

Coordination and Learning in Dynamic Global Games: Experimental Evidence*

Olga Shurchkov[†]

Current Version: August 2012

Abstract

Coordination problems are ubiquitous in social and economic life. Political mass demonstrations, the decision whether to join a speculative currency attack, investment in a risky venture, and capital flight from a particular country are all characterized by coordination problems. Furthermore, all these events have a dynamic nature which has been largely omitted from previous experimental studies. Here I use a two-stage variant of a dynamic global game to study experimentally how the arrival of information in a dynamic setting affects the relative aggressiveness of speculators. In the first stage, subjects exhibit excess aggressiveness, which appears to be driven by beliefs about others' actions rather than an intrinsic taste for attacking. However, following a failed first-stage attack, subjects learn to be less aggressive in the second stage. On the other hand, the arrival of new, more precise information after a failed attack leads to an increase in subjects' aggressiveness. Beliefs, again, play a crucial role in explaining how the arrival of information affects attacking behavior.

Keywords: Coordination, Learning, Dynamics, Global Games, Crises, Beliefs, Excess Aggressiveness, Experimental Economics

JEL Classification: C7, C9, D8, F3

*I am especially grateful to George-Marios Angeletos and Ernst Fehr for their invaluable insights and support throughout this project. Casey Rothschild has offered a tremendous level of guidance in the analysis and application of the theory. In addition, I would like to thank Ernst Fehr and the University of Zurich for providing the resources and the financial support without which this study would not be possible. I would also like to thank Daron Acemoglu, Miriam Bruhn, James Costain, Florian Ederer, Muhamet Yildiz, Akila Weerapana, and the seminar participants and organizers of the MIT macroeconomics workshop, Wellesley College Calderwood seminar, the IESE Conference on Complementarities and Information, the 2007 SAET conference, the conference of the French Economic Association on Behavioral Economics and Experiments, the ESA World Meeting 2007, and the 22nd World Congress of the EEA for valuable comments and discussion. All remaining errors are my own.

[†]Department of Economics, Wellesley College, 106 Central St., Wellesley, MA, USA.
e-mail: oshurchk@wellesley.edu.

1 Introduction

Coordination amongst economic agents is an essential element in many macroeconomic events. The ability of individual participants to agree on a specific course of action such as an attack on a currency peg, a bank run, or a riot can determine the ultimate outcome for an economy as a whole and may change the course of a nation's history.

Speculative attacks can lead to an enormous negative impact on economic growth and can cause political change and turmoil.¹ Recent laboratory experiments have shed light on the relative importance of macroeconomic fundamentals and attacking costs during a one-time decision whether to attack (Heinemann, Nagel, and Ockenfels (hereafter, HNO) 2004). However, real-world speculative attacks have a dynamic element. Even if a policy intervention succeeds in curbing a given attack, speculators may coordinate again as new information about the state of the economy arrives over time.

This paper uses an experimental approach to explore how the arrival of information in a dynamic setting affects the relative aggressiveness of speculators.²

The starting point is a two-stage variant on a dynamic global game developed by Angeletos, Hellwig, and Pavan (hereafter, AHP), 2007. This model captures the features of currency crises that seem to be essential for understanding these issues: (1) the coordination element of a speculative attack that arises due to strategic complementarities in agents' actions, (2) the heterogeneity of expectations about the underlying economic fundamentals among the agents, and (3) the fact that the agents' beliefs about their ability to induce a regime change may vary over time.³ The model consists of a large number of agents (15 subjects in the experiment) and two possible regimes, the status quo and an alternative to the status quo. The game continues into the second stage as long as the status quo is in place. In each stage, each subject can either attack the status quo (i.e., take an action that favors regime change) or not attack. The net payoff from attacking is positive if the status quo is abandoned in that stage and negative otherwise. Regime change, in turn, occurs if and only if the percentage of subjects attacking exceeds a threshold $\theta \in \mathbb{R}$ that parameterizes the strength of the status quo.⁴ The parameter θ captures the component of the payoff structure (the "fundamentals") that is never common knowledge, as is customary in the global games literature (Carlsson and van Damme 1993a, 1993b; Morris and Shin 1998). In the first stage, each subject receives a private signal about θ . If the game continues into the second stage, subjects may or may not receive more a more precise private signal about θ . In each stage, subjects are asked to state their beliefs about the expected size of the attack.

¹See IMF 2000 for examples from the Asian economic crisis of 1997-1998.

²Relative aggressiveness is defined as the ability to coordinate on successful attacks.

³Common-knowledge models of crises capture the first feature but abstract away from the second and the third features (Obstfeld 1996). Static global game models provide insight into the second feature but fail to incorporate the importance of learning and updating over time (Carlsson and van Damme 1993a, 1993b; Morris and Shin 1998).

⁴Within the context of currency crises, the fundamental θ represents the strength of the currency peg or the ability of the central bank to defend the peg. The agents are the speculators deciding whether to attack the currency. The cost of attacking can be interpreted as the interest rate. This framework has been applied to several macroeconomic phenomena: see Goldstein and Pauzner (2004) and Rochet and Vives (2004) for bank runs; Corsetti, Guimaraes and Roubini (2003) and Morris and Shin (2004) for debt crises; Atkeson (2000) for riots; Chamley (1999) for regime switches; and Edmond (2008) for political change.

In the laboratory, I conduct multi-round experiments that vary the strength of the fundamentals, the cost of attacking, and the availability of information in the second stage. There are four basic takeaways from the experiments. First, individuals display excess aggressiveness. Second, the extra aggressiveness appears to be driven by beliefs about the aggressiveness of others rather than an intrinsic taste for attacking. Third, individuals appear to successfully learn to be less aggressive following failed attacks, but, fourth, this learning is supplanted by the provision of additional information in the second stage.

The first-stage results are consistent with the previous literature on static global games in that the size of the attack/the individual likelihood of attacking decrease in the fundamental/the private signal, as well as in the cost of attacking (HNO 2004). Furthermore, I document a systematic deviation from the theoretical prediction toward aggressive attacking behavior, regardless of the cost, which is also consistent with previous literature (HNO 2004). However, while previous studies find weak or no evidence of strategic reasoning, I show that 81.4 percent of subjects best respond to their stated beliefs about others strategies.⁵ In particular, subjects expect others to attack aggressively and, in turn, respond with aggressiveness.⁶ The strategic complementarity feature of the environment then rewards this behavior with more successful attacks, thereby partially validating the aggressiveness of the original beliefs and resulting in subjects maintaining the same high level of aggressiveness in subsequent rounds of play. This self-fulfilling nature of beliefs partially explains the relative unimportance of the cost of attacking and the persistence of aggressiveness of beliefs as the game is repeated. Subjects not only believe that the status quo is more likely to be successfully overturned than the theory predicts, but they also experience this to be the case.

In contrast to the static global game framework, the two-stage dynamic setting allows me to observe how the arrival of new information affects subjects' ability to coordinate. If the experiment continues into the second stage in the no new information treatment, subjects learn only that the game has not ended. In the model, this implies that the fundamentals are sufficiently strong and that subjects should refrain from attacking again. Indeed, subjects exhibit this type of learning, since their aggressiveness is greatly reduced in the second stage of this treatment. Furthermore, the reduction in aggressiveness is matched by a significant decrease in the size of the believed attack relative to the first stage.

In the new information treatment, subjects receive a new, more precise, private signal in the second stage. In this case, subjects are still able to learn by observing that the game did not end after the first-stage attack, but they can also now learn by incorporating this more precise information into their decision of whether to attack the status quo. In the model, the arrival of this type of precise new information may result in an additional equilibrium where a new attack becomes possible. Experimentally, the arrival of new information significantly increases the probability of attack in the second stage relative to the treatment without new information. Learning, induced by the observation of a failed attack, alone makes subjects relatively less aggressive, but a new attack may become possible as the participants accumu-

⁵For example, HNO (2009) find that subjects often fail to best respond to stated beliefs. Costa-Gomes and Weizsäcker (2008) compare stated beliefs with actions in one-shot 3x3 games and find best response rates ranging from 55% to 73%.

⁶Since belief formation is unobservable by nature, I refrain from making conjectures about the sources of these aggressive beliefs. This type of "Level 2 reasoning" has been documented in other settings. See Cornand and Heinemann (2010) for an example with public information in beauty contest games.

late new information about the strength of the regime. Together, the second-stage findings imply that a policy-maker, having previously successfully defended the regime, cannot be assured that a crisis is averted.

Although there is a rich experimental literature on coordination games, this paper is the first to study a dynamic global coordination environment in an experimental setting. This study is also the first to characterize the role of self-fulfilling expectations in driving persistently aggressive attacking behavior. Furthermore, this paper shows that, while a failed attack can significantly reduce the aggressiveness of the speculators, the arrival of new private information about the state of the fundamentals can provoke new attacks to occur.

This paper is related to the empirical literature on coordination games which begins with experiments using common knowledge in a static environment (Cooper et al. 1990, 1992; Van Huyck, Battalio, and Beil 1990). Several studies have explored the predictions of common knowledge games in a dynamic environment (Cheung and Friedman 2009; Brunnermeier and Morgan 2010). Experimental papers that are most closely related to the present study test the predictions of static coordination games with private information (global games). Cabrales, Nagel, and Armenter (2007) study two-person games with random matching inspired by Carlsson and van Damme (1993a), while HNO (2004, 2009) examine the static speculative-attack model of Morris and Shin (1998). These papers focus on the static elements of coordination and do not provide insights into what happens after an unsuccessful attack or upon arrival of new information over time. Another closely related body of literature examines dynamic coordination games (Costain, Heinemann, and Ockenfels 2007; Schotter and Yorulmazer 2009). These papers provide important conclusions about the outcomes of sequential-move coordination games. However, they do not incorporate the effects of the arrival of new private information over time.

In addition, several studies have used field data to test the predictions of static global games (Prati and Sbracia 2002; Chen, Goldstein, and Jiang 2010; Danielsson and Peñaranda 2011). I view these studies as complementary to the experimental approach, although the data in these studies do not enable the researchers to identify the causal effects of the strength of the fundamentals, the cost of attacking, or the arrival of new information.

The rest of the paper is organized as follows. Section 2 discusses the experimental design: the theoretical environment and the parameters implemented in the lab, the treatments, and the procedures. Section 3 describes the forces behind attacking behavior in the first stage of the experiment, including the extent of subjects' ability to best respond to others' strategies and beliefs about others' strategies. Section 4 provides evidence that subjects learn to be less aggressive in the second stage of the no new information treatments. Section 5 focuses on the effects of the arrival of new information on subjects' aggressiveness. Section 6 concludes and discusses potential implications of the results.

2 Design of the Experiment

The experiment is based on a two-stage version of the model developed by AHP (2007).⁷ Below, I first briefly describe the theoretical environment that is parameterized and implemented in the laboratory. Next, I summarize the experimental treatments and procedures.

2.1 The Environment and Parameterization

In each round of the experiment, there are two possible regimes: the status quo and the alternative to the status quo. At the beginning of each round, a random number θ is drawn from a normal distribution $N(z, \sigma_z)$ which defines the initial common prior.⁸ θ parameterizes the exogenous strength of the status quo (or the quality of the economic fundamentals). A low value of θ represents a relatively weak state of the fundamentals, and a high value of θ represents a relatively strong state of the fundamentals.

Actions, Outcomes, and Payoffs. There are $N = 15$ agents⁹, indexed by i , simultaneously deciding between two possible courses of action. Agent i can either choose action A (“attack”), an action that favors regime change, or choose action B (“not attack”), an action that favors the status quo. An “attack” can be interpreted as a speculative run on a currency, large withdrawal of funds from the economy’s financial sector, or a political uprising against the government. Let us denote the regime outcome by $R_{t+1} \in \{0, 1\}$ where $R_{t+1} = 0$ refers to the survival of the status quo, while $R_{t+1} = 1$ refers to the collapse of the status quo. Similarly, the action of an agent is denoted by $a_{it} \in \{0, 1\}$, where $a_{it} = 0$ represents action B (“not attack”), while $a_{it} = 1$ represents action A (“attack”).

Action A is associated with an opportunity cost c . If action A is successful (i.e., $R_{t+1} = 1$), each agent choosing action A earns an income of $y = 100 > c$. If not (i.e., $R_{t+1} = 0$), then the agent choosing action A earns $0 < c$. Action B yields no payoff and has no cost. Hence, the utility of agent i is

$$u_{it} = a_{it}(yR_{t+1} - c).$$

Finally, the status quo is abandoned ($R_{t+1} = 1$) if and only if

$$A_t \geq \theta$$

where $A_t \equiv \frac{1}{N} \sum_{i=1}^{15} a_{it} \in [0, 1]$ denotes the mass of agents attacking at time t (*the aggregate size of the attack*).

If the regime change occurs ($R_{t+1} = 1$), there are no further actions to be taken in stage 2. However, if the status quo survives ($R_{t+1} = 0$), the agents again have the opportunity to attack the status quo in stage 2.

⁷Since I do not develop new theory in this paper, we refer the readers to AHP (2007) for detailed proofs of the propositions. For the sake of experimental tractability, I choose the simplest two-period variant of the multi-period model that can also be found in AHP (2007). Online supplemental materials to this paper define several auxiliary objects and characterize the equilibrium/equilibria in more detail.

⁸Note that z can be thought of as a public signal about θ that all agents receive.

⁹AHP (2007) assume a continuum of agents which enables them to produce closed-form solutions for the equilibrium. The conclusions in this paper are robust to using either the infinite N or the $N = 15$ case for the theory predictions.

Complementarity. Note that the actions of the agents are strategic complements, since it pays for an individual to attack if and only if the status quo collapses and, in turn, the status quo collapses if and only if a sufficiently large fraction of the agents attacks.

Information. In this setup (as in any global game), agents have heterogeneous information about the strength of the status quo. In particular, θ is never common knowledge, but each agent receives a private signal $x_{it} = \theta + \frac{1}{\sqrt{\beta_t}}\varepsilon_{it}$, where $\varepsilon_{it} \sim N(0, 1)$ is i.i.d. across agents and independent of θ and β_t is the precision of private information.¹⁰

Table 1 records the remaining parameters by session.¹¹

2.2 Treatments

Six sessions of the experiment were conducted at the experimental laboratory at the University of Zurich. Each session consisted of 40 independent rounds of play, with two stages in each round. There were different treatment conditions based on the cost of action A and on the information provided to the participants in the second stage of the experiment. The various treatment conditions are summarized in Table 2. In sessions 1 and 2, the experiment began with the Cost 20 (hereafter C20) treatments, followed by the Cost 50 (hereafter C50) treatments. In order to test for any potential order effects, the cost treatment order was reversed in sessions 3 and 4. Sessions 1-4 maintained the same information structure: no new information in the second stage. Sessions 5-6 maintained the cost at 60 (hereafter C60) and focused on changing the information structure. In particular, subjects in session 5 received an additional private signal in stage 2 during the last 20 rounds of the experiment. The order of the information treatments was reversed in session 6.

2.3 Procedures

Subjects in all sessions were students at the University of Zurich. The general procedures were kept the same throughout all six sessions. All sessions were computerized using the program z-Tree (Fischbacher, 2007). Upon the completion of the informed consent forms, subjects received paper copies of the instructions. Questions were answered in private, and subjects could not see or communicate with one another. At the end of the experiment, each participant filled out a computerized questionnaire, which asked subjects about their strategies, as well as their understanding of statistics and probability.¹² The final income of each subject was first given in points and then converted to Swiss Francs (CHF). The

¹⁰The information structure is parameterized by $\beta_t = \sigma_{x,t}^{-2}$ and $\alpha = \sigma_z^{-2}$, the precisions of private and public information, respectively, or equivalently by the standard deviations, $\sigma_{x,t}$ and σ_z . Thus, $\alpha + \beta_t$ is the overall precision of information. Subjects know the values of z , α , and β_t .

¹¹The prior about θ , z , was chosen to be high enough that a new attack becomes possible with the arrival of new information in the second stage (see Section 5). At the same time, in order to get a reasonable number of random draws within the critical interval of $[0, 100]$, I kept z sufficiently high and α sufficiently low. The standard deviation, β_1 , was chosen based on satisfying the criterion for stage-one uniqueness of equilibrium, namely $\beta_1 \geq \frac{\alpha^2}{2\pi}$. The standard deviation, β_2 was chosen to be sufficiently high to produce an equilibrium with a new attack in stage 2 (see section 5).

¹²Copies of the consent forms, instructions, and questionnaire questions in German or English are available upon request.

average income, including the show-up fee of about \$12 (15 CHF) across all sessions was about \$41 (51.6 CHF). Each session lasted approximately 1.5 hours.

Each of the six experimental sessions had 30 participants divided randomly into two groups of fifteen people. Each session consisted of 40 rounds of play, with each round corresponding to a new random number θ drawn from a normal distribution $N(z, \sigma_z)$.¹³ Thus, one can interpret each round as a new economy parametrized by the state of fundamentals, θ . Subjects were informed of the mean and the standard deviation of this distribution in the instructions. In addition, at the beginning of the round, each subject received a hint number (private signal, x) about the random number θ . In the instructions, subjects were informed that the hint number was centered around the true value of θ and were given its standard deviation ($\sigma_{x,1}$).

Each round consisted of one or two stages of decision-making. In each stage, each subject had to decide between actions A or B as described in section 2.1, yielding data on individual *action* as well as the aggregate *size of the attack*. Once all subjects had chosen their actions in each stage of every round, they were asked about their *beliefs*: “How many other members of your group do you think chose action A?”¹⁴ Next, each subject received the following information. If the game ended after stage 1, he or she found out that action A was successful,¹⁵ learned the value of the unknown number θ , the total number of subjects choosing action A, and the payoff for the round. If the status quo survived such that the game continued into the second stage, two scenarios were possible. In the treatment without new information (hereafter NNI), the subject only got a reminder of his or her first-stage hint number and received a notification that action A was not successful. In the treatment with new information (hereafter NI), subjects received a new, more precise hint number if the game continued into stage 2. Analogously to the first stage, subjects were informed that the second-stage hint number was centered around the true value of θ and were given its standard deviation ($\sigma_{x,2}$) in the instructions. At the end of the second stage, subjects learned whether action A was successful, the value of the unknown number θ , the total number of subjects choosing action A in the first and in the second stage, and the payoff for the round.

Summary statistics for the entire experiment are reported in Table A1 of Appendix A.

3 Aggressive Despite the Cost: First Stage Behavior

One would expect the cost to figure prominently in an individual’s decision to attack, since the cost is borne if and only if the subject chooses to attack. However, there are other factors

¹³Subjects were provided with several examples that familiarized them with the normal distribution to ensure their full understanding

¹⁴The belief variable takes on values 0-14. Belief elicitation was not incentivized. While belief accuracy is significantly higher when beliefs are incentivized (Gächter and Renner 2010, Wang 2011), incentivizing belief elicitation may allow risk-averse subjects to hedge with their stated beliefs against adverse outcomes of the other decisions (Blanco et al. 2008). Furthermore, in public goods experiments, incentivized beliefs tend to lead to further deviations from equilibrium play (higher contribution levels) than either non-incentivized beliefs or no beliefs at all (Gächter and Renner 2010). Since I find that in this experiment reported beliefs are highly correlated with actions, accuracy seems to be high even without incentives.

¹⁵Recall that the net payoff from attacking is positive if the status quo is abandoned and negative otherwise, which is the reason I choose to use the language of “successful” attacks.

that may play a role as well: the private signal which allows subjects to learn about the strength of the unobserved fundamentals, the subject’s expectations about others’ behavior, and the knowledge that the game would be repeated in the future. In this section, I confirm the results of previous studies, such as HNO 2004, which find that subjects follow threshold strategies, attacking whenever the fundamental state or signal is beyond some critical state or signal. Moreover, my data on subjects’ expectations enable me to provide new insights on the importance of self-fulfilling beliefs for coordination in this environment. In particular, I show that subjects attack more aggressively than the theory predicts in the high-cost treatments. This aggressiveness is not a result of learning to coordinate across rounds, but rather mostly due to a pre-existing belief that others will be more aggressive. In fact, subjects’ expectations that others will attack trump cost considerations. Finally, I take a closer look at whether this type of behavior is consistent with best-response strategies that characterize this environment.

3.1 Determinants of Attacks in the First Stage

Given the environment described in the previous section, there are three key determinants of attacking behavior in the first stage: the strength of the status quo, the cost of attacking, and beliefs about the size of attack. In addition to these within-round determinants that stem from the AHP (2007) model, the experiment also introduces a role for learning between rounds through playing experience. Below, I discuss the relationship between these factors and (1) the size of the attack and (2) the individual probability of choosing action A.

First, the model of AHP (2007) implies that the attack size, A , should decrease monotonically in θ , the unobserved true strength of the status quo. Figure 1 demonstrates that the data are consistent with this prediction. In particular, the figure plots the fitted lines from a locally weighted kernel regression of the average size of the attack for different realizations of θ for all three cost treatments. In the rounds with low draws of the parameter θ , the attack size is close to 1 (everyone choosing action A); in the rounds with high draws of θ , the attack size is close to 0 (everyone choosing action B). There is the intermediate range of fundamentals for which the size of the attack is decreasing in θ . The tipping point (the range of θ where subjects switch between actions A and B) depends on the cost of attacking. This result is consistent with previous findings, in particular HNO 2004, 2009.

Figure 1 demonstrates that the higher the cost of attacking, the less aggressive the subjects’ attacking behavior (i.e., they attack for a smaller range of the fundamental θ).¹⁶ However, Figure 2 shows that actual behavior is significantly less sensitive to the cost of attacking than theory predicts.¹⁷ In particular, Figure 2 compares the actual size of the

¹⁶Bootstrapping with 1000 randomly generated attack sizes based on each draw of θ and the corresponding actual size of the attack produces confidence intervals that make it possible to draw inference about the significance of cost at the aggregate level. I fail to reject the null that the attack sizes are the same when comparing the C20 and C50 and when comparing the C50 and C60 treatments even at the 90 percent confidence level. However, comparing the C20 with the C60 treatments does lead to the conclusion that the attack size is significantly smaller in the C60 treatment than in the C20 treatment for most realizations of θ (see Figure A1 in Appendix A).

¹⁷By contrast, HNO (2004) find support in favor of the effectiveness of the cost of attacking as a means to reduce the probability of a speculative attack. It is possible that the relatively low impact of cost changes here may also reflect the feature of the design that keeps the costs at a given level for the duration of 20

attack with the theory prediction for the entire range of realizations of θ in the C20 (Panel a) and C50 (Panel b) treatments (the C60 figure looks similar to the C50 figure).¹⁸ Each point represents a particular round associated with a single draw of θ and a corresponding attack fraction. The fitted smooth line derives from a locally weighted regression (same as in Figure 1 for the C20 and C50 treatments), while the dashed line represents the attack fraction predicted by the AHP (2007) model with 15 agents. In the C20 treatment, the actual size of the attack is actually slightly below the predicted size of the attack for the majority of θ realizations below 100 (although the difference is only statistically significant for a limited range of θ). By contrast, the actual size of the attack is significantly greater than predicted for the majority of realizations of θ in high cost treatments.

To get a better understanding of the origins of subjects' excess aggressiveness, I estimate the effects of the private signal, the cost of attacking, and beliefs on the individual probability of attacking, controlling for subject, round, and session fixed effects.¹⁹ The results from the regressions that use the first-stage data are reported in Columns 1 and 2 of Table 3. I confirm that the probability of attacking is decreasing in both the private signal and the cost of attacking. However, the effect of the cost is only statistically significant in the specification that excludes the expectation about the size of the attack (*Belief*). The introduction of the belief variable into the regression also reduces the magnitude of the coefficient on the private signal. Thus, subjects' expectations appear to play a crucial role in predicting the aggressive attacking behavior I observe on the aggregate level. The significant role of beliefs for individual decision-making is consistent with the strategic complementarity feature of the model. Interacting the independent variables with a dummy for the first 10 rounds of the experiment allows me to gauge whether learning across rounds significantly affects the results (Column 2). However, none of these additional controls are statistically significant in any of the specifications and their inclusion does not significantly change the coefficients on the main explanatory variables.

3.2 Is Aggressiveness Consistent with Best Response?

First, note that the payoff-dominant strategy in this setting is for all subjects to attack the status quo regardless of cost as long as the fundamental is below 100. If subjects were able to coordinate on this off-equilibrium behavior (in other words, if subjects were maximally ag-

periods. For example, HNO (2009) find that global games may overestimate the effect of the hurdle compared to responses to opportunity costs in some settings. One can address this issue using sessions that vary the costs each round, keeping the fundamentals constant. However, the large discontinuous change in the cost of attacking should draw more attention than small changes from round to round. (Subjects are made well aware of the change in the cost from treatment to treatment.) Therefore, finding much larger cost effects in this setting with repeated cost changes seems to be unlikely.

¹⁸The predicted size of the attack is calculated numerically for the $N = 15$ case using a Monte Carlo simulation to sample from the posterior distribution over θ and then calculate a threshold x^* such that subject i should attack if and only if $x_i < x^*$. Given this threshold, the size of the attack equals $\Pr(x_i < x^* | \theta) = \Phi(\sqrt{\beta_1}(x^* - \theta))$. The x^* for the $N = 15$ case is quantitatively very similar to the x^* for the infinite agent case of AHP (2007).

¹⁹The logistic specification is appropriate given the binary nature of the dependent variable. In order to obtain consistent estimates, conditional logit is used due to the panel structure of the data (Chamberlain 1980).

gressive), they would receive the greatest payoffs. However, unlike coordination games with public information which actually predict multiple equilibria, including a payoff dominant and a risk dominant equilibrium, a global game (a coordination game with private information about an unknown common fundamental) predicts existence and uniqueness only of the risk-dominant equilibrium. Thus, while several previous studies have found “aggressiveness” in the latter setting – subjects learning to play the payoff dominant equilibrium in the lab (Van Huyck et al. (1990), Schmidt et al. (2003)), my finding of excess aggressiveness in the high-cost treatments in the setting of a global game is more surprising. HNO (2004) find similar evidence; in this section, I explore further whether this observed aggressive behavior and the lack of learning over rounds are consistent with best response.

Recall that the estimates reported in Table 3 already suggest that beliefs play a crucial role for subjects’ attacking behavior. To test if subjects truly best respond to their beliefs, I construct the following criterion. For a given signal x_i , if subject i reports an expected fraction of others attacking to be 1 (all others attacking), then her best response is to attack if $x_i < \bar{x}$ and not to attack if $x_i > \bar{x}$, where \bar{x} is the value of the signal above which it is never optimal to attack regardless of others’ actions.²⁰ Analogously, if subject i reports an expected fraction of others attacking to be 0 (no other subject attacking), then her best response is not to attack if her $x_i > \underline{x}$ and to attack if $x_i < \underline{x}$, where \underline{x} is the value of the signal below which it is always optimal to attack regardless of others’ actions.²¹ For all other belief reports between 0 and 1, I take \tilde{x}_i to be the unique symmetric cutoff value consistent with the player’s reported beliefs (given her signal x_i). This \tilde{x}_i varies across individuals with different reported beliefs about the aggressiveness of the others in their group. Given \tilde{x}_i and x_i , subject i ’s best response is to attack if and only if the expected payoff is greater than the cost of attacking.

Given this criterion, I find that subjects act in accordance with their best-response to beliefs in about 81.4 percent of cases, averaging across cost treatments.²² Furthermore, learning over rounds does not seem to matter for best response behavior.²³ Note that, while I find evidence consistent with the majority of subjects best-responding to their beliefs, the criterion does not imply common knowledge of rationality. In other words, subjects still hold beliefs that are too aggressive relative to theory.

What remains to be explained is why the aggressive expectations are not corrected over rounds of play. While this study cannot offer a definitive explanation, one possibility is that, given these expectations, strategic complementarities blunt subjects’ incentives to correct mistaken beliefs, leading to the aggressive attacking behavior.²⁴ As Figure 3 demonstrates,

²⁰Note that θ can be greater than 1, in which case an attack always fails.

²¹Note that θ can be less than 1, in which case an attack always succeeds given that at least one subject attacks.

²²As compared to previous findings, such as Costa-Gomes and Weizsäcker (2008) and HNO (2009), this is a relatively high level of adherence to best response to individual beliefs.

²³See Table A2 in Appendix A for conditional logit estimates of the effect of rounds on the likelihood of individual behaving consist with best response, controlling for the private signal and cost.

²⁴The results therefore also relate to the literature on the role of persistent deviations from equilibrium play in environments characterized by strategic complementarity (Haltiwanger and Waldmann 1985, 1989; Fehr and Tyran 2005, 2008). This literature emphasizes the fact that under strategic complementarity the rational agents have an incentive to partially mimic the irrational agents which reinforces deviations from equilibrium play at the aggregate level. In this setting the mere belief that other players are deviating

subjects’s aggressiveness leads to a greater-than-predicted share of successful attacks (although only in the high-cost treatments).²⁵ Indeed, in the C50 treatment, 54 percent of all attacks were successful (compared to 40 percent in theory), and in the C60 treatment, 42 percent of all attacks were successful (compared to 26 percent in theory). The qualitative pattern displayed in Figure 3 also holds for the second 10 rounds of the experiment.

While aggressiveness in the first stage does not seem to be affected by repetition across rounds, I next show that learning between stages (within rounds) plays a crucial role for attacking behavior.

4 Learning to Be Less Aggressive in the Second Stage

AHP (2007) predict that the observation that the status quo has survived the first-stage attack would enable subjects to infer that the state of the fundamentals is not too weak, because otherwise it would have collapsed under the first attack. In particular, they learn that θ is above θ_1^* . This realization should cause a first-order-stochastic-dominance shift of beliefs upwards, causing subjects’ behavior to become less aggressive.²⁶

Do the experimental subjects exhibit this type of learning? First, consider Figure 4, which reports the average size of attack for rounds in which the game continued into the second stage. In the NNI treatment, the share of subjects attacking in the second stage decreases from 0.13 to 0.06, a reduction in aggressiveness that is statistically significant (t-test p-value < 0.001). Moreover, out of all the NNI rounds that continued into stage two, only 3 (4.8 percent) culminated in a successful attack (all in the C50 treatment).²⁷ Thus, subjects’ aggressiveness in the second stage is not sufficient to trigger successful attacks.

Second, consider the estimates reported in Columns 3 and 4 of Table 3 that are based on the data from all rounds that continued into the second stage of the NNI treatment. The significant negative effect of the *Stage* dummy implies that the survival of the status quo into the second stage significantly reduces aggressiveness (Column 3). Moreover, the inclusion of beliefs reduces the magnitude of the coefficient on the stage variable (although it remains statistically significant and negative; Column 4). This suggests that the mechanism through which subjects become less aggressive in the second stage of the NNI treatments is at least partly associated with their beliefs about others’ actions. Figure 5 plots reported individual beliefs about the size of the attack in stage one (in grey) and stage two (in black) for different realizations of the private signal for all the NNI treatments. The second-stage beliefs are

from the equilibrium frequency of attacking the status quo induces players to increase the probability of attacking, which partially reinforces the initial beliefs. This reinforcement may make it difficult to converge to the equilibrium. See also Izmailov and Yildiz (2010) for a general class of models where agents’ sentiments are believed to play an important role. Arifovic and Maschek (2012) build and simulate a model of a currency crisis where agents’ beliefs are the source of volatility that can lead to a devaluation.

²⁵The theoretical share of successful attacks for the N=15 case was calculated using a Monte Carlo simulation for each realization of θ to predict the number of attackers given the theoretical threshold x^* (see footnote 18).

²⁶In AHP (2007), this effect guarantees the existence of a unique equilibrium where no agent is willing to attack in the second stage (see Lemma 2 in AHP (2007) for proof). The robust prediction of the model is that the observation that the status quo has survived a first-period attack creates a large drop in the size of the attack in the second period.

²⁷See Figure 3 for the success rates in stage one by cost treatment.

everywhere below and to the left of the first-stage beliefs. The reduction in aggressiveness from the first to the second stage is statistically significant at the 95 percent confidence level.²⁸ This evidence is consistent with the predictions of AHP (2007).

However, in all cost treatments, I also find that beliefs about the size of the attack are significantly higher than zero. These aggressive beliefs may partially justify why some subjects still persist in attacking the status quo in the second stage. But why do beliefs about the number of attackers remain higher than predicted in the first place? One possible reason may be that subjects observe successful coordination in the second stage of the NNI treatments which reinforces their aggressive beliefs. However, I only observe 3 successful attacks (4.8 percent) in the second stage which all occur in the C50 treatment. Table 4 shows the impact of learning from experience: the attack size in the second stage decreases with rounds (Column 1) and, similarly, the second-stage attack is largest in the first 10 rounds of the experiment (Column 2). This is consistent with subjects converging toward the AHP (2007) prediction, although the significance of the effect disappears once I cluster standard errors at the group level.²⁹ For the sake of comparison, Table 4 also provides the estimates from the same kind of regressions using stage-one data. The attack size in the first stage increases after the first 10 rounds (Column 4) with subjects learning to coordinate early on in the experiment and after the first 10 rounds high level of aggressiveness does not depend on round. Thus, in both stages, subjects are exhibiting a tendency to converge toward behavior predicted by AHP (2007) over rounds.

The findings raise two follow-up questions: (1) Is the effect of learning that the survival of the status quo strong enough to prevent stage-two attacks even as new information arrives? (2) Is there an interaction between subjects' beliefs and learning with new information? I address these issues in the subsequent section.

5 Learning to Be More Aggressive in the Second Stage

This section explores what happens in the second stage of the game if subjects learn that the status quo has survived a first-stage attack, but at the same time also receive new information about the strength of the fundamental. That is, they get a new more precise private signal $x_{i2} = \theta + \frac{1}{\sqrt{\beta_2}}\xi_{i2}$, where $\xi_{i2} \sim N(0, 1)$. We already know that subjects become significantly less aggressive by learning that the status quo has survived the first-stage attack. Does this mean that they are always less aggressive in the second stage, or can this new information create an opportunity to coordinate on a new attack?

AHP (2007) predict that, for certain parameter values (see Table 1), the arrival of this new information, indeed, allows for an existence of a continuum of equilibria in which a new attack becomes possible. Note also that not attacking in the second stage always remains an equilibrium, even with new information. My experiment does not allow me to detect which equilibrium is selected in a given round.³⁰ Instead, I aim to answer a simpler question of

²⁸Confidence intervals were obtained by bootstrapping with 1000 randomly generated expected attack sizes based on each draw of x and the corresponding actual belief about the size of the attack.

²⁹Note that the number of rounds that continue into the second stage may simply be insufficient for subjects to achieve full convergence.

³⁰For more on multiple equilibria in this environment, see Shurchkov (2008).

whether new attacks become more likely with the arrival of new information in the second stage (which is a more robust and testable prediction).

Refer to the second part of Figure 4 which shows the average size of attack in the two stages for the NI treatment. In the first stage, the average attack fractions in the NNI and the NI treatments are not statistically different in magnitude: 0.13 and 0.14, respectively (t-test p-value = 0.7659). Therefore, the anticipation in the first stage that new information will arrive in the second stage does not alter behavior in the first stage. This fact is consistent with the theory. Furthermore, in otherwise similar rounds with new information, the second-stage average attack fraction is significantly higher than in the rounds without new information (t-test p-value = 0.0764). In fact, the average attacking frequency in the second stage of the NI treatment is not statistically different from the average attacking frequencies in the first stage of either NNI or NI treatments, with t-test p-values of 0.5853 and 0.5035, respectively.

Columns 5 and 6 of Table 3 report the estimates from the fixed-effects conditional logit regressions that show the effects of receiving new information on the individual probability of attacking. In all specifications, the probability of attack decreases significantly with the private signal. Also, attacking is significantly more likely in the second stage in the treatment with new information relative to the treatment with no new information. This evidence supports the above finding of the aggregate increase in aggressiveness with the arrival of new information in the second stage. Note that, in the second stage, beliefs about the attacking behavior of others are again highly significant in determining one's own likelihood of attacking. Moreover, once the belief variable is included in the regression (Column 6), the new information dummy becomes only marginally significant. Thus, new information seems to influence attacking behavior by affecting subjects' beliefs about how many others will be attacking. Finally, information treatment order and interactions of the control variables with the dummy for the first 10 rounds do not have a significant impact on behavior in the second stage. Thus, I conclude that the increase in aggressiveness in the second stage of the NI treatment relative to the NNI treatment cannot be attributed to learning through experience over rounds of play.

The persistent increase in aggressiveness in the second stage of the NI treatment relative to the NNI treatment occurs despite the fact that only one out of 47 attempted attacks (2.1 percent) was actually successful at overturning the status quo in the second stage of the experiment with new information. Once again, the individual-level regressions point to the importance of beliefs as at least a partial mechanism. Figure 6 plots second-stage beliefs for the NI (black line) and the NNI (grey line) treatments. Second-stage beliefs with new information lie above and to the right of the second-stage beliefs without new information for most realizations of x , and the difference is statistically significant at the 90 percent confidence level for the realizations of x above 80.³¹ The increase in the aggressiveness of beliefs in the second stage of the NI treatments partly justifies the associated increase in attacking behavior relative to the NNI case.

³¹Only a very small number of rounds survive into the second stage in the lower signal range which widens the confidence intervals for those realizations of x .

6 Conclusion

This study was motivated by certain open questions associated with events that require coordination, such as currency crises, capital flight episodes, or political uprisings. I exploit recent advances in dynamic global game theory to examine the impact of expectations and learning on attacking behavior in a laboratory experiment that mimics the essential features of such a coordination event.

In line with the theory and previous evidence, the attack is decreasing in both the underlying strength of the fundamentals and the cost of attacking. However, subjects are relatively insensitive to the changes in the cost of attacking and exhibit considerably more aggressive attacking behavior than the theory predicts when the cost of attacking is high. I show that the high level of aggressiveness is not simply the result of a random “taste for aggression,” but rather a best response to subjects’ aggressive expectations about the behavior of others. In fact, given their beliefs, the majority (81.4 percent) of subjects make attacking decisions that maximize their expected payoffs.

Subjects also exhibit a considerable degree of learning between the stages of the experiment. After a failed attack, the knowledge that the status quo has survived greatly reduces subjects’ likelihood of engaging in attacking behavior. However, if the agents receive new private information after a failed attack in the form of a more precise signal about the strength of the fundamentals, the probability of a new attack increases significantly. Information in the second stage appears to influence subjects’ aggressiveness by changing their beliefs about the behavior of others. On the other hand, experience with the game over the course of several rounds of play does not significantly improve subjects’ ability to coordinate.

In terms of policy, the results suggest that costly interventions, such as interest rate hikes, may not be particularly effective in thwarting speculative attacks. The policymaker’s ability to convince speculators that their fellow coordinators do not have an aggressive mindset may play a more important role in alleviating speculative pressures. More importantly, even if an initial attack has been successfully averted, economies with relatively widespread access to private information may be at a higher risk for repeat attacks.

The findings offer several directions for future research. Experimentally, the framework can be extended to have more than just two stages, in order to shed some light on the timing of speculative attacks. Furthermore, communication can play a role in equilibrium selection in the dynamic global game setting, and in multiplicity detection, in particular. Finally, the results point to the importance of belief formation and to the partially self-confirming nature of these beliefs (arbitrary beliefs induce attacking behaviors that partially confirm the initial beliefs). Thus, the range of behaviors that one can expect to occur may be considerably larger than the range of behaviors predicted by existing theory. A better understanding of the sources of the initial beliefs may constitute an exciting challenge for modelling coordination behavior in the future.

References

- [1] Angeletos, G-M., Hellwig, C., & Pavan, A. (2007). Dynamic global games of regime change: learning, multiplicity, and timing of attacks. *Econometrica*, 75(3), 711-756.
- [2] Arifovic, J. & Maschek, M. (2012). Currency crisis: evolution of beliefs, real world data and policy experiments. *Journal of Economic Behavior and Organization*, 82, 131-150.
- [3] Atkeson, A. (2000). Discussion of Morris and Shin's 'Rethinking multiple equilibria in macroeconomic modelling'. *NBER Macroeconomics Annual*.
- [4] Blanco, M., Engelmann, D., Koch, A. K. & Normann, H-T. (2008). Belief elicitation in experiments: is there a hedging problem? IZA Discussion Papers 3517.
- [5] Brunnermeier, M. K. & Morgan, J. (2010). Clock games: theory and experiments. *Games and Economic Behavior*, 68(2), 532-550.
- [6] Cabrales, A., Nagel, R., & Armenter, R. (2007). Equilibrium selection though incomplete information in coordination games: an experimental study. *Experimental Economics*, 10(3), 221-234.
- [7] Carlsson, H. & van Damme, E. (1993a). Global games and equilibrium selection. *Econometrica*, 61(5), 989-1018.
- [8] Carlsson, H. & van Damme, E. (1993b). Equilibrium selection in stag hunt games in *Frontiers of Game Theory*, ed. by K. Binmore, A. Kirman, & P. Tani. Cambridge, MA: MIT Press, 237-253.
- [9] Chamberlain, G. (1980). Analysis of covariance with qualitative data. *Review of Economic Studies*, 47, 225-238.
- [10] Chamley, C. (1999). Coordinating regime switches. *The Quarterly Journal of Economics*, 114(3), 869-905.
- [11] Chen, Q., Goldstein, I. & Jiang, W. (2010). Payoff complementarities and financial fragility: evidence from mutual fund outflows. *Journal of Financial Economics*, 97(2), 239-262.
- [12] Cheung, Y-W. & Friedman, D. (2009). Speculative attacks: a laboratory study in continuous time. *Journal of International Money and Finance*, 28(6), 1064-1082.
- [13] Cooper, R. W., DeJong, D. V., Forsythe, R., & Ross, T. W. (1990). Selection criteria in coordination games: some experimental results. *American Economic Review*, 80(1), 218-233.
- [14] Cooper, R. W., DeJong, D. V., Forsythe, R., & Ross, T. W. (1992). Communication in coordination games. *Quarterly Journal of Economics*, 107(2), 739-771.

- [15] Cornand, C. & Heinemann, F. (2010). Measuring agents' reaction to private and public information in games with strategic complementarities, CESifo Working Paper Series 2947.
- [16] Corsetti, G., Guimaraes, B., & Roubini, N. (2003). International lending of last resort and moral hazard: a model of IMF's catalytic finance. NBER Working Paper 10125.
- [17] Costain, J. S., Heinemann, F., & Ockenfels, P. (2007). Multiple outcomes of speculative behavior in theory and in the laboratory. Bank of Spain working paper.
- [18] Costa-Gomes, M. & Weizsäcker, G. (2008). Stated beliefs and play in normal form games. *Review of Economic Studies*, 75, 729-762.
- [19] Danielsson, J. & Peñaranda, F. (2011). On the impact of fundamentals, liquidity and coordination on market stability. *International Economic Review*, 52(3), 621-638.
- [20] Edmond, C. (2008). Information revolutions and the overthrow of autocratic regimes. Mimeo, New York University.
- [21] Fehr, E. & Tyran, J-R. (2005). Individual irrationality and aggregate outcomes. *Journal of Economic Perspectives*, 19(4), 43-66.
- [22] Fehr, E. & Tyran, J-R. (2008). Limited rationality and strategic interaction – the impact of the strategic environment on nominal inertia. *Econometrica*, 76(2), 353-394.
- [23] Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171-178.
- [24] Gächter, S. & Renner, E. (2010). The effects of (incentivized) belief elicitation in public goods experiments. Discussion Papers 2010-12, The Centre for Decision Research and Experimental Economics, University of Nottingham.
- [25] Goldstein, I., & Pauzner, A. (2004). Contagion of self-fulfilling financial crises due to diversification of investment portfolios. *Journal of Economic Theory*, 119(1), 151-183.
- [26] Haltiwanger, J. C. & Waldman, M. (1985). Rational expectations and the limits of rationality: an analysis of heterogeneity. *American Economic Review*, 75(3), 326-40.
- [27] Haltiwanger, J. C. & Waldman, M. (1989). Limited rationality and strategic complements: the implications for macroeconomics. *Quarterly Journal of Economics*, 104(3), 463-83.
- [28] 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.
- [29] Heinemann, F., Nagel, R., & Ockenfels, P. (2009). Measuring strategic uncertainty in coordination games. *Review of Economic Studies*, 76, 181-221.
- [30] IMF. (2000). Recovery from the Asian Crisis and the role of the IMF.

- [31] Izmalkov, S. & Yildiz, M. (2010). Investor sentiments. *American Economic Journal: Microeconomics*, 2(1), 21-38.
- [32] Morris, S. & Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88(3), 587-597.
- [33] Morris, S. & Shin, H. S. (2004). Coordination risk and the price of debt. *European Economic Review*, 48, 133-153.
- [34] Obstfeld, Maurice. 1996. "Models of currency crises with self-fulfilling features." *European Economic Review*, 40(3-5), 1037-1047.
- [35] Prati, A. & Sbracia, M. (2002). Currency crisis and uncertainty about fundamentals. Economic working papers 446, Bank of Italy.
- [36] Rochet, J. & Vives, X. (2004). Coordination failures and the lender of last resort: was Bagehot right after all? *Journal of the European Economic Association*, 2(6), 1116-1147.
- [37] Schotter, A. & Yorulmazer, T. (2009). On the dynamics and severity of bank runs: an experimental study. *Journal of Financial Intermediation*, 18(2), 217-241.
- [38] Schmidt, D., Shupp, R., Walker, J. M., & Ostrom, E. (2003). Playing safe in coordination games: the roles of risk dominance, payoff dominance, and history of play. *Games and Economic Behavior*, 42, 281-299.
- [39] Shurchkov, O. (2008). Effects of one-sided communication on coordination and equilibrium selection in dynamic global games: experimental evidence. Mimeo, MIT.
- [40] Van Huyck, J. B., Battalio, R. C., & Beil R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review*, 80(1), 234-48.
- [41] Wang, S. W. (2011). Incentive effects: the case of belief elicitation from individuals in groups. *Economic Letters*, 111, 30-33.