OXFORD REVIEW OF ECONOMIC POLICY, VOL. 18, NO. 4

EXCHANGE-RATE ATTACK AS A COORDINATION GAME: THEORY AND EXPERIMENTAL EVIDENCE

FRANK HEINEMANN
Ludwig-Maximilians-Universität München[1]

This paper compares theoretical predictions for a coordination game, used to explain the onset of a currency crisis, with observations from laboratory experiments. Theories that assume full rationality suggest that public information may destabilize an economy by creating self-fulfilling belief equilibria, while private information leads to a unique equilibrium. In experiments, differences in behaviour for these two kinds of information are small. Public information increases efficiency and coordination among players, and there is no evidence for destabilizing effects owing to self-fulfilling beliefs.

I. INTRODUCTION

Coordination games with strategic complementarities are a common structure of economic decision problems. Examples are investment decisions with super-additive efforts, purchasing network commodities, financial market activities, and market intermediation. In these games, the payoff to some action is positively related to the number of players who take the same action. Strategic complementarities may result in multiple equilibria with self-fulfilling beliefs. If a player believes that all other players take a particular action, his or her best response is taking the same action. If *all* players believe that all others take some action A, then they all take action A, thus, beliefs are correct and establish a Nash equilibrium.

This paper reviews research on a particular coordination game, used to explain the onset of a currency crisis, and compares theoretical predictions with evidence from laboratory experiments. A speculative attack on a currency with a fixed exchange rate is a typical example of a coordination game with strategic complementarities. Traders, who expect a devaluation, short sell the currency and create pressure on the exchange rate. If many traders take

[1] I would like to thank Stephen Morris, David Myatt, Rosemarie Nagel, Peter Ockenfels, Hyun Shin, and an anonymous referee for helpful comments and stimulating discussions.

the same position, market pressure may force the central bank to abandon the peg and expectations are correct. If traders expect the peg to hold, even small transaction costs prevent short selling, and the lack of market pressure allows the retention of the peg.

Multiple equilibria make it hard to predict the aggregate outcome for any combination of exogenous parameters or to predict the effects of a change in these parameters even qualitatively. Unrelated events may lead to a change in individual strategies. News without economic content (sunspots) may have an impact on the economy, if players only believe in it. An economy that is best described by a model with multiple equilibria may be unstable. Beliefs may change in an unpredictable way and collective moods may trigger a sudden crisis that has no fundamental economic reason and cannot be resolved by economic instruments.

Recently, the theory of global games, developed by Carlsson and van Damme (1993*a*,*b*) and Morris and Shin (2002), has provided a convincing solution to the problem of multiplicity. If players do not have common knowledge,[2] but only private information about the payoff function, coordination games may have a unique equilibrium. This method is convincing for two reasons: first, because common knowledge is a very demanding requirement and almost impossible to achieve in practice[3] and second, because strategic complementarity allows the derivation of the equilibrium (if it is unique) by iterative elimination of dominated strategies, a concept with firm decision-theoretic foundation (Brandenburger, 1992).

These game-theoretic developments have stimulated numerous research projects on coordination games such as the speculative-attack problem. Fukao (1994) and Morris and Shin (1998) show that private information on the state of the economy generates a unique equilibrium, in which traders attack a currency if and only if they believe the state of the economy to be worse than a unique equilibrium threshold. Heinemann (2000) and Heinemann and Illing (2002) analyse how the probability of speculative attacks is related to underlying fundamentals and to the precision of private information. Morris and Shin (1999) and Hellwig (2001) show that uniqueness of the equilibrium requires private information to be sufficiently precise when compared to public information.[4] These results have triggered a debate on the stabilizing effects of information. Public dissemination of information may lead to common knowledge and multiple equilibria and, thereby, destabilize an economy. With private information, on the other hand, there is a unique equilibrium and the economy is not in danger of being moved by self-fulfilling beliefs.

Equilibrium theories, including the iterative elimination of dominated strategies, require agents to follow an infinite number of levels of reasoning about the reasoning of other players about the reasoning of other players, and so on. Experimental evidence, however, shows that most real subjects fail to follow more than, at most, three levels of reasoning. For any limited number of levels of reasoning, private information is not sufficient for a unique prediction. But the theory of global games also provides a plausible story about the formation of beliefs under bounded rationality. It might be used as a selection criterion. However, there are other theories on belief formation that compete with the theory of global games.

It is an empirical question as to whether public information destabilizes an economy by creating self-fulfilling beliefs or whether any selection theory gives a good prediction of actual behaviour. Cabrales *et al*. (2002) and Heinemann *et al*. (2002) test the predictions of the theory of global games in laboratory experiments with public and private information. They find only small differences in behaviour between sessions with public and private information. In particular, there is no evidence supporting the view that public information may destabilize an economy by creating self-fulfilling beliefs. In both information conditions, subjects coordinate on threshold strategies that are fairly predictable.

While Cabrales *et al*. (2002) observe convergence of behaviour towards the risk-dominant equilibrium, as predicted by the theory of global games, the experiment by Heinemann *et al*. (2002) suggests

[2] A statement is common knowledge if everybody knows it, everybody knows that everybody knows it, and so on.

[3] For the problems associated with establishing common knowledge see the article by Stephen Morris in this issue.

[4] Public information is a statement that is common knowledge among players. A player's information is private if others do not know it.

that real subjects settle on an equilibrium somewhere between the payoff-dominant and the risk-dominant equilibrium. Although statistical tests reject the prediction of the theory of global games numerically, they find strong support for the qualitative comparative statics of the global-game solution with respect to parameters of the payoff function: increasing transaction costs or capital controls reduce the probability of a currency crisis.

The next section describes the exchange-rate attack game. Section III gives theoretical predictions for this game and section IV describes the design of the experiment by Heinemann *et al*. (2002). Section V reports results and section VI collects the arguments in favour of and against public information. Section VII compares observed behaviour with predictions from theories of equilibrium selection. The last section concludes and raises questions that are open for future research.

II. EXCHANGE-RATE ATTACK AS A COORDINATION GAME

Consider a central bank that pegs the exchange rate of its currency to some other currency or currency basket. Reasons may be the desire to join a monetary union, attempts to reduce exchange-rate uncertainty, or a commitment to earning high credibility for domestic monetary policy. However, there may be good reasons to give up the peg in response to major economic shocks. Obstfeld (1986) introduced a model of a balance-of-payments crisis, where a realignment of the exchange rate is associated with fixed costs. Fluctuations in economic fundamentals call for adjustments of the exchange rate, and keeping the currency peg is associated with a welfare loss that is rising in the difference between the fixed rate and the shadow rate that would prevail in a regime of flexible exchange rates. If this difference is small, the costs of realignment exceed the welfare loss associated with the fixed rate, and the central bank keeps the peg. But large shocks to the economy may be associated with a welfare loss that exceeds the costs of realignment, so that the central bank prefers to give up the peg and the exchange rate is adjusted towards the shadow rate.

Agents who expect a devaluation borrow the currency, sell it, and hope to be able to buy it back at a lower rate when they have to repay their debt: they 'attack' as we say. The payoff to an attack is the difference in exchange rates minus transaction costs that include the difference in interest rates between domestic and foreign currency. Speculative attacks increase the supply of the currency on the foreign-exchange market and thereby raise the costs to the central bank of maintaining the peg. In order to defend a fixed rate, the central bank must absorb excess supply of its currency and sell foreign reserves. The larger the amount of capital involved in an attack, the more reserves must be sold to keep the peg. The costs of defending a fixed exchange rate add to the welfare loss stemming from shocks to economic fundamentals and, therefore, the decision to devalue may crucially depend on the expectations of private agents about future exchange rates.

The payoff function in this game depends on the state of the economy: the description of a state may contain productivity measures, foreign reserves, trade balance, or public deficit. The state of the economy has implications for the payoff from a successful speculative attack and also for the degree of coordination that is needed to enforce a realignment. The better the state of the economy, the higher is the shadow exchange rate and the smaller are the gains from a potential devaluation. In addition, in stronger economies, the central bank is in a better condition to fight attacks. The amount of capital needed to enforce a devaluation rises with better states.

If the economy is in good shape, the shadow exchange rate is close to or even above the fixed rate and the potential gain from a realignment is smaller than transaction costs of an attack. An attack is inevitably associated with a negative payoff to those who attack. Here, it is a dominant strategy not to attack, and there is a unique equilibrium in which the peg is maintained. If the economy is in poor condition, the welfare loss of keeping the peg exceeds fixed costs of a realignment and a devaluation is inevitable. An attack of any size is successful and yields positive profits and, hence, attacking is a dominant strategy.

Figure 1
The Exchange-rate Attack Game

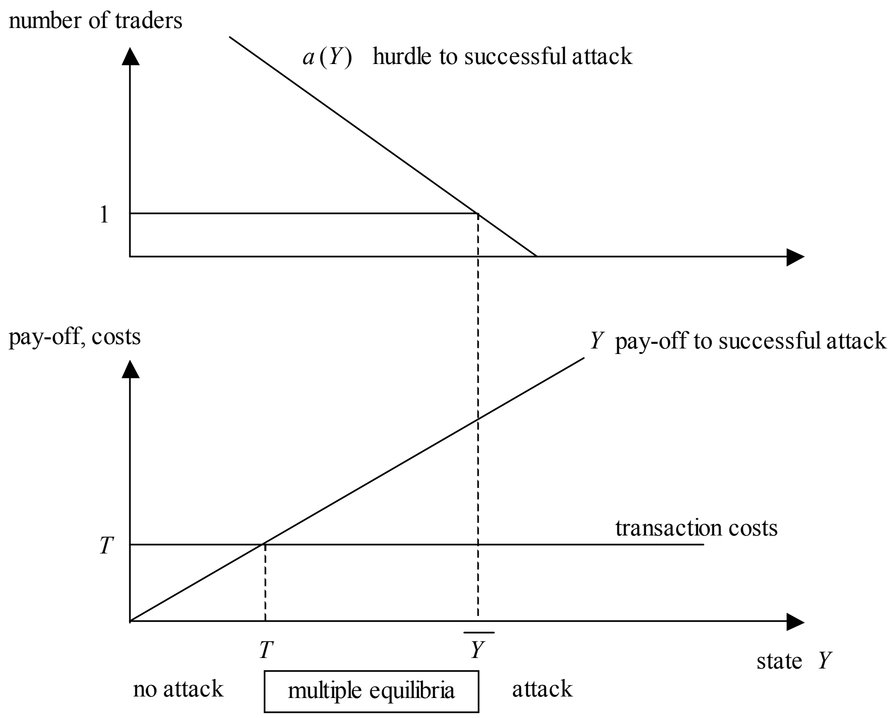

However, if the country is in an intermediate position, the model has multiple equilibria. If traders expect a devaluation and attack the currency, the outflow of reserves is large, the central bank prefers a realignment, and traders' expectations are correct. If traders expect the exchange rate to hold, they do not attack. Market pressure is lower and the rate can be sustained as expected. There are two pure-strategy equilibria with self-fulfilling beliefs. Either all agents believe in devaluation, attack, and are rewarded accordingly, or no agent attacks and a single attacking player would only lose on transaction costs, because a single agent is too small to enforce a devaluation.

Let us consider a simple version of this game, based on Morris and Shin (1998). There is a finite number of traders n, who may decide whether to attack or not. An attack is associated with transaction costs T. If the currency is devalued each attacking agent earns an amount Y. Y is the difference between the currency peg and the shadow exchange rate. The lower is Y, the higher is the shadow rate and the better is the state of the economy. An attack is successful if and only if a sufficient number of traders decides to attack. The hurdle to success is higher for better states of the economy and therefore it is a decreasing function in Y. Let $a(Y)$ be the number of players who are needed to enforce a devaluation and assume that for $Y = T$, $1 < a(Y) \leq n$ holds.[5] Figure 1 shows the hurdle to successful attacks and the payoff function for this game. If at least $a(Y)$ traders attack, attacking traders receive a payoff $Y - T$. Otherwise they lose T.

If $Y < T$, potential gains from an attack are smaller than transaction costs. The economy is in good shape and it is a dominant strategy not to attack.

If $Y > \overline{Y}$, a single trader can enforce a devaluation and yield a reward that covers transaction costs. $\overline{Y}$ is defined by $a(\overline{Y}) = 1$. Here, the economy is so weak that devaluation is inevitable. An attacking trader is certain to gain and risks nothing. For these states attacking is a dominant strategy.

For intermediate states $T < Y < \overline{Y}$, an attack is rewarding if successful, but success requires a

[5] The assumption that $a(T) > 1$ ensures that $T < \overline{Y}$. With $a(T) \leq n$, an attack is enforceable by a coordinated attack of all players, whenever this is profitable.

coordinated attack by a sufficient number of traders. In these intermediate states we have two equilibria: if all players attack, it is a best response to attack. If no player attacks, it is a best response not to attack.

As this game has multiple solutions, players are uncertain about the strategies of other players. Depending on subjective beliefs, the risk of losing transaction costs if an attack fails must be weighted against the gain from an attack succeeding. Probabilities of success and failure are subjectively determined. They cannot be deduced from rationality and common knowledge of the state. But, we can build theories to explain these beliefs that can be tested in experiments.

III. THEORETICAL PREDICTIONS

The theory of global games,[6] introduced by Carlsson and Van Damme (1993*a*,*b*) and applied to the exchange-rate-attack problem by Fukao (1994) and Morris and Shin (1998), recognizes that perfect information about the state of the world is an unrealistic assumption. Suppose that the true state Y is subject to some probability distribution and players lack perfect information on Y. Instead, they get private signals $X^i = Y + u^i$, where idiosyncratic noise terms u^i are subject to identical and independent probability distributions with $E(u^i) = 0$.

A strategy is now a function, assigning either of two actions (attack or not) to any possible signal X^i. Even if players know the strategies of other players, they cannot calculate the exact payoff to a speculative attack, but only a conditional distribution depending on their private signal.

Morris and Shin (1998) assume that signals have a uniform distribution in an ε-surrounding of the true state and Y has a uniform distribution with sufficiently wide support, so that for all relevant signals the conditional distribution of Y given X^i is uniform in an interval $[X^i \pm \varepsilon]$. As it turns out, and is formally proven by Morris and Shin (1998), the analysis of equilibria can be restricted to consider 'cut-off' strategies. A 'cut-off' strategy is characterized by a 'cut-off' value $\widetilde{X}$, such that a player attacks the currency if and only if his or her private signal X^i exceeds $\widetilde{X}$. Furthermore, equilibria are symmetric in the sense that in equilibrium all players have the same 'cut-off' value. An equilibrium is characterized by the property that a player who gets the marginal ('cut-off') signal is indifferent towards attacking. Morris and Shin show that the conditional expected payoff to a player who gets some signal $\widetilde{X}$, given that other players attack if and only if they receive signals above $\widetilde{X}$, is an increasing function in $\widetilde{X}$. Hence, there is a unique X^* for which this expected payoff equals transaction costs and makes the marginal player indifferent.

To understand this result, suppose that all players follow some 'cut-off' strategy, characterized by $\widetilde{X}$. The probability that an attack by some player i is successful at state Y equals the probability that at least $a(Y) - 1$ of the other players get signals above $\widetilde{X}$ and is given by

$$p(Y,\widetilde{X}) = \begin{cases} 1 & \text{if } \widetilde{X} \le Y - \varepsilon \\ 1 - Bin\left(\hat{a}(Y) - 2, n-1, \dfrac{Y - \widetilde{X} + \varepsilon}{2\varepsilon}\right) & Y - \varepsilon < \widetilde{X} < Y + \varepsilon \\ 0 & \widetilde{X} \ge Y + \varepsilon, \end{cases}$$

where $\hat{a}(Y)$ is the round-up of $a(Y)$ and Bin is the cumulative binomial distribution. The conditional expected payoff to some attacking player with signal X^i is then given by

$$\frac{1}{2\varepsilon} \int_{X^i - \varepsilon}^{X^i + \varepsilon} Y\, p(Y,\widetilde{X})\; dY - T.$$

Now, consider a player who receives the marginal signal $X^i = \widetilde{X}$. His or her expected payoff from attacking is

$$U(\widetilde{X}) = \frac{1}{2\varepsilon} \int_{\widetilde{X} - \varepsilon}^{\widetilde{X} + \varepsilon} Y\, p(Y,\widetilde{X})\; dY - T,$$

[6] The term 'global game' refers to the interpretation that the true game characterized by Y is drawn randomly from some larger class of games.

which is monotonically increasing in $\widetilde{X}$. Hence, there is a unique solution to $U(X^*) = 0$ that defines the equilibrium 'cut-off' signal X^*.

With a uniform distribution of Y, the probabilistic environment generates an additional equilibrium condition that is sufficiently strong to characterize a unique equilibrium, in which each player attacks if and only if he or she gets a signal above the equilibrium threshold X^*. For normal distributions, Morris and Shin (1999) and Hellwig (2001) show that uniqueness requires the variance of private information to be sufficiently small in comparison to the prior variance of Y. Frankel *et al*. (2001) show that for any probability distribution on states and signals, the set of equilibria converges towards a single strategy profile if the variance of private information approaches zero. Thus, uniqueness of equilibrium is a general property of global games provided that private information is sufficiently precise.

The speculative-attack game has a unique equilibrium when information is private and multiple equilibria when there is public information. This has triggered a discussion on optimal ways to release central bank information. If the central bank releases its own information publicly, it may risk triggering an attack out of self-fulfilling beliefs, especially at states where attacks would not occur if information was private. Transparency may be harmful if it destabilizes an economy by making attacks unpredictable. The question here is whether actual players behave in the way that is suggested by the theory of global games.

In a Nash equilibrium, players act as if each of them knows the strategies of all other players, although these strategies are formed simultaneously. In fact, knowledge of other players' strategies is an assumption inherently associated with the concept of Nash equilibrium (Aumann and Brandenburger, 1995).[7] In games with strategic complementarities, the limits to the set of Nash equilibria can also be calculated by using the process of iterative elimination of dominated strategies. This requires players not to play a dominated strategy. Furthermore, they are assumed to believe that other players do not play a dominated strategy, that others believe others not to play a dominated strategy, and so forth. The assumed mental process consists of an infinite number of levels of players' reasoning about each others' beliefs. From various experiments, however, we know that real players behave as if they apply only a limited number of levels of reasoning. This casts doubts on the descriptive power of the Nash equilibrium concept, independent of whether it predicts a single or multiple equilibria.

There is another interpretation of the global game solution that is an *ad-hoc* assumption on beliefs, but does not require a high-order reasoning process. As Morris and Shin (2002) point out, for ε converging to zero, players behave as if they believe that the proportion of attacking agents has a uniform distribution in [0,1]. This can be interpreted as a Laplacian belief on aggregate behaviour.[8] It defines a unique threshold state Y^*, such that players attack if and only if $Y > Y^*$. Heinemann *et al*. (2002) show that this threshold is the unique solution to

$$Y(n - \hat{a}(Y) + 1) = nT.$$

This interpretation does not require private information and may serve as a theory of equilibrium selection in the game with public information.

Different theories of equilibrium selection predict different critical states at which attacks occur. Harsanyi and Selten (1988) suggest the payoff-dominant equilibrium, which is the combination of strategies that yields the highest payoff for all traders. This theory says that players attack whenever a coordinated attack is rewarding, that is for $Y > T$. Payoff dominance is a normative concept that requires coordination on efficient strategies. It does not consider risk associated with strategic uncertainty.

The Maximin strategy is the one that yields the highest payoff in the case of the most unfavourable behaviour of other players. In our game, it is the strategy to attack if and only if a player is sure that

[7] In this issue, Oliver Board gives an account of the epistemic conditions behind the Nash equilibrium concept.

[8] For situations with insufficient information on probabilities, Laplace suggested a distribution that places equal probabilities on all possible events. Although this is very sensitive to the definition of events, it has become common to refer to uniform distributions as an expression of the Laplacian suggestion.

a devaluation will occur, that is for $Y > \bar{Y}$. This theory assumes players to be extremely pessimistic about the behaviour of others and extremely risk averse.

Theories that consider the risk associated with strategic uncertainty predict thresholds somewhere in between T and $\bar{Y}$. The consideration expressed in the theory of global games is just one of several conceivable theories that meet this task.

The risk-dominant strategy, also defined by Harsanyi and Selten (1988), is the best reply of a player who believes that other players believe that the probability of success has a uniform distribution in [0,1]. The associated threshold to attack is some intermediate state close to the solution of the global game[9] and formally defined by the unique solution to

$$Y\left[1 - Bin\left(\hat{a}(Y) - 2, n-1, 1 - T/Y\right)\right] = T.$$

Another way to model Laplacian beliefs is to assume that a player attaches probability ½ to every other player attacking. In contrast to the solution of the global game, this is a belief about individual behaviour. It predicts a threshold, formally given by the unique solution to

$$Y\left[1 - Bin\left(\hat{a}(Y) - 2, n-1, ½\right)\right] = T,$$

that is close to the state where half of all players are needed for success.

A more general approach is a theory that assumes individual beliefs to differ according to some distribution. The critical state is then defined by a player who believes that at the critical state other players attack with probability p. Parameter p can be estimated in experiments using maximum-likelihood techniques.

A possible answer to the question of which theory gives the best prediction of actual behaviour can be found in experiments that test these theories. It is an empirical question as to whether public information destabilizes an economy by creating self-fulfilling beliefs or whether any of the selection theories gives a good prediction. Properly designed laboratory experiments with real subjects can measure the impact of strategic uncertainty and illuminate the process by which beliefs are formed and altered by experience and whether and how players achieve coordination.

IV. AN EXPERIMENT ON THE EXCHANGE-RATE-ATTACK GAME

Heinemann *et al*. (2002) present an experiment that imitates the speculative attacks model described above and compare behaviour in sessions with common (public) and private information. In each session there are 15 students interacting with each other. During a session subjects play 16 periods. In each period they are faced with ten independent situations in which they have to decide between two alternatives, A and B. The experiment avoids any connotation that might be associated with 'speculation' or 'attacks'. Action A is introduced as a secure alternative, yielding a positive and constant payoff T that may be interpreted as avoided costs of a speculative attack. Action B is the risky action and can be interpreted as attack, yielding a payoff of Y if the number of subjects choosing B exceeds a hurdle $a(Y)$, and zero otherwise.

In each situation, state Y is randomly selected with a uniform distribution from the interval [10, 90]. In sessions with common information (CI), all subjects are informed about Y and know that all of them share the same information. In sessions with private information (PI) each subject receives a private signal X^i. Signals are randomly selected with a uniform distribution from the interval $[Y-10, Y+10]$ for each player, separately. Here, subjects know that all of them get only private information and they share knowledge of the underlying random process. Each session consists of two stages, each with eight independent periods. In each period, subjects are confronted with ten situations that differ by randomly selected states and signals (in PI-condition).

A subject who chooses A gets a payoff of T with certainty. The payoff for B is Y, if at least $a(Y) = 15(80 - Y)/Z$ subjects choose B, and zero otherwise. The formula is given in the instructions, but also explained via an example and a table. The hurdle to

[9] In games with two players and two actions, the risk-dominant equilibrium and the global-game solution coincide.

Figure 2
Number of *B*-players Needed for Success of Playing *B* at State *Y*

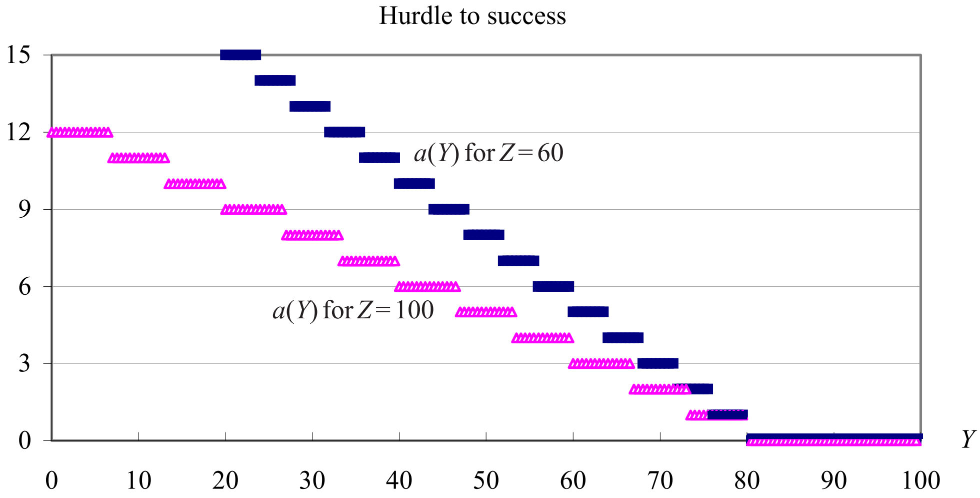

success $a(Y)$ may be interpreted as the number of traders who must attack to enforce a devaluation. The rules of this game precisely match those of the exchange-rate-attack model introduced in section II above. Z is a parameter that determines the steepness of this function. In our experiment, we use $Z = 60$ and $Z = 100$ in different sessions. These hurdle functions are shown in Figure 2. For the steeper function ($Z = 60$) the hurdle increases to 15 at $Y = 20$. This says that all 15 players must attack (choose B) to get a positive payoff from the attack if the true state is $Y = 20$.

The two stages of each session differ by the amount T that is paid for alternative A. In half of all sessions we started with $T = 20$ and switched to $T = 50$ in the second stage. In other sessions we reversed the order.

The experiment was run in computer laboratories using z-tree software (Fischbacher, 1999). Subjects could not communicate or see others' screens. In each period, subjects were first shown a table with ten situations: the left-hand side displayed state Y in the CI condition or private signal X^i in the PI condition for each of the ten situations. At the right-hand side, subjects had to decide between A and B by clicking on either of two boxes. There was no presetting. Decisions could be changed until subjects clicked an OK-button at the lower end of the screen.

Once all players had completed their decisions in one period, they were informed for each situation about Y, how many people had chosen B, whether decision B was successful or not, and their own payoff. Furthermore, they were reminded of their own signals (in PI-condition) and their own choices. Other information about other players was withheld. After all players had left the information screen, a new period started and information about previous periods could not be revisited. Subjects were allowed to take notes and many of them did.

After completion of all 16 periods, subjects answered a questionnaire and were then paid in private, converting their total points into cash. The experiment was run in Frankfurt and Barcelona and the payoffs were converted into local currency. The applied conversion rate in sessions with $Z = 60$ was 1 DM or 70 pta. for 200 points. In sessions with $Z = 100$ subjects received 1 DM for 250 points.[10] Average payment per subject varied across sessions from €14.30 to €22.50. A session took about 90 to 120 minutes. In total, there were 23 sessions with eight different conditions and 345 participating students. Table 1 gives an overview of the different sessions.[11]

[10] This condition has only been used in Frankfurt.
[11] More details, such as sample screens and instructions, can be found in Heinemann *et al.* (2002).

Table 1
Session Overview

Z	Secure payoff *T*	Location	Number of sessions with: public information	private information
100	1st stage 20/2nd stage 50	Frankfurt	1	1
100	50/20	Frankfurt	1	1
60	20/50	Frankfurt	1	2
60	20/50	Barcelona	3	3
60	50/20	Frankfurt	2	2
60	50/20	Barcelona	3	3
Total number of sessions			11	12

V. EXPERIMENTAL EVIDENCE ON THE EFFECTS OF PRIVATE VERSUS PUBLIC INFORMATION

This section summarizes the main results of the experiment by Heinemann *et al*. (2002) to the extent needed to answer the questions of whether public information destabilizes an economy by creating self-fulfilling beliefs and how it influences the likelihood of currency crises.

(i) Threshold Strategies

Most subjects chose *A* for low states or signals and *B* for high states or signals. In the first period 72 per cent of all subjects behaved in a way that is consistent with a personal threshold, so that a subject attacks if and only if the state (in CI-games) or his or her signal (in PI-games) is above this threshold. The proportion of subjects who applied threshold strategies gradually rises up to 96.5 per cent in the last period (see Figure 3).

The number of players using threshold strategies does not depend on the information condition. On the one hand, it is not surprising that most participants played threshold strategies, since the hurdle for success of *B* is decreasing in *Y* and the payoff to *B* in case of success is increasing. Common sense tells us that we should choose *B* for high states and *A* for low states or signals. On the other hand, deductive reasoning needs very strong assumptions to get this result: in games with private information, theory predicts threshold strategies but requires

Figure 3
Percentage of Subjects Whose Behaviour was Consistent with Undominated Threshold Strategies

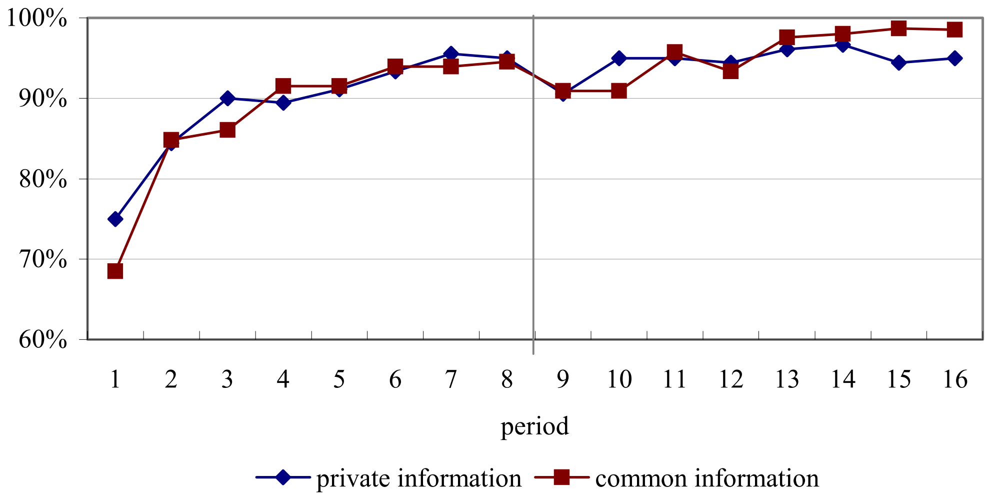

Figure 4
Percentage of Decisions in which Subjects Could Have Achieved a Higher Payoff by a Different Decision

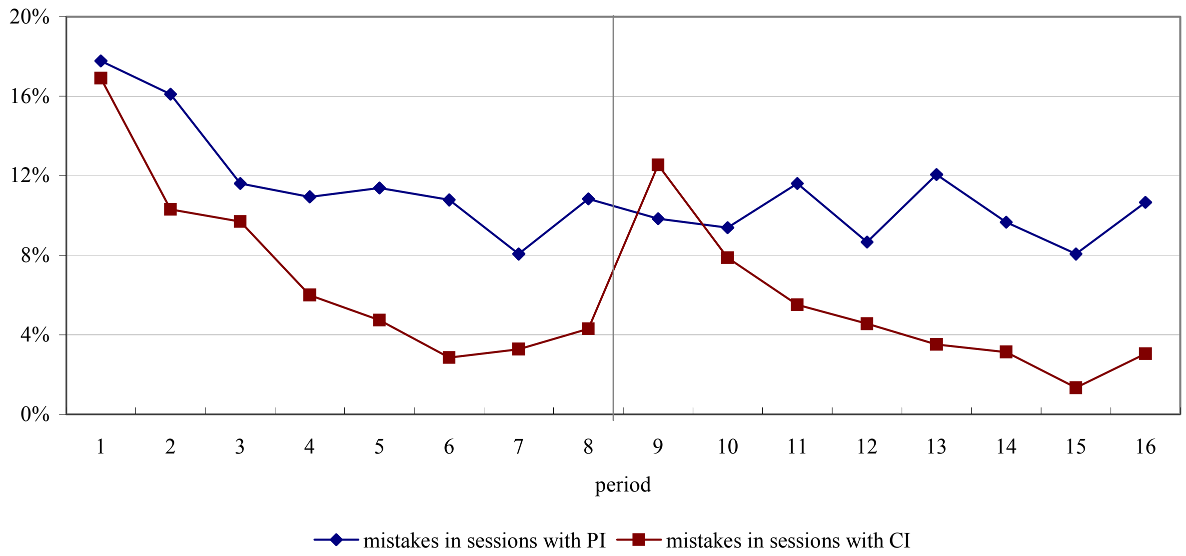

common knowledge of the game structure. As we know from other experiments, for example by Stahl and Wilson (1994) and Nagel (1995), most real subjects fail to reason more than at most three levels of beliefs over beliefs. In games with common information, non-threshold strategies may even occur in Nash equilibria.

The strength of threshold strategies lies in their robustness. If a subject expects others to play threshold strategies or randomize, his or her best response is a threshold strategy. Even though other strategies might form an equilibrium in common knowledge games, the best response to any reasonable belief deviating from common knowledge is a threshold strategy. As there is strategic uncertainty at least in the first rounds of a treatment, threshold strategies are a natural way to play. Once a sufficient number of subjects plays threshold strategies, the best response is again a threshold strategy. Other strategies may be an equilibrium under common knowledge, but they are not robust against even the slightest deviations from common knowledge.

(ii) Coordination

During the eight periods of a treatment, subjects coordinate their behaviour. In the information phases they see, for each selected state, how many players have chosen to attack and whether the attack succeeded or how close it was to success. Players who deviate from coordination lose payoffs either by waiving a potential profit from an attack or by attacking without success. If all subjects are perfectly coordinated, none of these *ex-post* mistakes occurs.

Figure 4 shows the average number of regrettable decisions per subject. In sessions with CI, it decreases over time as subjects coordinate their actions. Within a stage, most subjects coordinate on a common threshold within a few periods. Once coordination by a large number is achieved, the threshold never changes thereafter.[12] When the secure payoff to action A is changed in period 9, coordination failures are high again, but decrease until subjects find a new coordination point. The decreasing number of regrettable decisions shows that players gradually learn to predict whether an attack will be successful at a particular state.

In sessions with PI, it is much more difficult to predict precisely the actions of others, because they depend on random signals. Here, subjects can learn the probability with which an attack will be successful at a particular state and given a particular signal. Depending on personal risk aversion, subjects decide on different thresholds. Therefore in PI games, players cannot always avoid attacking without suc-

[12] Van Huyck *et al.* (1990, 1991) also observed high persistence of coordination points in a series of experiments on coordination games.

Figure 5
Frequency of *B* Play for Each of the 80 Randomly Selected States of the 8 Periods in Stage 1 of a PI Session

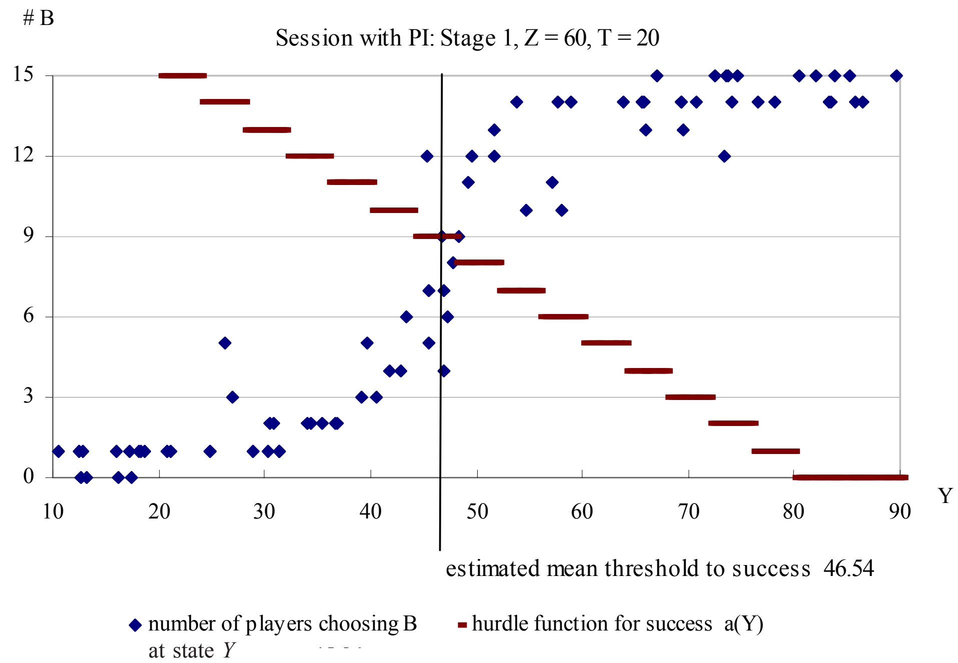

Note: All data points on or above the hurdle function are successful attacks. Largest *Y* without success for *B*, 47.72; smallest *Y* with success for *B*, 45.36.

cess or abstaining from an attack that is successful. They may get a particular signal that is much better or worse than the true state and thus leads them to a wrong decision. Due to the random process of private signals, the expected quota of mistakes is at least 6–7 per cent. Even with perfect coordination, subjects cannot avoid these mistakes. The observed quota of mistakes in sessions with PI is decreasing for the first three periods and bounces around 10 per cent for the remainder of a session (see Figure 4). Surprisingly, there is no rise in the percentage of regrettable decisions with the change of treatments in period 9.

(iii) Aggregate Behaviour

During one stage of a session, 80 different states are randomly selected from the interval [10,90]. For each state, we see how many subjects chose action *B*. Combining these data gives a very intuitive picture of aggregate behaviour. The higher the state, the larger the number of subjects who play *B*. The proportion of attacking agents looks like a cumulative distribution function and its intersection with the hurdle function determines a critical state below which most attacks fail and above which they mostly succeed. Figure 5 shows aggregate data of frequency of *B* play for each randomly selected state from one stage of a session with private information.

Heinemann *et al*. (2002) estimate the mean thresholds to successful attacks for each session and treatment by calculating the mid-points between the largest state without success for action *B*, and the lowest state at which *B* is successful. This gives them 46 data points as measures for the session-specific thresholds at which attacks start being successful.

(iv) Probability of Crises

The higher the threshold from which point on attacks are successful, the smaller the set of states at which attacks succeed, and the lower the prior probability of an attack being successful. Therefore, thresholds to successful attacks are directly connected to the probability of crises.

Table 2
Estimated Coefficients of Controlled Variables and According t-values

Intercept γ_0	T γ_1	Z γ_2	TZ γ_3	Loc γ_4	$Info$ γ_5	Ord γ_6	TO γ_7	R^2
			(t-values)					Adjusted R^2
43.57	0.878	–0.3825	–12.89	1.35	2.45	5.10	–4.29	0.89
(11.12)	(10.92)	(8.23)	(–5.26)	(1.27)	(2.63)	(3.87)	(–2.31)	0.87

Heinemann *et al*. (2002) use a linear regression to analyse how these thresholds depend on controlled variables. The employed regression equation is

$$Y^* = \gamma_0 + \gamma_1 T + \gamma_2 Z + \gamma_3 TZ + \gamma_4 Loc + \gamma_5 Info + \gamma_6 Ord + \gamma_7 TO + u,$$

where dummy variables are introduced: $Loc = 1$ for sessions in Frankfurt, $Info = 1$ for sessions with PI, $Ord = 1$ for sessions starting with $T = 20$, $TZ = 1$ if and only if $T = 50$ *and* $Z = 60$, and $TO = 1$ if and only if $T = 50$ *and* $Ord = 1$. TZ and TO are interaction variables and u is the residual. Table 2 gives the results of this regression based on the analysis of 46 data points, one for each stage in each session.

Controlled variables explain 89 per cent of the data variation. The payoff to the secure action T has the largest impact. With higher opportunity costs, thresholds to attack rise significantly. A higher hurdle to success (lower Z) also increases thresholds. These effects are in line with the theory of global games and may be interpreted as the effectiveness of transaction costs and capital controls as means of reducing the probability of currency crises.

The most interesting result is the effect of information. With private information, mean thresholds tend to be 2.45 units higher than with common information. This means that public information reduces the thresholds to attack. In the random environment of the experiment, this is associated with a rise in the probability of successful attacks by 3.1 per cent owing to public information. The difference is small but statistically significant.

The location dummy is not significant, but the order dummy is. In sessions starting with $T = 50$, thresholds tend to be lower than in sessions with the reverse order. This can be explained by a numerical inertia in increments that subjects use to find their personal thresholds. The strong impact of interaction variables is due to the non-linear structure of the payoff function.

(v) Predictability of Attacks

The main objection against public information comes from fears that behaviour could be driven by self-fulfilling beliefs. Multiplicity of equilibria supports the view that speculative attacks are hardly predictable. In the experiment, however, predictability of attacks was higher in sessions with CI than in sessions with PI.

Heinemann *et al*. (2002) estimate thresholds separately for sessions with common and private information. In both subsets, more than 90 per cent of data variation can be explained by other controlled variables. But, the standard deviation of residuals is still 2.84 in sessions with CI and 4.63 in sessions with PI. If thresholds have a normal distribution, they can be predicted within an interval of four standard deviations with a probability of 95 per cent. We see clearly that this interval is much larger with private than with common information.

Within a session, predictability of attacks given past observations is also higher with common than with private information. This is partly due to randomness of private signals and partly to differences in individual strategies, as subjects cannot fully coordinate their actions when they do not know other players' signals. With common information, subjects coordinate on a common threshold within a few periods. Often this threshold can be identified as one of the steps of the hurdle function shown in Figure 2 above.

VI. PROS AND CONS OF PUBLIC INFORMATION

A central bank may commit to publication of information that is relevant to prediction of its reactions to the foreign-exchange market. Y is the size of a potential devaluation and it also determines whether the central bank will keep the peg or not for any possible size of an attack. A transparent central bank shares its own information and its goals with the public and thereby communicates its potential reactions to markets.

Theory suggests that with public information, traders share knowledge of the payoff function and attacks may be triggered by self-fulfilling beliefs. With private information, there is a unique equilibrium in threshold strategies, but actual behaviour is determined by private signals that are randomly chosen and cannot be controlled by the central bank. This puts a limit on the predictability of attacks in an environment with private information.

Experimental evidence shows that real players behave very similarly in situations with common and private information.[13] Strategic uncertainty is the prevailing source of risk and subjects deal with this uncertainty in similar ways for both information conditions. They all use threshold strategies and average behaviour in groups of 15 players is fairly predictable. Public information helps players to coordinate their strategies and thereby reduces efficiency losses from uncoordinated actions.

Predictability of attacks seems higher with public than with private information. The variation of coordination points in sessions with public information is smaller than the variation of critical states to successful speculative attacks in sessions with private information. Thus, predictability of attacks is higher with public information. This holds for predictions conditioned on the state (predictions by the central bank) as well as for predictions conditioned on experience by players (predictions by market participants).

On the one hand, transparency increases predictability and reduces efficiency losses from uncoordinated activities. On the other, public information raises the probability of attacks: the set of states for which the peg can be maintained is smaller with public than with private information. Public information not only helps players to coordinate, but also seems to raise their trust in the ability of a group to coordinate on a fairly efficient strategy. Players attack at states at which they do not yet attack when they have only private information. Public information moves the threshold down towards the payoff-dominant equilibrium. Even though this is an efficient outcome for players, one must bear in mind that speculative traders play against the central bank and an efficient coordination on attacks may not be in the interest of the central bank. In order to maximize the set of states for which the peg can be maintained, the central bank has an incentive to withhold information and leave traders to rely on less precise outside information.

VII. TESTING THEORIES OF EQUILIBRIUM SELECTION

For both information conditions, thresholds to crises are fairly predictable. This raises the question of whether and which equilibrium theories succeed in predicting observed thresholds. In the game with private information there is a unique equilibrium with a threshold X^*, such that players attack if and only if they get a signal above X^*. In the game with common information there are various refinement theories predicting different thresholds Y^*. In the experiment, individual thresholds differ between subjects. Heinemann *et al*. (2002) calculate the means of individual thresholds for each session and test whether these means accord with theoretical predictions of various refinement theories. A comparison of theoretical equilibrium thresholds with average observations is summarized in Table 3.

In sessions with PI, observed thresholds are close to the unique equilibrium in treatments with $T = 20$. But in treatments with $T = 50$, observed thresholds are always clearly below the theoretical prediction. The hypothesis that subjects play equilibrium thresholds could be rejected in a two-sided F-test. However, observations followed the qualitative comparative

[13] Cabrales *et al*. (2002) conducted another experiment on global games. In their design, there were only five possible states and signals and iterative elimination of dominated strategies singled out the risk-dominant equilibrium with only four steps of elimination in the PI-condition. Behaviour converged towards this strategy for both common and private information.

Table 3
Observed Mean Thresholds of Individual Strategies and Theoretical Equilibrium Thresholds

Treatment	$T = 20, Z = 100$	$T = 20, Z = 60$	$T = 50, Z = 100$	$T = 50, Z = 60$
Private information game:				
Observed mean threshold	29.76	41.96	55.22	57.02
Unique equilibrium	32.36	41.84	60.98	66.03
Common information game:				
Observed mean threshold	26.30	37.84	52.95	52.56
Payoff-dominant equilibrium	20	20	50	50
Global-game solution	33.33	44.00	60.00	64.00
Risk-dominant equilibrium	34.55	44.00	62.45	67.40
Maximin strategy	73.33	76.00	73.33	76.00
Laplacian beliefs on individual behaviour	33.07	48.00	51.48	56.00
Subjective belief $p = 2/3$ for every other subject attacking	23.52	40.00	50.04	52.00

statics of the unique equilibrium in response to changes in T and Z.

In sessions with CI, observed thresholds are always somewhere between the payoff-dominant equilibrium and the solution of the global game with diminishing variance. In treatments with $T = 20$, observations are closer to the global-game solution; in treatments with $T = 50$, they are closer to the payoff-dominant equilibrium. The reason is the difference in the hurdle to success at critical states for these two treatments. To coordinate on and sustain a lower threshold, a larger number of subjects is needed to play B at critical states. Payoff dominance neglects the associated risk of coordination failure that is higher for states close to 20 than for states close to 50. The global-game solution seems to overestimate the effect of the hurdle function. Subjects seem to have more trust in efficient outcomes than is expressed by a Laplacian belief on the proportion of attacking agents. Two-sided F-tests reject both payoff dominance and the global-game solution. The risk-dominant equilibrium predicts even higher thresholds than the global-game solution and can also be rejected. The maximin concept is clearly off track.

Laplacian beliefs on individual behaviour define a threshold that is a best response to the belief that each other agent independently randomizes between A and B with probability ½. The resulting threshold deviates from the global-game solution into the direction where the hurdle to success requires half of all players to attack. In the experiment, it could also be rejected, although it gave a good prediction for treatments with $T = 50$. But for $T = 20$ it overestimates thresholds. It seems to rely on a too pessimistic assumption on beliefs.

All of these theories are non-parametric. They can either fit the data or be rejected, but there are no parameters representing the degree of strategic uncertainty. Parametric theories allow us to estimate one or more parameters in such a way that theoretical predictions give a better fit to observed data. If the belief that other players attack with probability ½ is too pessimistic, a higher probability may better explain observations.

Generalizing the theory of independent beliefs, one may ask whether there is any p such that subjects behave as if they believe other players choose B with probability p. For $p = 1/2$ this is the Laplacian approach that could be rejected. Trying different values for p, we find that this theory cannot be rejected for $p \in (0.6, 0.7)$. Table 3 above gives the equilibrium thresholds consistent with each player believing that others independently attack with probability 2/3. It is pretty close to actual observations in sessions with $Z = 60$. For the sessions with a low hurdle function ($Z = 100$), it predicts a threshold that may be too close to payoff dominance. Not rejecting this theory may be due to the limited range and number of observations. There are other parametric theories that may give a better explanation. This, however, requires additional experiments of similar structure to get a more precise measure of strategic uncertainty and discriminate between different parametric theories.

Besides numerical tests, one can also compare comparative static properties of refinement theories with observations from the experiment. The payoff-dominant equilibrium predicts the threshold to be *T*, independent of the hurdle parameter *Z*. Maximin strategies are independent of *T* and are exclusively determined by the hurdle function. Observed thresholds, however, react to both parameters. Thresholds rise with rising *T* and with a rising hurdle (lower *Z*). The other refinement theories predict changes in the direction we observed. Even if the global-game solution or risk dominance cannot be used as a numerical predictor, these theories predict the right qualitative comparative static results in response to changes of the payoff function.

The theory of global games suggests treating strategic uncertainty as random differences in private beliefs. In this view, the attacks game with CI may be seen as a game with PI, where the variance of information is smaller than in a game where information is really private. In games with PI, the dispersion of individual beliefs adds to the variance in the random process. Heinemann and Illing (2002) show that in the global game the probability of a successful attack rises with rising variance of private information. In the experiment, the probability of attacks is higher with common than with private information. This is another reason to reject the hypothesis that strategic uncertainty can be measured by a variance in private beliefs.

VIII. CONCLUSION

This paper compared theoretical and empirical results on the exchange-rate-attack model that is a leading example for coordination games with strategic complementarities. The game has multiple Nash equilibria when players have public information on the payoff function. If players have only private information on some payoff-relevant parameter, there is a unique equilibrium. This raises the question of whether public information may be destabilizing an economy by evoking self-fulfilling beliefs.

The experiment by Heinemann *et al*. (2002) mimics the exchange-rate-attack game. There are sessions with public information and otherwise equal sessions with private information. There is no evidence that behaviour is less predictable with public than with private information. The destabilizing effect that public information may have, owing to the possibility of self-fulfilling beliefs, is less severe than theory predicts. With public information, subjects can better coordinate their behaviour and reach an equilibrium that is more efficient than the outcome in sessions with private information. For the exchange-rate-attack game, public information increases the probability of successful attacks, but also increases predictability and the degree of coordination.

The observation that public and private information lead to similar results in laboratory experiments may be explained by bounded rationality of real players. Assuming full rationality leads to an apparent discontinuity in the set of equilibria when the variance of private information goes to zero. For a small positive variance of private signals, there is a unique equilibrium. But, if the variance of private signals is zero and all players share the same information, there is a wide range of states for which both attack and not-attack are equilibria.[14] The discontinuity at this point is due to the assumption that players are rational and rationality is common knowledge. Players are assumed not to play a dominated strategy. Furthermore, they are assumed to believe that other players do not play a dominated strategy, that others believe others not to play a dominated strategy, and so forth. The set of equilibria can be defined by iterated elimination of dominated strategies and the assumed mental process requires an infinite number of levels of players' reasoning about each others' beliefs. If players have a limited number of levels of reasoning about each others' beliefs, the set of consistent strategies is continuous at $\varepsilon = 0$. With *k* levels of reasoning, the region of indeterminacy is an interval in the state space with borders that are at

[14] Hellwig (2001) shows that the set of equilibria is continuous along another dimension, called common p-beliefs, where p = 0 is the case of private information and p = 1 is public information.

most $(1 + k)\varepsilon$ units to the interior of $[T, \bar{Y}]$, which is the multiplicity region of the game with public information. With $\varepsilon \to 0$ there is a continuous approach to full indeterminacy. This may explain why experimental evidence is similar for public and private information, but it does not explain why observations are so easy to predict.

Apparently, real subjects have ways to deal with strategic uncertainty that are fairly predictable on some aggregate level. An explanation might be provided by heuristic models of belief formation. Theories of equilibrium selection such as payoff dominance, risk dominance, or maximin strategies, might serve this purpose. Although statistical tests could reject the hypotheses that average behaviour follows any of these theories of equilibrium selection, changes in exogenous parameters move observed coordination points in the same direction as predicted by refinement theories that consider strategic uncertainty, such as global games or risk dominance. Most importantly, most subjects use threshold strategies as required by these refinement theories.

Real subjects coordinate on equilibria somewhere in between the payoff-dominant equilibrium and the solution of the global game (or the risk-dominant equilibrium). Parametric theories of belief formation cannot be easily rejected, because they allow us to fit parameters to the data of the model. In our experiment, mean thresholds were close to optimal thresholds of a subject who believes that (at the critical state) every other player attacks with probability 2/3. However, there are other parametric theories of belief formation and the experiment produced too few data and was not really designed for discriminating between those theories. It is an open question as to whether there are any simple theories of belief formation that give a robust description of actual behaviour, valid for generic coordination games.

In most theories, it is assumed that all players have the same beliefs. Different theories select different equilibria, but each of them is defined by beliefs that are identical for all players. Real subjects may hold different beliefs that can be ordered according to the state or signal at which a subject switches between actions. The outcome of a coordination game is determined by beliefs of a marginal player. The marginal player is indifferent at the critical state, where the number of attackers equals the hurdle to success. However, the marginal player is not always the same. When only few players are needed to achieve a successful attack (e.g. for $T = 50$, see Figure 2 above), the marginal player is one of those who are more optimistic about the likely success of an attack. If the hurdle is higher at the critical state (e.g. for $T = 20$), the outcome is determined by more pessimistic subjects.

Strategic uncertainty has long been thought to be a Knightian uncertainty to which probability measures cannot be attached. Experiments such as the one described in this paper have shown that there are patterns of behaviour that allow us to test parametric theories of belief formation. It is a promising task for future research to construct and test theories of behaviour in situations of strategic uncertainty. While it is impossible to attach probabilities to strategic uncertainty by purely deductive theories, laboratory experiments may help us to develop measures of strategic uncertainty.

REFERENCES

Aumann, R. J., and Brandenburger, A. (1995), 'Epistemic Conditions for Nash Equilibrium', *Econometrica*, **55**(1), 1–18.

Brandenburger, A. (1992), 'Knowledge and Equilibrium in Games', *Journal of Economic Perspectives*, **6**(4), 83–101.

Cabrales, A., Nagel, R., and Armenter, R. (2002), 'Equilibrium Selection through Incomplete Information in Coordination Games: An Experimental Study', Working Paper, available at http://www.econ.upf.es/cgi-bin/onepaper?601.

Carlsson, H., and van Damme, E. (1993*a*), 'Global Games and Equilibrium Selection', *Econometrica*, **61**(5), 989–1018.

— — (1993*b*), 'Equilibrium Selection in Stag Hunt Games', in K. Binmore, A. Kirman, and P. Tani (eds), *Frontiers of Game Theory*, Cambridge, MA, MIT Press,

Fischbacher, U. (1999), 'z-Tree 1.1.0: Experimenter's Manual', University of Zurich, Institute for Empirical Research in Economics, available at http://www.iew.unizh.ch/ztree/index.php

Frankel, D., Morris, S., and Pauzner, A. (2001), 'Equilibrium Selection in Global Games with Strategic Complementarities', *Journal of Economic Theory*, forthcoming.
Fukao, K. (1994) 'Coordination Failures under Incomplete Information and Global Games', Discussion Paper Series A No. 299, The Institute of Economic Research, Hitotsubashi University, Kunitachi, Tokyo.
Harsanyi, J. C., and Selten, R. (1988), *A General Theory of Equilibrium Selection in Games*, Cambridge, MA, MIT Press.
Heinemann, F. (2000), 'Unique Equilibrium in a Model of Self-fulfilling Currency Attacks: Comment', *American Economic Review*, **90**(1), 316–18.
— Illing, G. (2002), 'Speculative Attacks: Unique Equilibrium and Transparency', *Journal of International Economics*, **58**(2), 429–50.
— Nagel, R., and Ockenfels, P. (2002), 'Speculative Attacks and Financial Architecture: Experimental Analysis of Coordination Games with Public and Private Information', Working Paper, available at http://www.sfm.vwl.uni-muenchen.de/heinemann/publics/hno-1.htm
Hellwig, C. (2001), 'Public Information, Private Information and the Multiplicity of Equilibrium in Coordination of Games', *Journal of Economic Theory*, forthcoming.
Morris, S., and Shin, H. S. (1998), 'Unique Equilibrium in a Model of Self-fulfilling Currency Attacks', *American Economic Review*, **88**(3), 587–97.
— — (1999), 'A Theory of the Onset of Currency Attacks', in P.-R. Agenor, M. Miller, D. Vines, and A. Weber (eds), *The Asian Financial Crisis: Causes, Contagion and Consequences*, Cambridge, Cambridge University Press, 230–55.
— — (2002), 'Global Games: Theory and Applications', forthcoming in M. Dewatripont, L. Hansen, and S. Turnovsky (eds), *Advances in Economics and Econometrics, the Eighth World Congress*, Cambridge, Cambridge University Press.
Nagel, R. (1995), 'Unraveling in Guessing Games: An Experimental Study', *American Economic Review*, **85**(5), 1313–26.
Obstfeld, M. (1986), 'Rational and Self-fulfilling Balance-of-Payments Crises', *American Economic Review*, **76**(1), 72–81.
Stahl, D. O., and Wilson, P. W. (1994), 'Experimental Evidence on Players' Models of Other Players', *Journal of Economic Behavior and Organization*, **25**(3), 309–27.
Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1990), 'Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure', *American Economic Review*, **80**(1), 234–48.
— — — (1991), 'Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games', *Quarterly Journal of Economics*, **106**(3), 885–910.