pi/\left(b_1\sqrt{3}\right)$. Table 4 reports both coefficients and values for periods 76 to 100.[11] Robust standard errors are clustered at the cohort-period level.

The average thresholds are 429.11 and 446.22 for the incomplete and complete information treatments, respectively. It is notable that the two values close to the risk-dominant threshold of 300 occurred in the incomplete information treatment and the one value essentially at the payoff-dominant threshold of 500 occurred in the complete information treatment.[12] However, comparing the 9 cohorts in each treatment, we cannot reject the hypothesis that both treatments were drawn from the same distribution ($p = 0.83$ for a Mann-Whitney test). Utilizing an alternate regression specification, similar to equation (2), that calculates treatment-specific coefficients rather than cohorts (see overall values in Table 4), we do not see a significant difference in thresholds ($p = 0.24$).

Most estimated thresholds are around 450 regardless of treatment. A stochastic steady state appears to have emerged for most cohorts with a threshold in the interval [400,500]. These thresholds are cohort-specific and would seem difficult to predict on an *a priori* basis.[13]

The dispersion measures across cohorts are generally higher in the incomplete information treat-

[11] We omit cohort 10 because it has a perfect threshold. In the last 25 periods, all subjects selected A when $Q \leq 486$ and B when $Q \geq 502$. There was no value of Q between 486 and 502; we used their average (494) as a threshold.

[12] In the last 50 periods, all eight subjects in that cohort played A when Q was in the interval $[0, 500)$ and played B when Q was in the interval $[500, 600]$ in every period. There is another cohort under the complete information treatment that produced an almost perfect step function but this cohort switched a strategy at $Q = 400$.

[13] A similar logistic regression (not shown) using a time trend as in equation (3) does not change the qualitative results of thresholds.

Cohort	b_0	b_1	estimated threshold b_0/b_1	estimated dispersion $\pi/\left(b_1\sqrt{3}\right)$
Incomplete Information: 9 cohorts, 72 subjects, 1800 observations				
1	18.193 (4.159)	-0.056 (0.012)	326.877 (3.77)	32.589 (7.308)
2	13.618 (2.241)	-0.037 (0.006)	369.282 (6.969)	49.185 (8.016)
4	14.552 (2.980)	-0.033 (0.007)	436.831 (7.489)	54.448 (11.340)
6	29.868 (6.443)	-0.068 (0.015)	439.802 (7.777)	26.708 (5.705)
3	97.968 (27.769)	-0.216 (0.059)	454.523 (5.910)	8.415 (2.294)
8	24.569 (3.032)	-0.054 (0.007)	455.457 (5.970)	33.624 (4.078)
9	76.794 (29.358)	-0.167 (0.063)	458.800 (5.359)	10.836 (4.041)
7	28.753 (3.936)	-0.062 (0.008)	460.112 (6.890)	29.025 (3.933)
5	31.069 (6.004)	-0.068 (0.014)	460.250 (4.772)	26.869 (5.414)
Overall	13.495 (2.662)	-0.032 (0.005)	424.665 (15.714)	57.078 (9.227)
Complete Information: 9 cohorts, 72 subjects, 1800 observations				
15	211.372 (13.316)	-0.527 (0.033)	401.085 (0.132)	3.442 (0.218)
16	18.248 (3.154)	-0.043 (0.007)	422.385 (6.739)	41.985 (6.743)
12	9.789 (4.200)	-0.023 (0.004)	423.908 (8.980)	78.547 (13.403)
18	27.263 (4.259)	-0.063 (0.010)	435.904 (3.704)	29.001 (4.720)
14	23.004 (7.484)	-0.051 (0.016)	452.832 (9.079)	35.705 (11.149)
13	32.422 (4.200)	-0.072 (0.009)	453.410 (4.940)	25.366 (3.322)
11	41.826 (8.430)	-0.091 (0.018)	460.581 (5.240)	19.973 (3.964)
17	68.300 (5.027)	-0.145 (0.011)	471.847 (0.716)	12.530 (0.920)
10	-	-	494 -	0 -
Overall	18.982 (3.202)	-0.043 (0.007)	446.385 (9.551)	42.653 (6.777)

Tab. 4: **Estimated Cohort Threshold Values from Logit Regressionfor the Last 25 Periods.** Since each cohort contains 8 subjects each cohort threshold is based on 200 total subject decisions. Standard errors are clustered at the cohort-period level for cohort estimates. A separate regression estimates the overall treatment thresholds with robust standard errors clustered at the cohort level.

ment, though not significantly so. The result contradicts the theoretical comparative static of global games—that would predict lower dispersion in the incomplete information treatment—because the incomplete information game contains a unique theoretical threshold. It also shows that observing incomplete information as in the global game setting could not solve the problems of the equilibrium multiplicity of the underlying game with complete information. In fact, it performs worse in predicting a unique threshold.

4.3 Individual Thresholds

While some subjects adjusted thresholds over time, others were "exact threshold players," individuals who always chose A for Q_i below a certain value and always chose B for Q_i above that value. For the last 25 periods, 76 percent of the subjects in the incomplete information treatment and 81 percent of the subjects in the complete information treatment were exact threshold players. In order to more precisely measure the heterogeneity of each subject, we estimated the following logit model for each subject:

$$\text{logit}\left(\Pr(A_{ijt}|Q_{it}, j)\right) = b_{0i} - b_{1i} \times Q_{it} + b_{2i} 1/t, \tag{3}$$

all three coefficients b_{0i}, b_{1i}, b_{2i} are specific for subject i. We also add a time trend, $1/t$, for period t, to capture rapid learning at the beginning of a session and smaller afterward. For subjects whose thresholds cannot be measured using the logit model, we use the middle of the switching point as the threshold.[14] The average estimated threshold, b_{0i}/b_{1i} (i.e., when $1/t$ approaches zero) are 432.894 and 448.855 for the incomplete and complete information treatments, respectively. This difference is not large (16), but it is marginally significant ($p = 0.107$ for a Mann-Whitney test).

The range of the estimated thresholds for subjects under incomplete information is wider than for subjects under complete information.[15] This further supports our idea that observing incomplete

[14] If a subject used an exact threshold, that is, he always played A when $Q_i \leq w$ and played B when $Q_i \geq z$ when there was no Q_i between w and z in any periods, we will use $\frac{w+z}{2}$ as a threshold. For example, if a subject always played A when $Q_i \leq 497$ and played B when $Q_i \geq 502$, that subject's threshold is 499.5. There are seven and eight exact threshold players under complete and incomplete information, respectively. If there are some errors, that is, a subject did not always play A when $Q_i \leq w$ or B when $Q_i \geq z$ for any values of w and z, we select w and z which yield the least errors. There are eight and seven players under complete and incomplete information whom we use this method, respectively. Because of their random behavior, we also exclude two players under complete information whose estimated thresholds from the logit model are above 600.

[15] The largest and smallest individual thresholds under incomplete information are 511 and 289, while under complete information they are 504 and 344. Similar to the cohort thresholds, the dispersion of estimated thresholds across individuals is higher in the incomplete information treatment. The standard deviations of thresholds are 51.1 and 37.6 for incomplete and complete information, respectively.

Earnings ($)	Incomplete	Complete
$Q^* = 100$ (Expected)	$18.17	$18.55
$Q^* = 300$ (Expected)	$23.17	$23.60
$Q^* = 500$ (Expected)	$24.83	$25.44
Average (Actual) Earnings	$24.03	$24.49

Tab. 5: Expected earnings when all subjects use the same thresholds of 100, 300, and 500 and average earnings in the experiments, by treatment (N=144).

information could not solve the problems of equilibrium multiplicity.

4.4 Efficiency Loss from Incomplete Information

There are two reasons for lower payoffs under incomplete information than under complete information. First, observing different private signals of Q could lead to mis-coordination even when players use the same threshold. For example, if two players use the same threshold of 300 but player 1 observes his private signal of 310 while player 2 observes 290, player 1 would play B and player 2 would play A, which results in mis-coordination. The payoff loss in our experiments from this type of mis-coordination alone would be approximately 58 cents. Second, incomplete information leads to a unique threshold equilibrium of 300. It is less efficient than a threshold of 500, a payoff-dominant threshold, which is one of the equilibrium thresholds under complete information. The expected payoff difference between the thresholds of 300 and 500 is $1.66 in our experiments.

Table 5 reports the expected earnings when subjects use various thresholds and the average earnings from our experiments under each information condition. The average earnings of each of the 72 subjects under incomplete information and complete information treatments are $24.03 and $24.35, respectively. This difference, at the subject level, is statistically significant ($p < 0.024$ for a t-test and $p = 0.0184$ for a Mann-Whitney test), but it is relatively small (only 32 cents). It is smaller than the difference between the average earnings of subjects under incomplete information and the theoretical prediction, which is 86 cents (the average earnings are also significantly greater than the expected earnings from using a threshold of 300, $p < 0.001$ for both treatments). This result is quite positive from an efficiency standpoint since subjects under incomplete information earn significantly greater than the theoretical prediction and the efficiency loss from observing imprecise information is not large.

4.5 Debriefing Questionnaire

After the 100 choices were made, the subjects were asked to complete a debriefing questionnaire consisting of four questions. The answers given clarified and confirmed the results we observed on subject coordination and decision-making.

The first question was: "What strategy did you use while playing this game? Please include details about what led you to choose A or B." Seventy-two percent of the subjects in the incomplete information treatment and ninety-two percent of the subjects in the complete information treatment reported clearly that they were using a threshold.[16] Twenty-five different exact thresholds were mentioned in the responses, ranging from 300 to 500.[17] Ten percent of the subjects reported what we call a fuzzy threshold. They would chose A for sure if Q was less than w and B for sure if Q was more than z, where $w < z$, and sometimes one or the other for Q in $[w, z]$.[18] This confirmed that subjects used thresholds in their decision-making process.

The second and the third questions were: "Did you change your strategy over time?" and "If you changed your strategy, what made you change it?" The typical response involved the behavior of the other players, in particular, the need to coordinate on the same threshold. For example, a subject wrote, "I was initially choosing the highest number of all those provided, so that was typically A unless B was a higher number. However, through the experiment other participants stopped choosing the highest number ($A = 500$) when B became more than 400." Our interpretation of this quote is that the participant started using what might be called a wishful thinking strategy (Maximax) because they wrote that the payoff to A was 500. Over time they learned that the group was coordinating on a threshold of 400 and this led them to change their behavior.[19]

[16] For example, a subject wrote, "I chose B when the odds were that Q was greater than 500. I used the estimate to decide this."

[17] The most frequent exact threshold strategy was to choose A if Q is less than 500 and B otherwise. It was chosen by nineteen percent of the subjects. Other popular choices were thresholds at 450, 400, and 440 to 445 in order of decreasing popularity. Self-reported and estimated thresholds in our analysis for each subject were similar.

[18] Subjects using a fuzzy threshold seem to be engaged in fast and slow thinking popularized by Kahneman (2011). Schotter and Trevino (2017) exploit the difference in measured response time to accurately predict observed individual thresholds in a global game. We could also view these subjects as using pure strategies for low and high signals, and mixed strategies for signals between w and z.

[19] Reading the debriefing answers from the cohort that perfectly coordinated on the payoff-dominant threshold of 500, cohort 10, we are now convinced that subjects initially started with a wishful thinking strategy rather than any equilibrium concepts such as payoff dominance. The answers from subjects in cohorts with significantly lower thresholds, cohorts 1 and 2, revealed that many subjects in these cohorts started by using low thresholds. Once other subjects observed that they could not coordinate on high thresholds, they had to change to lower thresholds. In contrast, more subjects in other cohorts started from using high thresholds.

5 Discussion

The experiment of Rankin et al. (2000) is intriguing in two ways. First, its design produces a level of efficient play in stag hunt experiments not seen in other design. Second, its use of perturbed payoffs brings the design functionally very close to a global game (as much as a game could ever be under complete information). Given that the global games prediction is equivalent to the criterion of risk-dominance and not payoff dominance, it is peculiar that these two properties are found in the same experimental design.

Our experiment replicates the design of Rankin et al., but provides a second treatment—that varies only slightly—to create a global game. Having only a minor functional difference between treatments is desirable for many reasons; one such reason is that we may directly compare threshold strategies. The thresholds in both treatments are quite similar and very close to the predictions of payoff dominance, but not risk dominance.

It is tempting to conclude that this is a victory for payoff dominance and that the theory of global games has little predictive power in stag hunt games. Indeed, other global games experiments using the more complicated currency attack game (often with further modifications) have reached similar conclusions. (e.g., Heinemann et al., 2004). But Cabrales et al. (2007) follow a very similar approach in their research and reach a very different conclusion. They find the global games threshold, equivalently risk dominance, is predictive in stag hunt games of long-run play. However, it is also similarly predictive under complete information, in a setting that does not follow the perturbed payoff structure of Rankin et al.. In both our paper and theirs, the results in the global game treatment appear dependent on the complete information analogue.

It is interesting to speculate on what may have caused the differences in results between our experiments. A natural starting point is to investigate the payoffs in our experiment. One possibility is that the low stakes in each period could lead subjects to risk more and learn more slowly. However, the accumulation of earnings from several periods would result in several dollars lost, and subjects would easily have been able to conclude this and adjust accordingly to maximize payoffs from each round. Another explanation could be that the gains from playing best response are small compared to the efficiency loss from moving to the equilibrium, and, as such, players would rather try to coordinate on high thresholds. Previous work suggests subjects are reluctant to play best response when it requires them to move to less efficient outcomes (e.g., Van Huyck et al., 1997). Our own back-of-the-envelope calculations suggest, in our experiment, the efficiency loss is more than nine times the monetary incentive to best respond. Because this value is highly dependent to the optimization premium (Battalio et al., 2001), the value of best responding, and Cabrales et al. (2007) find that subjects in the treatment with a higher optimization premium are more likely to converge to the theoretical prediction than those in the treatment with a lower optimization premium, there may be some validity to this story.

Another possibility may be that the structure of our game puts global games threshold is in the middle of the possible parameter realizations $Q^* = 0.5$, separating realizations into cases where payoff and risk dominance coincide $Q^* < 0.5$ and differ $Q^* > 0.5$. In Cabrales et al. (2007) all five possible states were above the risk dominance threshold, meaning there was no parameter realization where the two criterion coincided. In other words, in their experiment risk dominance would always select the secure option. Stahl and Van Huyck (2002) investigate this very issue in complete information using the Rankin et al. (2000) design, comparing two treatments with different ranges of experience: one with $0 < Q < 1$ and one with $0.5 < Q < 1$. They find that subjects in the first treatment coordinate on higher thresholds, providing possible plausibility to this story.

It is difficult to say what these results may mean for the theory of global games. It may be as Crawford et al. (2013) argue that the theory is not a good predictor of human behavior because it requires many steps iterated strategic thinking, and experimental evidence repeatedly finds that individuals fall short of the many steps required for this level reasoning. This argument could explain why subjects in the incomplete information treatments in our study deviate from the equilibrium prediction and use thresholds close to the payoff-dominant threshold. It also may explain the general lack of an effect of incomplete information in global games experiments.

Alternatively, there is the possibility many of these experiments (ourselves included) are testing global games incorrectly. Payoff perturbations may be present in every game—even those with complete information—because of strategic uncertainty. In such case, the meaning and identification of treatment effects between incomplete and complete information treatment is quite dubious; both treatments would actually be quite similar. Assuming the noise from strategic uncertainty fits the appropriate theoretical constraints, one would test global games simply by observing how predictive risk-dominance is as a equilibrium selection device under complete information. As we have previously stated, the evidence of this is quite mixed.

There is frustratingly very little we know for certain in this domain. As desirable as it would be to have a model to pick the "right" equilibrium in the stag hunt game, we cannot even test such model until we decide what is the “right” experimental framework in which to test this model.

References

Al-Ubaydli, O., Jones, G., and Weel, J. (2013). Patience, cognitive skill, and coordination in the repeated stag hunt. *Journal of Neuroscience, Psychology, and Economics*, 6(2):71.

Battalio, R., Samuelson, L., and Van Huyck, J. (2001). Optimization incentives and coordination failure in laboratory stag hunt games. *Econometrica*, 69(3):749–764.

Brindisi, F., Celen, B., and Hyndman, K. (2014). The effect of endogenous timing on coordination under asymmetric information: An experimental study. *Games and Economic Behavior*, 86:264–281.

Büyükboyacı, M. (2014). Risk attitudes and the stag-hunt game. *Economics Letters*, 124(3):323–325.

Cabrales, A., Nagel, R., and Armenter, R. (2007). Equilibrium selection through incomplete information in coordination games: An experimental study. *Experimental Economics*, 10:221–234.

Carlsson, H. and van Damme, E. (1993a). Equilibrium selection in stag hunt games. In Binmore, K., Kirman, A., and Tani, P., editors, *Frontiers of Game Theory*, pages 237–253. The MIT Press.

Carlsson, H. and van Damme, E. (1993b). Global games and equilibrium selection. *Econometrica*, 61(5):989–1018.

Clark, K., Kay, S., and Sefton, M. (2001). When are nash equilibria self-enforcing? an experimental analysis. *International Journal of Game Theory*, 29(4):495–515.

Clark, K. and Sefton, M. (2001). Repetition and signalling: experimental evidence from games with efficient equilibria. *Economics Letters*, 70(3):357–362.

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

Cooper, R., DeJong, D., Forsythe, R., and Ross, T. (1992). Communication in coordination games. *The Quarterly Journal of Economics*, 107(2):739–771.

Cornand, C. (2006). Speculative attacks and informational structure: An experimental study. *Review of International Economics*, 14(5):797–817.

Crawford, V., Costa-Gomes, M., and Iriberri, N. (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1):5–62.

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

Dubois, D., Willinger, M., and Van Nguyen, P. (2012). Optimization incentive and relative riskiness in experimental stag-hunt games. *International Journal of Game Theory*, 41(2):369–380.

Duffy, J. and Feltovich, N. (2002). Do actions speak louder than words? an experimental comparison of observation and cheap talk. *Games and Economic Behavior*, 39(1):1–27.

Duffy, J. and Feltovich, N. (2006). Words, deeds, and lies: strategic behaviour in games with multiple signals. *The Review of Economic Studies*, 73(3):669–688.

Duffy, J. and Ochs, J. (2012). Equilibrium selection in static and dynamic entry games. *Games and Economic Behavior*, 76:97–116.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125.

Harsanyi, J. C. and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. The MIT Press.

Heinemann, F., Nagel, R., and Ockenfels, P. (2004). The theory of global games on test: Experimental analysis of coordination games with public and private information. *Econometrica*, 72(5):1583–1599.

Heinemann, F., Nagel, R., and Ockenfels, P. (2009). Measuring strategic uncertainty in coordination games. *The Review of Economic Studies*, 76(1):181–221.

Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.

Kawagoe, T. and Ui, T. (2010). Global games and ambiguous information: an experimental study.

Morris, S. and Shin, H. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88:587–597.

Morris, S. and Shin, H. (2003). Global games: Theory and applications. In Dewatripont, M., Hansen, L., and Turnovsky, S., editors, *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, pages 56–114. Cambridge University Press.

Rankin, F., Van Huyck, J., and Battalio, R. (2000). Strategic similarity and emergent conventions: Evidence from similar stag hunt games. *Games and Economic Behavior*, 32(2):315–337.

Schmidt, D., Shupp, R., Walker, J. M., and Ostrom, E. (2003). Playing safe in coordination games:: the roles of risk dominance, payoff dominance, and history of play. *Games and Economic Behavior*, 42(2):281–299.

Schotter, A. and Trevino, I. (2017). Is response time predictive of choice? an experimental study of threshold strategies. Mimeo.

Shurchkov, O. (2013). Coordination and learning in dynamic global games: experimental evidence. *Experimental Economics*, 16(3):313–334.

Shurchkov, O. (2016). Public announcements and coordination in dynamic global games: Experimental evidence. *Journal of Behavioral and Experimental Economics*, 61:20–30.

Stahl, D. and Van Huyck, J. (2002). Learning conditional behavior in similar stag hunt games. Mimeo.

Szkup, M. and Trevino, I. (2015a). Costly information acquisition in a speculative attack: Theory and experiments. Mimeo.

Szkup, M. and Trevino, I. (2015b). Information acquisition in global games of regime change. *Journal of Economic Theory*, 160:387–428.

Van Huyck, J., Cook, J., and Battalio, R. (1997). Adaptive behavior and coordination failure. *Journal of Economic Behavior and Organization*, 32:483–503.

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

Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *The Quarterly Journal of Economics*, 106(3):885–910.

A Mathematical Appendix

A.1 Proof of Proposition 1

According to Harsanyi and Selton, the players' prior belief S_i about player i's strategy should coincide with the prediction of an outside observer who reasons in the following way about the game:

(i) Player i believes that his opponents will either all choose A or that they all choose B; he assigns a subjective probability z_i to the first event and $1 - z_i$ to the second.

(ii) Whatever the value of z_i, player i will choose a best response to his beliefs.

(iii) The beliefs (i.e. the z_i) of different players are independent and they are all uniformly distributed on [0.1].

From (i) and (ii), the outside observer concludes that player i chooses A if $z_i > Q$, and that he chooses B if $z_i < Q$. Hence, using (iii), the outside observer forecasts player i's strategy as $S_i = (1 - Q)A + QB$, with different S_i being independent. Harsanyi and Selten assume that the mixed strategy vector $S = (S_1, ..., S_n)$ describes the players' prior expectations in the game. Since S is not a Nash equilibrium, this expectation is not self-fulfilling, and, thus, has to be adapted. In a stag hunt game, the situation is symmetric, either all players will have A as the unique best response against S in which case all-A is the distinguished equilibrium, or they will all have B as the unique best response against S and in which case all-B is the distinguished equilibrium.

We can write player i's expected payoff associated with A when each player chooses A with probability t as:

$A_n^{\pi}(t) = \sum_{k=1}^{n} \frac{(n-1)!}{(k-1)!(n-k)!} t^{k-1}(1-t)^{n-k}\pi(k,n)$,

where $\pi(k,n)$ is the payoff of selecting A when k players including player i select A from a total of n players.

If the players' prior S is $(1-Q)A+QB$, then the expected payoff associated with A is $A_n^{\pi}(1-Q)$. Each player's best response against S is A if $A_n^{\pi}(1-Q) > Q$ and B if $A_n^{\pi}(1-Q) < Q$. When $\pi(k,n) = \frac{k-1}{n-1}$, $Q = A_n^{\pi}(1-Q)$ when $Q = 0.5$, we can conclude that A risk dominates B when $Q < 0.5$ and B risk dominates A when $Q > 0.5$.

A.2 Proof of Proposition 2

Consider a $2X2$ stag hunt game with incomplete information as shown in Table 2 where q is uniform on some interval $[a,b]$ where $a < 0$ and $b > 1$. Each player receives a signal $Q_i = Q + \epsilon_i$ that provides an unbiased estimate of Q. The ϵ_i is uniformly distributed within $[-E,E]$ where $E \leq -\frac{a}{2}$ and $E \leq \frac{b-1}{2}$. The signals Q_1 and Q_2 are independent. After observing the signals, the players choose actions simultaneously and get payoffs corresponding to the game in Table 2 with the actual value of Q. It is understood that the structure of the class of games and the joint distribution of Q, Q_1 and Q_2 are common knowledge.

It is easily seen that player i's posterior of q will be uniform on $[Q_i - E, Q_i + E]$ if he observes $Q_i \in [a+E, b-E]$, so his conditionally expected payoff from choosing B will simply be Q_i. Moreover, for $Q_i \in [a+E, b-E]$, the conditional distribution of the opponent's observation Q_j will be symmetric around Q_i and have support $[Q_i - 2E, Q_i + 2E]$. Hence, the probability that $Q_j < Q_i$ is equal to the probability that $Q_j > Q_i$, which is 0.5.

Now suppose player i observes $Q_i < 0$, his conditionally expected payoff from choosing B is negative and smaller than the minimum payoff from choosing A, which is 0. Hence A is conditionally strictly dominate B for player i when he observes $Q_i < 0$. It should be clear that iterated dominance arguments allow us to get further. For instance, if player j is restricted to playing A when observing $Q_j < 0$, then player i, observing $Q_i < 0$, must assign at least probability of 0.5 that player j would select A. Consequently, player i's conditionally expected payoff from choosing A will be at least 0.5, so choosing B (which yields 0) can be excluded by iterated dominance for $Q_i = 0$.

Let Q_i^* be the smallest observation for which A cannot be established by iterated dominance. By symmetry, obviously, $Q_1^* = Q_2^* = Q^*$. Iterated dominance requires player i to play A for any $Q_i < Q^*$, so if player j observes Q^* he will assign at least probability 0.5 to player i's choosing A and, thus, player j's expected payoff from choosing A will be at least 0.5. Since player j's expected payoff from choosing B equals Q^*, we must have $Q^* \geq 0.5$, for otherwise iterated dominance would require player j to play A when he observes Q^*.

We can proceed in the same way for large values of Q. B is dominant for each player if $Q_i > 1$ because the expected payoff from choosing B is larger than the maximum payoff of choosing A. Let Q^{**} be the largest value for which B cannot be established by iterated dominance. If player j observes $Q_j = Q^{**}$, his expected payoff from choosing A will be at most 0.5 given that player i conforms to iterated dominance. Since Q^{**} equals player j's expected payoff from choosing B, we can conclude that $Q^{**} \leq 0.5$. Combine these two findings with the fact that $Q^{*} \leq Q^{**}$, we get $Q^{*} = Q^{**} = 0.5$.

When $n > 2$, player i's payoff is the average payoff from the matches with the other $n-1$ players or $\pi(k,n) = \frac{k-1}{n-1}$. Since the game is symmetric, each player has the same Q^{*}. The probability that for each player $j \neq i$, $Q_j < Q_i$ is equal to the probability that $Q_j > Q_i$, which is 0.5. We can use the same argument as the case where $n = 2$ to conclude that A strictly dominates B at high values of Q_i and B strictly dominates A at low values of Q_i. We need to find a critical value of Q^{*} in which the expected payoff from selecting A and B are the same when player i observe Q^{*}. When subject i observes Q^{*}, his expected payoff is $\sum_{k=1}^{n} \frac{\pi(k,n) \times \frac{(n-1)!}{(k-1)!(n-k)!}}{2^{(n-1)}}$. This expression is equal to 0.5 which is the same to the expected payoff in a standard $2X2$ game. That is Q^{*} does not change and equals to 0.5. Using the mean matching protocol preserves the expected payoff which results in the same Q^{*} of 0.5.