

Equilibrium selection through incomplete information in coordination games: an experimental study



Antonio Cabrales · Rosemarie Nagel ·
Roc Armenter

Abstract We perform an experiment on a pure coordination game with uncertainty about the payoffs. Our game is closely related to models that have been used in many macroeconomic and financial applications to solve problems of equilibrium indeterminacy. In our experiment, each subject receives a noisy signal about the true payoffs. This game (inspired by the “global” games of Carlsson and van Damme, *Econometrica*, 61, 989–1018, 1993) has a unique strategy profile that survives the iterative deletion of strictly dominated strategies (thus a unique Nash equilibrium). The equilibrium outcome coincides, on average, with the risk-dominant equilibrium outcome of the underlying coordination game. In the baseline game, the behavior of the subjects converges to the theoretical prediction after enough experience has been gained. The data (and the comments) suggest that this behavior can be explained by learning. To test this hypothesis, we use a different game with incomplete information, related to a complete information game where learning and prior experiments suggest a different behavior. Indeed, in the second treatment, the behavior did not converge to equilibrium within 50 periods in some of the sessions. We also run both games under complete information. The results are sufficiently similar between complete and incomplete information to suggest that risk-dominance is also an important part of the explanation.

A. Cabrales (✉)
Universidad Carlos III de Madrid, Madrid, Spain
e-mail: antonio.cabrales@uc3m.es

R. Nagel
Universitat Pompeu Fabra, Barcelona, Spain
e-mail: rosemarie.nagel@upf.edu

R. Armenter
Federal Reserve Bank of New York, New York, USA
e-mail: roc.armenter@ny.frb.org

Keywords Global games · Risk dominance · Equilibrium selection · Common knowledge

JEL Classification C72 · C91 · C92 · D82 · G10 · G21

1 Introduction

There are many important economic phenomena that can be modeled as coordination games, that is, games with multiple equilibria. Often, these equilibria are Pareto-ranked. For example, unemployment or underemployment (Diamond 1982; Hart 1982; Bryant 1983), currency crises (Obstfeld 1996), and bank runs (Diamond and Dybvig 1983)¹ have all been modeled in this way. Some of the equilibria in these models look more plausible than others, and policy implications also depend on which equilibrium most closely represents reality.

But, without further information, it is hard to know what determines which equilibrium will be observed. In principle, even a sunspot, rather than any fundamental information, can be the explanation. This is not the way that people involved in these games usually perceive the situation,² and theoretically it is not very satisfactory. Furthermore, with multiplicity of equilibria, even comparative statics exercises are not straightforward (a change of policy parameters could, for example, trigger a change of equilibrium).

A recent approach to this problem has been to introduce incomplete information into the model. In this way, a unique equilibrium can arise. This equilibrium is often the one that was thought to be the most economically plausible. What is perhaps surprising is that *even small departures from complete information* (“almost common knowledge”) can generate unique equilibria, when under common knowledge there is more than one. This is clear in the e-mail game of Rubinstein (1989).³ Similarly, Carlsson and van Damme (1993) “slightly” perturb the players’ information in a coordination game.⁴ This leads to an incomplete information game with a unique solution by iterated deletion of *strictly* dominated strategies. Morris et al. (1995b) and Kajii and Morris (1997) further explain and generalize the logic behind the result of Carlsson and van Damme (1993).

Many researchers have applied this technique to solve equilibrium indeterminacy in the important economic situations we mentioned above. For example, Burdzy et al. (2001), and Frankel and Pauzner (2000) apply this technique in dynamic settings. Goldstein and Pauzner (1999) apply it to bank runs and Goldstein (1999) to banking and currency crises. Heinemann (2000), Heinemann and Illing (2002), and Morris

¹Cooper and John (1988) show that many of these models share the feature that there are spillovers between different players’ strategies. These spillovers create strategic complementarities, which often cause a multiplicity of equilibria.

²Most newspaper accounts of the recent currency crisis (with an added bank run element) in Argentina focus on fundamental factors rather than on sunspots.

³Other related examples can be found in Myerson (1991, p. 66) or Binmore (1992, p. 445).

⁴These games have two Nash equilibria in pure strategies, and one mixed-strategy equilibrium.

and Shin (1998) deal with speculative attacks. Morris et al. (1995a) study bubbles, and Shin (1996) asset trading.

Despite their theoretical and practical attractiveness, these ideas face some potential empirical difficulties. The logic of the result requires that players do many rounds of deletion of dominated strategies, each of which requires sophisticated Bayesian reasoning. It is well-known that real-life players are not particularly good at performing either of these tasks. Sefton and Yavaş (1996), for example, study experimentally a game that is derived from a coordination game. The unique solution of the transformed game, which also involves many rounds of deletion of strictly dominated strategies, is not the most-observed experimental outcome.⁵ More generally, experiments have shown that subjects apply low levels of reasoning (Stahl and Wilson 1995; Nagel 1995; Costa-Gomes et al. 2001; Costa-Gomes and Crawford 2002).

In this paper, we conduct an experimental test of a game inspired by Carlsson and van Damme (1993). In the baseline treatment under incomplete information, we find that the theoretical solution is the most frequent outcome after 50 periods. Furthermore, the frequency of strictly dominated actions, whose optimality does not require assumptions about other agents' rationality, decreases quickly. For actions that are optimal only if other (dominated) actions are played infrequently, the frequency of play decreased at a slower speed. However, our subjects clearly do not arrive at the solution by careful introspection,⁶ but rather adapt their play in response to their observations about the other players' behavior. When we repeat the game with common knowledge about the state of the world, the observed behavior converges sometimes to the payoff-dominant equilibrium.

To test the hypothesis that learning is the explanation for our results, we ran a different treatment, inspired by the work of Battalio et al. (2001) who analyze experimentally coordination games with *complete* information. In the new treatment, the basins of attraction of the equilibria in the underlying complete information game are identical to those in the baseline treatment, but the difference of payoffs between different strategies is lower (the *optimization premium* in the language of Battalio et al. 2001). Indeed, in the second treatment, the behavior in the lab did not converge to the unique equilibrium of the incomplete information game after 50 periods in some of the sessions, neither with complete nor with incomplete information about the state of the world. Our results are even more surprising than those of Battalio et al. (2001), since there is a unique equilibrium in the incomplete information version of our game.

We also performed simulations of the agents' behavior using the experience-weighted attraction learning model of Camerer and Ho (1999) for the two treat-

⁵In a betting game by Sonsino et al. (1999) and Søvik (2000), which is also solved by iterated deletion of strictly dominated strategies, the empirical solution does not coincide with the theoretical solution. It is possible, though, that the number of periods was not enough for equilibrium convergence, as the trend in her data suggests.

⁶Both their actions and their responses in the comments sheets suggest that they do not grasp the subtle Bayesian issues involved. Charness and Levin (2005) show that when Bayesian updating and reinforcement learning yield the same predictions, most subjects behave in accordance with the theory. However, when Bayesian-updating and reinforcement learning point into different directions, less than half of the subjects use Bayesian updating correctly.

Table 1 The 2×2 game, the true values, signals, conditional probabilities and expected utilities

True value X	90		80			70			60			50	
Signal S	T	U	T	U	W	U	W	Y	W	Y	Z	Y	Z
$P(X)$	1/5		1/5			1/5			1/5			1/5	
$P(S X)$	1/2	1/2	1/3	1/3	1/3	1/3	1/3	1/3	1/3	1/3	1/3	1/2	1/2

ments.⁷ The simulations have a qualitatively good fit with the experimental results, which further strengthens the evidence in favor of the hypothesis that learning is the explanation of behavior in this game.⁸

The remainder of the paper is structured as follows: in Sect. 2 we discuss theoretically the game that is played by the experimental subjects. Section 3 describes the experimental design. Section 4 collects the experimental results, Sect. 5 discusses the related experimental literature, and Sect. 6 concludes.

2 The game

2.1 Incomplete information games

(a) In the baseline treatment, the experimental subjects face an incomplete information game in which there are 2 players, each of whom has 2 actions, A and B. The payoffs for each action profile are given in the following matrix

	A	B
A	X, X	$X, 0$
B	$0, X$	$80, 80$

The random variable X can take 5 possible values, 90, 80, 70, 60 and 50 (row 1 in Table 1). Each of these variables have the same probability $p = 1/5$ (row 3 of Table 1). The players do not know the value of X , but independently receive a private signal about the value of X . The signals S_i can be: T, U, W, Y, Z. The relationship between signals S and the value of X are shown in row 1 and 2 of Table 1, e.g. if the true value is 90, then the signal can be only T or U or if the signal is T, the true value can be only 90, or 80. The conditional probabilities $P(S|X)$ are shown in row 4. The strategy for all players is a function from the set of signals to the set of actions.

Proposition 1 *The only strategy in this game that survives the iterated deletion of strictly dominated strategies is to play action A irrespective of the signal received.*

⁷Available at: <http://www.eco.uc3m.es/acabrales/research/pdf/cabnagelarmfinal.pdf>

⁸Camerer and Chong (2002) provide parameter estimates for several learning models using our data for game 1 with incomplete information.

Proof For complete proof, see Appendix A in the working paper version of this paper.⁹

Let us give some intuition here. The proof is done in four steps of iterated reasoning.

Step 1 (elimination of strictly dominated strategy): A rational player with signal T or U will never play B, since the expected utility of action A is greater than that of action B, irrespective of what a player with signal W or Y does.

Step 2 Knowing that a player with signal T or U will play A, a rational player with signal W will never play B, irrespective of what a player with signal Y, or Z does.

Step 3 Knowing that players with signal T, U, or W will never play B, a rational player with signal Y will never play B, irrespective of what a player with signal Y or Z does.

Step 4 Knowing that players with signal T, U, W, or Y will never play B, a rational player with signal Z will never play B, irrespective of what a player with signal Z does. \square

The signal received by the agents in Carlsson and van Damme (1993) is distributed continuously in a small interval around the value of X . We chose to modify the game in order to make the theoretical reasoning simpler. In particular, one can see in the proof of the proposition that the equilibrium can be reached after only four rounds of deletion of dominated strategies. Even so, the task still seems to be challenging.

It is also worth noting that the proposition assumes that the figures in the payoff matrix are Von Neumann–Morgenstern utilities. However, we do not control for risk preferences in the laboratory. It is not difficult to see that, if agents are risk loving, one can find other equilibria in this game in which an action is signal dependent. We could have tried to control for risk preferences, but Selten et al. (1999) have cast some doubts on procedures to control for risk aversion. Our results can be better explained by assuming that some agents are mildly risk loving and do not have a good assessment of play probabilities at the beginning of the game.

(b) *For the alternative control treatment*, the payoffs are

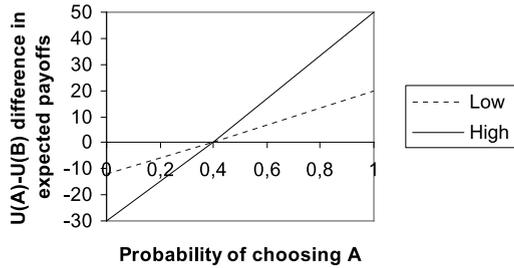
	A	B
A	X, X	$R, 0$
B	$0, R$	$80, 80$

The payoffs for B remain the same as in the baseline treatment (0 and 80, respectively). The pair of random variables (X, R) can take 5 possible values (36, 84), (32, 80), (28, 76), (24, 72) and (20, 68). The information structure is the same as in the baseline treatment. Each of these pairs of variables has the same probability $p = 1/5$. The players do not know the value of (X, R) , but independently receive a private signal about the value of (X, R) , and the signals S_i can be: T, U, W, Y, or Z.

In Fig. 1 we can see the difference in payoffs between strategy A and B as a function of the mixed-strategy of an opponent. We plot this (for the complete-information analog of the game) with $X = 50$ in the baseline treatment (high), and

⁹Available at: <http://www.econ.upf.edu/docs/papers/downloads/601.pdf>

Fig. 1 Expected payoff difference between strategies A and B in the complete information games (*full line*: C-base for $X = 50$, *broken line*: C-control for $(X, R) = (20, 68)$), for different probabilities of an A choice by the opponent



for $(X, R) = (20, 68)$ in the alternative control treatment (low). This figure will be useful in the following discussion.

Our alternative treatment was inspired by the work of Battalio et al. (2001), who did experiments with two (sets of) complete-information coordination games. Both games had identical structures in terms of the *sizes* of the basins of attractions,¹⁰ but in each set of games the *depth* of the basin of attraction was different. For a given equilibrium strategy pair, this *depth* is the difference in payoffs between the equilibrium strategy and the alternative.¹¹ This difference determines the strength of the incentives to choose a best-response. Battalio et al. (2001) showed that in one of the games the behavior in the laboratory always converges to the risk-dominant equilibrium, but in the alternative game, the behavior sometimes converges to the payoff-dominant equilibrium. They explain this phenomenon by resorting to learning dynamics.

As we show in Proposition 2, the new treatment has the same equilibrium structure as the baseline one. There is still a unique strategy profile that survives the iterated deletion of strictly dominated strategies. But, given the results of Battalio et al. (2001), one would expect markedly different behavior in the two treatments if our results are a consequence of learning. Indeed, we would like to interpret the presence or absence of different convergence behavior in both treatments as a “test” of the learning explanation for our results.

Proposition 2 *The only strategy in the alternative game that survives the iterated deletion of strictly dominated strategies is to play action A irrespective of the signal received.*

Proof See Appendix A in the working paper version of this paper.¹² □

¹⁰In this kind of game, the “basin of attraction” of an equilibrium is the set of mixed-strategy profiles for which it is a best-response for both players to play the equilibrium strategy. Its “size” is the area of that set. In Fig. 1 this is represented by the length of the segments of the horizontal axis to the right and left of its intersection with the payoff-difference function.

¹¹As the game is symmetric, this difference is the same for both players. The depth in Fig. 1 is the vertical distance to the payoff-difference function at 0 and 1.

¹²Available at: <http://www.econ.upf.edu/docs/papers/downloads/601.pdf>

2.2 Complete information games

For both payoff structures (control and alternative treatments) we also ran experiments with closely related complete information games. That is, the exact payoffs were given and this was common knowledge. Therefore, in the baseline treatment we have 5 possible games (randomly chosen in each period) with payoffs for $X = \{50, 60, 70, 80, 90\}$.

- If $X = 90$, the strategy combination AA is a unique equilibrium which is Pareto optimal and in dominant strategies.
- For $X = 80, 70, 60$, or 50 , the strategy combinations AA and BB are the two pure-strategy equilibria.

Similarly, for the alternative treatment, we have 5 possible games, with the possible realizations for (X, R) being $(36, 84)$, $(32, 80)$, $(28, 76)$, $(24, 72)$ or $(20, 68)$. For the first realization there exists only one equilibrium in dominant strategies, AA. In the other four games there are two pure equilibria: AA or BB. We do not consider mixed equilibria.

3 Experimental design

The experiments were run with undergraduates of all faculties in the Leex (Laboratori d'economia experimental) of the Universitat Pompeu Fabra. No subject could participate in more than one session. Upon arrival, students were randomly assigned to their seats. One of the instructors read the instructions (see Appendix B in the working paper version of this paper¹³) aloud and questions were answered in private. Sessions lasted about 1 hour including the reading of the instructions. All sessions presented here were run by computer using z-tree tools.

We present in this paper a 2×2 design of treatments as shown in the next table (in parentheses is the name we refer to later in the text).

	Complete info (C)	Incomplete info (IC)
High <i>optimization premium</i> (baseline)	3 sessions with two independent groups in each session (C-base called complete game 1)	2 sessions with two independent groups in each session (IC-base called incomplete game 1)
Low <i>optimization premium</i> (control)	2 sessions with two independent groups in each session (C-control called complete game 2)	4 sessions with two independent groups in each session (IC-control called incomplete game 2)

Each session consisted of 16 subjects, which were divided into 2 independent subgroups of 8 subjects. In this way, we obtain two independent observations for

¹³Available at: <http://www.econ.upf.edu/docs/papers/downloads/601.pdf>

each session. Each subject was randomly matched with an opponent in each round. In the instructions we state that subjects are randomly matched, though they don't know that there are two subgroups.

Each session consisted of 50 rounds.¹⁴ In each round, a true value out of the five possible values (X being 50, 60, 70, 80 or 90 for the baseline and (X, R) being (36, 84), (32, 80), (28, 76), (24, 72) and (20, 68) for the control treatment) was drawn randomly and independently for each pair. In the complete information game, the players in a pair were informed about the true value. In the incomplete information treatment, subjects were not informed of the true value, but instead each subject received an independent signal depending on the true value. The relationship between the value X and the signal as shown to the subjects is given in the first two rows in Table 1 for the baseline treatment. A similar table was given for the alternative treatment (see instructions). In the instructions, we mentioned the probabilities of 20% for each value X (X, R). Furthermore, we stated that each possible signal for a given X (X, R) had equal probability without specifying these probabilities. We did not state probabilities of X (X, R) conditional on these signals. We informed the players that the signal for each player was independently drawn from the same distribution. We emphasized several times that the underlying variable X (X, R) was the same within a pair but the signal could be different.

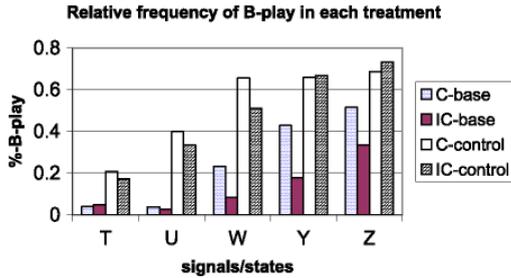
After each period, each player was informed about the choice of the other player and both players' payoffs. In the incomplete information treatments, he received the information about the true value of X (X, R). The signal of the other player was never revealed.

4 Experimental results

The main objective of this study is to test experimentally the effect of the introduction of payoff uncertainty in a coordination game. With this general goal in mind, we first compare the behavior in complete and incomplete information games. Second, we also observe the differences in behavior, or in speed of convergence to a particular behavior, after different signals when information is incomplete. The reason for this interest is that different levels of reasoning (rounds of deletion of dominated strategies) are required to compute the equilibrium behavior after different signals. Finally, we compare behavior between the baseline and control treatments in order to test whether convergence to the unique equilibrium depends on the payoff difference between the equilibrium strategy and the alternative strategy (that is, on the *depth* of the basin of attraction).

¹⁴Initially, we started with hand-run experiments of 15 rounds of the baseline treatment with incomplete information and one complete information game, in which X is always restricted to 50, that is the lowest value of all possible X s in the incomplete information treatments. The lowest value of X was chosen, since the signals of that value most likely induced a B choice in the uncertainty treatment. Since the behavior did not converge, neither with incomplete nor with complete information, we extended the rounds to 50. The results of the 15 period sessions can be found in the working paper version. Available at: <http://www.econ.upf.edu/docs/papers/downloads/601.pdf> and they are similar to the first 15 periods in the 50 period sessions.

Fig. 2 Relative frequency of B-play given signal/state in each treatment



In Fig. 2, we show the relative frequencies of B choices pooled over all rounds conditional on each possible signal/state separately for each treatment. Clearly, the lower the signal/state, the higher the relative frequency with which B is chosen within a treatment: only in 4 out of 22 independent observations do we observe fewer B-choices for a lower signal or state than for a higher signal or state. Thus, the null hypothesis of equal B-frequencies for all signal/states can be rejected in favor of the alternative hypothesis that a lower signal/state triggers more B-play than a higher signal/state at the 0.2% level, ($z = -2.98$) using a one-sided binominal test. For signal or state T, B-play amounts to 4.1% for C-base, 5% for IC-base, 21% for C-control and 17% for IC-control. There is no difference in frequency of B-play for the T signal between the complete or incomplete versions within the same game. The relative frequencies for B-play given T or U (strictly dominated strategies) are similar as in other studies on iterated reasoning (see e.g. Søvik 2000). With signal or state Z, B-play amounts to 62% for C-base, 33% for IC-base¹⁵ and 71% for C-control and 73% for IC-control.

The dynamics of behavior over time is shown in Figs. 3a–d. Here we separate the relative frequency of B-choices in 5 period blocks for each possible signal, pooled over all sessions per treatment. Subjects learn to avoid a B choice given a signal T or U in all treatments, though convergence is much slower in the control treatment. Only in IC-base does a signal Z trigger almost all A-play (on average 12%) in the final periods. In (both of) the control treatments and the C-base treatment, for signals or states Z and Y there is no convergence on average across sessions to playing A within 50 periods. In the IC-control game there is high dispersion of behavior over time for the Z-signal. In 4 out of the 8 groups of independent observations there is convergence to the risk dominant equilibrium, while in 2 there is convergence to all B-play and in the remaining two it is 60% and 66% B-play. Under complete information in both games, 2 out of 10 groups converge to the risk dominant equilibrium when the worst state is chosen (Z-state), while in the remaining groups B-play amounts to 77%. If we compare the frequency of B-play for the first 5 periods for signal/state Z between the complete and incomplete information treatments within the same game, there is no significant difference at the 10% level ($U = 15.5$ game 1; $U = 24$ game 2) using the Mann–Whitney U test. This holds true also for the two control treatments in the

¹⁵This difference in frequency of B-play between C-base and IC-base is significant at the 10% ($U = 19$), using the Wilcoxon–Mann–Whitney U-test.

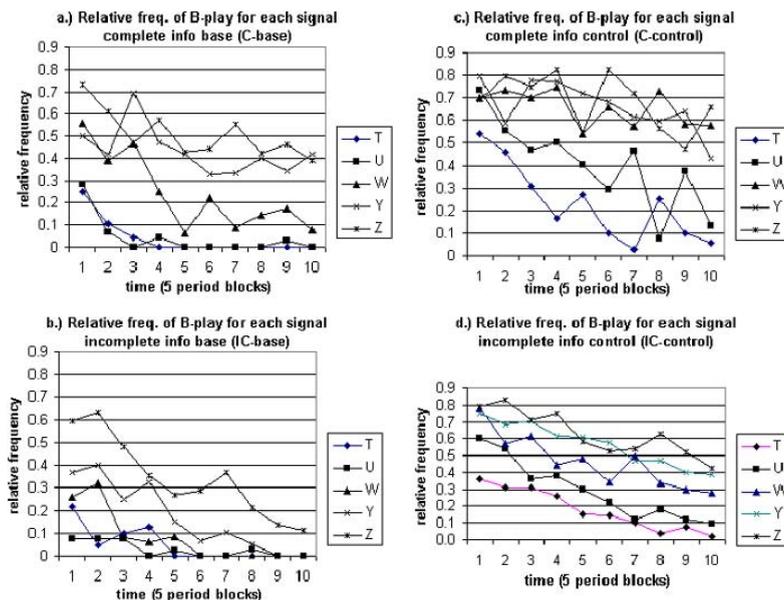


Fig. 3 Relative frequency of B-play over time, separately for each signal/state and treatment

last five periods ($U = 20$). However, the frequency of B-play is significantly (at the 5% level, $U = 20.5$) larger in the C-base than in the IC-base treatment.

To sum up: For signals or states T and U with “high” expected value of choice A (or low level of reasoning in incomplete information games) subjects quickly learn to avoid the B-choice in all treatments. However, only in the IC-base do players learn to play the A-choice for low states, and thus achieve the global game equilibrium as proposed by Carlsson and van Damme, which takes about 30 periods. Do subjects apply the concept of iterated elimination of dominated strategies? The comments give no indication that they do. Rather, it seems that they learn to avoid B-choices when they receive zero payoffs. This is already conjectured in Carlsson and van Damme (1993).¹⁶ In the working paper version,¹⁷ we explore more thoroughly whether learning is a good explanation for the observed behavior and the driving force for such a learning process.

¹⁶Carlsson and van Damme (1993) say: “This [common knowledge of rationality] justification of our solution concept, however, may not be totally convincing since a major motivation for our approach has been the wish to relax the common knowledge assumption. Hence it would be desirable to find alternative justifications which [...] dispense with it altogether. The kind of stories that naturally come to one’s mind are those were the strategy choices, instead of being determined by strictly rational considerations, result from some learning or evolutionary process.”

¹⁷Available at: <http://www.eco.uc3m.es/acabrales/research/pdf/cabnagelarmfinal.pdf>

5 Related experimental literature

There are only a few papers which test or refer to the theory of global games in the experimental literature. Heinemann et al. (2004) is the one most closely related to ours. These authors depart from the model of currency crises of Obstfeld (1996) and Morris and Shin (1998). Like us, Heinemann et al. (2004) also find a small, but significant, difference in behavior between complete and incomplete information games, with the equilibrium of the incomplete information game describing the observed behavior well.

There are several differences in the design and behavior of our paper and Heinemann et al. (2004). In the experiment we report here, the subjects play a 2 person game with changing opponents, and signals are taken from a discrete distribution. In Heinemann et al. (2004), the subjects repeatedly play a 15 player game, and signals are drawn from a uniform distribution. The advantage of a discrete distribution is that the equilibrium can be computed through a limited (and low) number of rounds of deletion of strictly dominated strategies. Furthermore, in our case the equilibrium outcome coincides on average with the risk dominant equilibrium of a related complete information game. In Heinemann et al. (2004) the unique equilibrium lies in between the risk dominant and Pareto dominant equilibria of the related complete information game. In Heinemann et al. (2004), the observed behavior does not change much over time, and is essentially explained by the initial choices, which are close to the global game solution. In our paper, on the other hand, behavior changes significantly over time and it can be well-explained by learning dynamics.

Having two players and a discrete signal space instead of many players and a continuous signal space allows us to better understand the determinants of behavior, and the influence of learning becomes much clearer. In this way, it becomes easy to know under which conditions convergence will become much slower (thus, for practical purposes, failing to occur). We can thus construct a control treatment where equilibrium behavior is shown empirically to be much harder to achieve.

Cornand (2006) extends the paper by Heinemann et al. (2004) by introducing a public signal in addition to the private signal and a treatment in which the agents receive two public signals. She finds that players rely more on the public signal than the private signal. The main implication is that a central bank should disclose a common signal if the agents have a private signal in order to decrease the probability of attack and the incidence of coordination failure.

Heinemann et al. (2007) analyze behavior in coordination games with complete information. They show that the global game equilibrium of the complete information game is a good description of the behavior in that game. As a consequence, it is also a good recommendation for playing the game. Cheung and Friedman (2005) test speculative attack models with continuous time, size asymmetries, and varying amounts of public information. Costain et al. (2005) report bimodal outcomes in sequential move coordination games with complete and incomplete information, as predicted by the model in Costain (2004). Brunnermeier and Morgan (2004) study “clock games” that end when the third of six players exits, and the first three players receive a payoff that increases continuously in the exit time. Consistent with the unique symmetric pure strategy Nash equilibrium, players exit sooner when they have better information about other players’ choices and clock settings.

6 Discussion

In this paper, we have tested experimentally the theoretical predictions of the model in Carlsson and van Damme (1993). Rubinstein (1989) analyzes a related game precisely to demonstrate that common knowledge predictions may not be robust because they require too much rationality from the agents.¹⁸ He argues that the theoretical prediction for his e-mail game would not hold in real-life.

Our results show that Rubinstein's (1989) conjecture was partly right. The theoretical results of Carlsson and van Damme (1993) do not hold under all circumstances. They certainly do not hold for players without experience in the game, and in some cases even after a relatively long experience. More interestingly, we have systematic evidence for when they are more likely to hold. When the strength of incentives (the *optimization premium* in the language of Battalio et al. 2001) is high (baseline), we are more likely to observe convergence to the theoretically predicted outcome. The Battalio et al. (2001) results, which inspired our alternative treatment, are suggested by learning theories. We have performed simulations of learning models with our game which indicate that learning is indeed a good candidate to explain our results.

We also find that the complete information version of the game presents features quite similar to those of the game with incomplete information (less so, but still remarkable, in the baseline treatment). It is thus plausible to think that that risk-dominance plays an important role in the result. One could perhaps interpret this similarity between complete and incomplete information as saying that the "controllable" part of incomplete information is not the only thing that explains the results. There may be a more fundamental lack of common knowledge about preferences, which, despite our efforts cannot really be controlled (degree of risk aversion, preference for "fairness" and so on). And even with common knowledge of preferences, there could still be lack of common knowledge about the equilibrium, which is why Harsanyi et al. (1988) introduced risk dominance. In their words: "risk dominance is important only in those situations where the players would be initially uncertain whether the other players would choose one or the other equilibrium."

Acknowledgements We thank Urs Fischbacher for teaching us to program this experiment and for providing his program, z-tree, Martin Menner and Aniol Llorente for technical support of the simulations, Frank Heinemann, and Nick Vriend for helpful comments. Cabrales and Nagel acknowledge the financial support of Spain's Ministry of Education under grant BEC2000-1029 and PB98-1076, the Generalitat de Catalunya under grant 1999SGR-00157, and thank the Barcelona Economics Program of CREA. Correspondence: Rosemarie Nagel, ICREA and Department of Economics, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain.

References

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

¹⁸Rubinstein (1989) asserts: "The sharp contrast between our intuition and the game-theoretic analysis is what makes this example paradoxical. The example joins a long list of games such as the finitely repeated Prisoner's Dilemma, the chain store paradox, and Rosenthal's game, in which it seems that the source of the discrepancy is rooted in the fact that in our formal analysis we use mathematical induction, while human beings do not use mathematical induction when reasoning."

- Binmore, K. (1992). *Fun and games: a text on game theory*. Lexington: Heath.
- Brunnermeier, M. K., & Morgan, J. (2004). *Clock games: theory and experiments* (Mimeo).
- Bryant, J. (1983). A simple rational-expectations Keynes-type model. *Quarterly Journal of Economics*, 98, 525–529.
- Burdzy, K., Frankel, D. M., & Pauzner, A. (2001). Fast equilibrium selection by rational players living in a changing world. *Econometrica*, 69, 163–189.
- Camerer, C. F., & Chong, J. C. (2002). *Functional EWA: a one-parameter theory of learning in games* (Mimeo).
- Camerer, C. F., & Ho, T. H. (1999). Experience-weighted attraction learning in games: a unifying approach. *Econometrica*, 67, 827–874.
- Carlsson, H., & van Damme, E. E. (1993). Global games and equilibrium selection. *Econometrica*, 61, 989–1018.
- Charness, G., & Levin, D. (2005). When optimal choices feel wrong: a laboratory study of Bayesian updating, complexity, and affect. *American Economic Review*, 95, 1300–1309.
- Cheung, Y. G., & Friedman, D. (2005). *Speculative attacks: a laboratory study in continuous time* (Working paper 606). Dept. of Economics, University of California, Santa Cruz.
- Cooper, R., & John, A. (1988). Coordinating coordination failures in Keynesian models. *Quarterly Journal of Economics*, 103, 441–463.
- Cornand, C. (2006). Speculative attacks and informational structure: an experimental study. *Review of International Economics*, 14, 797–817.
- Costa-Gomes, M., & Crawford, V. P. (2002). *Cognition and behavior in two-person guessing games: an experimental study*. Manuscript in preparation, University of California, San Diego, and University of York.
- Costa-Gomes, M., Crawford, V. P., & Broseta, B. (2001). Cognition and behavior in normal-form games: an experimental study. *Econometrica*, 69, 1193–1235.
- Costain, J. (2004). *A herding perspective on global games and multiplicity* (Working Paper 03-29 (08), May 2003). Department of Economics, Universidad Carlos III Madrid. Revised May 2004.
- Costain, J., Heinemann, F., & Ockenfels, P. (2005). *Multiple outcomes of speculative behavior in theory and in the laboratory*. Manuscript, Department of Economics, Universidad Carlos III Madrid.
- Diamond, P. (1982). Aggregate demand management in search equilibrium. *Journal of Political Economy*, 90, 881–894.
- Diamond, D. W., & Dybvig, P. H. (1983). Bank runs, deposit insurance and liquidity. *Journal of Political Economy*, 91, 401–419.
- Frankel, D. M., & Pauzner, A. (2000). Resolving indeterminacy in dynamic settings: the role of shocks. *Quarterly Journal of Economics*, 115, 285–304.
- Goldstein, I. (1999). *Interdependent banking and currency crises in a model of self-fulfilling beliefs* (Mimeo). Tel-Aviv University.
- Goldstein, I., & Pauzner, A. (1999). *Endogenous probability of bank runs in a rational expectations model* (Mimeo). Tel-Aviv University.
- Harsanyi, J. C., et al. (1988). *A general theory of equilibrium selection in games*. Cambridge: MIT Press.
- Hart, O. (1982). A model of imperfect competition with Keynesian features. *Quarterly Journal of Economics*, 97, 109–138.
- Heinemann, F. (2000). Unique equilibrium in a model of self-fulfilling speculative attacks: comment. *American Economic Review*, 90, 316–318.
- Heinemann, F., & Illing, G. (2002). Speculative attacks: unique sunspot equilibrium and transparency. *Journal of International Economics*, 58, 429–450.
- Heinemann, F., Nagel, R., & 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., & Ockenfels, P. (2007). *Measuring strategic uncertainty in coordination games* (Working paper). Universitat Pompeu Fabra.
- Kajii, A., & Morris, S. (1997). The robustness of equilibria to incomplete information. *Econometrica*, 65, 1283–1310.
- Morris, S., & Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88, 587–597.
- Morris, S., Postlewaite, A., & Shin, H. S. (1995a). Depth of knowledge and the effect of higher order uncertainty. *Economic Theory*, 6, 453–467.
- Morris, S., Rob, R., & Shin, H. S. (1995b). *p*-dominance and belief potential. *Econometrica*, 63, 145–157.
- Myerson, R. B. (1991). *Game theory: analysis of conflict*. Cambridge: Harvard University Press.

- Nagel, R. (1995). Unraveling in guessing games: an experimental study. *American Economic Review*, 85, 1313–1326.
- Obstfeld, M. (1996). Models of currency crisis with self-fulfilling features. *European Economic Review*, 40, 1037–1047.
- Rubinstein, A. (1989). The electronic mail game: strategic behavior under ‘almost common knowledge’. *American Economic Review*, 79, 385–391.
- Sefton, M., & Yavaş, A. (1996). Abreu–Matsushima mechanisms: experimental evidence. *Games and Economic Behavior*, 16, 280–302.
- Selten, R., Sadrieh, A., & Abbink, K. (1999). Money does not induce risk neutral behavior, but binary lotteries do even worse. *Theory and Decision*, 46, 211–249.
- Shin, H. S. (1996). Comparing the robustness of trading systems to higher-order uncertainty. *Review of Economic Studies*, 63, 39–59.
- Sonsino, D., Erev, I., & Gilat, S. (1999). *On rationality, learning and zero-sum betting—an experimental study of the no-betting conjecture* (Mimeo). Technion University.
- Søvik, Y. (2000). *Strength of dominance and depths of reasoning* (Mimeo). University of Oslo.
- Stahl, D. O., & Wilson, P. W. (1995). On players’ models of other players: theory and experimental evidence. *Games and Economic Behavior*, 10, 218–254.