Coordination and Learning in Dynamic Global Games: Experimental Evidence*

Olga Shurchkov†

Current Version: February 27, 2007

First Version: September 31, 2006

Abstract

This paper uses a laboratory experiment to test the predictions of a dynamic global game designed to capture the self-fulfilling nature of speculative attacks. The game has two stages and a large number of heterogeneously informed agents deciding whether to attack a status quo. In the first stage, the equilibrium size of the attack is decreasing in both the underlying strength of the status quo and the agents' cost of attacking. In the second stage, the knowledge that the status quo has survived the first-period attack decreases the incentive to attack, implying that a new attack is possible only if agents receive new information. Our experimental evidence supports these theoretical predictions in both stages. However, we also find that the subject's actions are overly aggressive relative to the theory's predictions. We further find that the excess aggressiveness in actions stems from the aggressiveness of their beliefs about others' actions.

*I am especially grateful to Ernst Fehr for his invaluable insights, continuous support, and help in planning and implementing the experiment. He also provided the resources and the financial support without which this study would not be possible. I am indebted to my advisor, Marios Angeletos, for his time, helpful comments, and guidance throughout this project. I would also like to thank my advisor, Daron Acemoglu, Miriam Bruhn, Florian Ederer, Muhamet Yildiz, and the participants of the MIT macroeconomics and theory field lunches for valuable comments and discussion. Finally, I would like to acknowledge the hospitality of the Institute for Empirical Research in Economics where the experimental sessions were conducted. All remaining errors are my own.

†Department of Economics, Massachusetts Institute of Technology, 50 Memorial Drive, MA 02139. e-mail: oshurchk@mit.edu.

1 Introduction

In recent years, there has been a growing literature on global games – coordination games of incomplete information – and their applications. By incorporating strategic uncertainty into agents' actions and introducing heterogenous information that agents receive over time, these games can reduce the set of equilibrium outcomes, thereby delivering concrete predictions about equilibrium selection.

Global coordination games can be applied to a variety of crises situations: a speculative attack against a currency peg, a bank run, or a riot. It is currency crises, in particular, that have received much attention in recent history. Within the affected country, a crisis can have an enormous negative impact on economic growth and can cause political change and turmoil.[1] Furthermore, in recent years, currency crises have been increasing in frequency, while at the same time remaining largely unpredictable and unexpected. It is this last fact that makes the theory of global games appealing, since it can help us understand the way in which people form expectations and the way in which they coordinate on different courses of action.

While it would be desirable to test the theory of global games with field data, this proves to be a rather difficult task. Field data contain additional forces not captured by the model that can limit the ability of standard econometric techniques to identify the model. Furthermore, one of our goals is to understand the structure of agents' equilibrium strategies. However, the available data on crises are at the aggregate level and contain no information about individual behavior. Finally, field data contain only the information about people's actions and provide no insight into people's expectations that cause these actions.

In this paper, we use an experimental approach to analyze agents' behavior and formation of expectations under the conditions of a speculative attack, which allows us to shed light on the reasons behind the onset and timing of currency crises. For this, we restrict our attention to the class of models that capture the following three features of currency crises: (1) the coordination element of currency crises that arises due to strate-

[1] For example, in Indonesia, the Asian economic crisis of 1997-1998 caused the growth rate of real GDP to fall from 8.2 percent in 1996 to 1.9 percent in 1997 to -14.2 percent in 1998. The figures are similar in other affected countries (IMF, 2000). Furthermore, in Indonesia, the crisis played a role in ending Suharto's long period of authoritarian rule when riots and demonstrations caused his political isolation, finally compelling him to resign. In Thailand and South Korea, democratic elections were held and opposition parties came to power for the first time since political liberalization (Freedman, 2004). The effects of these crises go beyond the borders of just one nation. More and more often, crises spread from the country of origin to other economies, threatening to cause worldwide contagion (Forbes, 2000, Boston, 2003).

gic 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. While the so-called "second generation" models of currency crises (Obstfeld,1996) focus on the first feature, the literature on static global games (Morris and Shin, 1998) concentrate on the second. More recently, Angeletos, Hellwig, and Pavan (2006) captured the third feature of crises by extending the global games framework to incorporate dynamics.

Model. Our experiment is based on a two-period variant of the latter paper. The model consists of a large number of agents and two possible regimes, the status quo and an alternative. The game continues into the second period as long as the status quo is in place. In each period, each agent 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 period and negative otherwise. Regime change, in turn, occurs if and only if the percentage of agents attacking exceeds a threshold $\theta \in \mathbb{R}$ that parameterizes the strength of the status quo. The parameter θ captures the component of the payoff structure (the "fundamentals") that is never common knowledge. In the first period, each agent receives a private signal about θ. If the game continues into the second period, agents may or may not receive more private information about θ.

The model admits a variety of interpretations. For example, it can be applied to self-fulfilling bank runs, currency crises, or political change. In these contexts, regime change occurs, respectively, when a large run forces the banking system to suspend its payments, when a large speculative attack forces the central bank to abandon the peg, or when a large number of citizens decide to take actions to subvert a repressive dictator or some other political establishment.

Model Predictions. In the first period (stage) of the game, agents follow monotone threshold strategies, attacking the status quo if their individual signal is below a certain threshold and refraining from attacking otherwise. This implies that the size of the attack is monotonically decreasing in the strength of the economic fundamentals, θ, and, equivalently, the agents' strategies are decreasing in their individual private signals. The second testable prediction in the first stage is that both the individual and aggregate thresholds are decreasing in the cost of attacking.

The second-period predictions depend on the information structure. If the agents do not receive a second private signal in the second stage, then not attacking is the unique equilibrium. This result arises from the fact that agents have learned that the game survived a past attack, and hence the fundamental θ must be good enough (which can be viewed as "endogenous learning"). On the other hand, in addition to a no-attack equilibrium, a new attack becomes possible, if agents receive sufficiently precise new information in the second stage and have a relatively "lenient" initial prior about the state of the fundamentals, i.e., the first-period attack is not too aggressive. Multiplicity originates from the interaction between two elements: the knowledge that the regime survived past attacks and the arrival of new private information over time (i.e., the interaction between "endogenous" and "exogenous learning"). In particular, "endogenous learning" reduces the probability of attack in the second period, but "exogenous learning" can make a new attack possible.

Experimental Results. To test the predictions of the model, we conduct several treatments of a laboratory experiment where we vary the strength of the fundamentals, the cost of attacking, and the availability of information in the second stage.

In the first stage of the experiment, we find that agents' strategies are consistent with the theory that predicts a unique equilibrium in monotone strategies. However, we observe that the subjects behave more aggressively relative to the theory prediction for the treatments in which the cost of attacking is set to be relatively high. Furthermore, we find that the agents' behavior is less responsive to the changes in the cost of attacking than in theory.

In the second stage of the experiment, we detect the effects of the interaction between learning and coordination. In order to distinguish between the effects of endogenous and exogenous learning, we first run a treatment where the subjects do not receive any new private information. In this case, the only additional information that the subjects receive in stage two is that the game has not ended, which can be interpreted as "endogenous learning." We find that, in the second stage of this treatment, the subjects are in fact learning from the outcome of the first stage since the probability of attack is greatly reduced. However, the endogenous learning effect is not as strong as in the model, and the magnitude of the effect depends on the cost of attacking.

In order to examine the interaction between endogenous and exogenous learning, we then run a treatment where the subjects receive an additional more precise private signal in the second stage of the experiment. In this case, the subjects are still able to learn endogenously through their observation that the experiment proceeded into stage two, but in addition, they can now learn exogenously by incorporating this more precise information into their decision to attack the status quo. We find that the probability of attack in the second stage now increases significantly relative to the treatment with endogenous learning only.

While the theory predictions regarding the monotonicity of strategies and the importance of learning are supported by the experimental evidence, as mentioned above, we find that the subject's actions are overly aggressive relative to the theory's predictions. Therefore, we test two hypotheses about the aggressiveness of subjects' behavior. The first hypothesis postulates that, given their beliefs of what other agents will do, the subjects make "mistakes" or "too aggressively." An alternative hypothesis is that agents do not make "mistakes" given their beliefs of what others will do and that it is their beliefs that are more aggressive that the theory predicts. We find evidence that leads us to reject the first hypothesis: given the subjects' aggressive expectations relative to the model predictions, their actions are mostly consistent with best-response strategies. That is, subjects' aggressive behavior stems from their aggressive beliefs, rather than from "mistakes".

Related Literature. The theoretical literature on coordination games, applied to currency crisis, starts with a seminal paper by Obstfeld (1996). Obstfeld's model is a coordination game with perfect information that yields multiple equilibria. While these so-called "second generation" models capture the self-fulfilling aspect of currency crises that can happen without any apparent change in macroeconomic fundamentals, they can also be viewed as incomplete and weak in delivering policy prescriptions or predictions about the short-run timing of crises. Seeking to resolve the indeterminacy, Carlsson and van Damme (1993a), (1993b) and later Morris and Shin (1998) relax Obstfeld's assumption of common knowledge. They show that, under certain restrictions on the information structure, multiplicity of equilibria can be eliminated by assuming that agents receive heterogenous private information about the state of the fundamentals. This result has already been applied to several macroeconomic phenomena: see Goldstein and Pauzner (2001) 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 (2005) for political change.

To capture the interaction between the dynamics of expectations about the economic fundamentals and the dynamics of coordination, Angeletos, Hellwig, and Pavan (2006) extend the static benchmark of Morris and Shin by allowing agents to take actions in multiple periods and accumulate information over time. This extension emphasizes the idea that speculators have the option to take multiple shots against the currency peg and may also accumulate information over time and learn from past outcomes. The authors show that dynamics may sustain multiple equilibria. Multiplicity originates from the interaction between two types of learning: the knowledge that the regime survived past attacks ("endogenous learning") and the arrival of new private information over time ("exogenous learning").

The experimental literature that tests the predictions of the above models begins with studies that are based on coordination games with perfect information, in which equilibria are Pareto-rankable (as is the case in Obstfeld's model). Cooper et al. (1990) find that the observed pattern of play is accurately predicted by the Nash equilibrium concept and that coordination failures can emerge in which the outcome is a Pareto-inferior Nash equilibrium. Van Huyck, Battalio, and Beil (1990) find similar results when they study a class of tacit pure-coordination games with multiple equilibria. In particular, their experimental results suggest that the Pareto-dominant outcome is extremely unlikely either initially or in repeated play and that coordination failures arise due to strategic uncertainty. In a follow-up experiment, Cooper et al. (1992) study coordination games with nonbinding, pre-play communication. They find that in coordination games with a cooperative strategy, one-way communication increases play of the Pareto-dominant equilibrium relative to the no-communication baseline.

Cabrales, Nagel, and Armenter (2003) test the *global* coordination game theory in two-person games with random matching inspired by Carlsson and van Damme (1993a). They find that, with private information about the payoffs, the subjects' behavior converges to the theoretical prediction after enough experience has been gained. In a recent paper, Heinemann, Nagel, and Ockenfels (2004) test the predictions of the static speculative-attack model of Morris and Shin in a laboratory experiment. The authors compare sessions with private and public information, and conclude that in all sessions subjects used threshold strategies, i.e., attacked whenever the state of the fundamentals or the signal was beyond some critical state or signal, respectively. The authors point to these results as evidence in support of the theory which predicts a unique equilibrium under certain parameter restrictions.

This study is the one of the first to test the predictions of *dynamic* global game theory in a laboratory experiment. Brunnermeier and Morgan (2004) examine "clock games" that end when the third of six players exits, and those three players receive a payoff that increases continuously in the exit time. The authors report that, consistent with the unique symmetric pure strategy Nash equilibrium, players exit sooner when they have better information about other players' choices and clock settings. In a more recent study, Cheung and Friedman (2006) examine speculative attacks with varying amounts of public information featuring continuous time and focus on size asymmetries (i.e., the effect of a large player on behavior and outcomes). They find that weaker (or more rapidly deteriorating) fundamentals increase the likelihood of successful speculative attacks and hasten their onset, and that public access to information about either the net speculative position or the fundamentals also enhances success. The presence of a larger speculator further enhances success.

To our knowledge, our experiment is the first to put to the test a dynamic global game and focus on the theoretical predictions of uniqueness in dynamic global games, to detect learning over time, and to address the questions of subjects' aggressiveness and rationality.

The rest of the paper is organized as follows. Section 2 describes the dynamic two-period model of a speculative attack and discusses theoretical predictions to be tested. Section 3 describes the experimental procedures and treatments. Section 4 describes the results of the data analysis. Section 5 concludes and discusses possible extensions.

2 The Model

Our model is a simple two-period version of the model developed by Angeletos, Hellwig, and Pavan (2006) in which there are two regimes, the status quo and the alternative. The agents, indexed by i, decide simultaneously 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. The status quo collapses if the mass of agents choosing action A ("aggregate size of the attack"), exceeds θ, which parametrizes the strength of economic fundamentals. A low value of θ thus represents a relatively weak state of the fundamentals, and a high value of θ represents a relatively strong state of the fundamentals. We will 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. Action A is associated with an opportunity cost c. If action A is successful (i.e., the status quo is abandoned), each agent choosing action A earns an income of $y > c$. If not (i.e., the status quo prevails), then the agent choosing action A earns 0. Action B yields no payoff and has no cost.[2] The payoff of an individual agent can be written as

$$u_{it} = U(a_{it}, A_t, \theta) = \begin{cases} a_{it}(y - c) \text{ if } A_t \geq \theta \\ -a_{it}c \text{ if } A_t < \theta \end{cases} .$$

where $a_{it} \in \{0, 1\}$ denotes the action chosen by agent i at time t ($a_{it} = 1$ represents attacking and $a_{it} = 0$ represents not attacking) and A_t denotes the aggregate size of the attack at time t.

[2]Note that the payoff to the agent does not depend on θ. Or it only depends on it in the following way: if θ is so low that the regime collapses, the payoff to the agent choosing action A is y, but if the regime survives, the payoff is always the same (0).

Note that Obstfeld's key assumption that delivers multiplicity is that the state of the fundamentals, θ, is common knowledge among the agents. In that case, each agent's best response function is

$$g(A_t, \theta) = \underset{a_{it} \in \{0,1\}}{\arg\max} U(\cdot) = \begin{cases} 1 \text{ if } A_t \geq \theta \\ 0 \text{ if } A_t < \theta \end{cases} .$$

We can see that, for all $\theta \in [0, 1]$, there are two pure-strategy Nash equilibria in this game, namely that either all agents choose action A or all agents choose action B; $[0, 1]$ is the region where multiplicity is possible.

However, in our setup, agents have heterogeneous information about the strength of the status quo. Nature draws θ from a normal distribution $N(z, 1/\alpha)$ which defines the initial common prior about θ. Note that z can be thought of as the public signal that all agents receive. In addition to receiving a public signal, each agent then receives a private signal $x_{it} = \theta + \xi_{it}$, where $\xi_{it} \sim N(0, 1/\beta_t)$ is i.i.d. across agents and independent of θ and β_t is the precision of private information. [3] The status quo is in turn abandoned if and only if the measure of agents choosing action A, which is denoted by A_t, is greater than or equal to θ.

2.1 First-Period Predictions

Let us first focus on the equilibrium in the first period of the game. Note that it is strictly dominant to choose action A for sufficiently low signals – namely for $x_1 < \underline{x}$, where $\underline{x}$ solves $Pr(\theta \leq 0|\underline{x}) = c/y$ – and to choose B for sufficiently high signals – namely for $x_1 > \overline{x}$, where $\overline{x}$ solves $Pr(\theta \leq 1|\overline{x}) = c/y$.[4] This suggests that we should look for monotone Bayesian Nash equilibria in which the agents' strategy is non-increasing in x_1.

Next, we characterize the equilibrium in the first period. Suppose that there is a threshold, x_1^*, such that an agent will choose action A if and only if he gets a private signal below this cutoff ($x_1 \leq x_1^*$). The measure of agents choosing action A is then decreasing in θ and is given by $A_1(\theta) = Pr(x_1 \leq x_1^*|\theta) = \Phi(\sqrt{\beta_1}(x_1^* - \theta))$, where Φ is the c.d.f. of the Standard Normal distribution. It follows that the status quo is abandoned if and only if $\theta \leq \theta_1^*$, where θ_1^* solves $\theta_1^* = A_1(\theta_1^*)$, or equivalently

$$\theta_1^* = \Phi(\sqrt{\beta_1}(x_1^* - \theta_1^*)). \tag{1}$$

[3]The 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. The agents know the values of z, α, and β_t.

[4]Note that, for notational tractability, we supress the individual subscript, i, from now on.

The posterior probability of regime change for an agent with signal x_1 is then simply $Pr(R_1 = 1|x_1) = Pr(\theta \leq \theta_1^*|x_1)$. Since the latter is decreasing in x_1, each agent finds it optimal to choose action A if and only if $x_1 \leq x_1^*$, where x_1^* solves $Pr(\theta \leq \theta_1^*|x_1^*) = c/y$. Since posteriors about θ are normally distributed with mean $\frac{\beta_1}{\beta_1+\alpha}x_1 + \frac{\alpha}{\beta_1+\alpha}z$ and variance $\frac{1}{\beta_1+\alpha}$ (precision $\beta_1 + \alpha$), this condition is equivalent to

$$\Phi\left(\sqrt{\beta_1+\alpha}(\theta_1^* - \frac{\beta_1}{\beta_1+\alpha}x_1^* - \frac{\alpha}{\beta_1+\alpha}z)\right) = c/y. \tag{2}$$

If we solve the system of the two equations (1) and (2), we get the equilibrium values (x_1^*, θ_1^*) in the first period of the model. Such a solution always exists and is unique for all z if and only if $\beta_1 \geq \frac{\alpha^2}{2\pi}$. (See Appendix A for the proof.) Moreover, iterated elimination of strictly dominated strategies implies that, when the monotone equilibrium is unique, there is no other equilibrium.

Figure 1 illustrates the equilibrium in the first period. It plots the aggregate size of the attack, A_t, against θ.

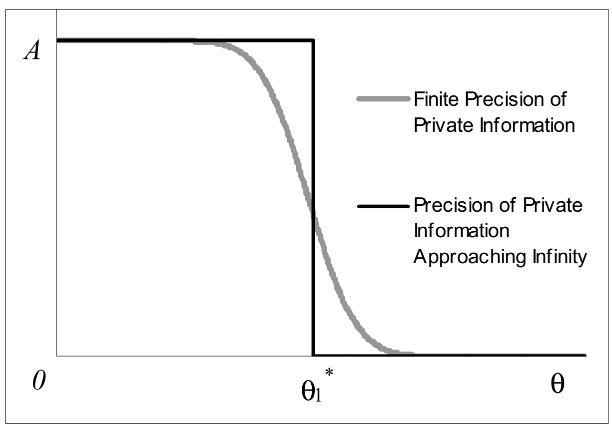

Figure 1: Size of the Attack, A_t, vs. θ for Finite $\beta_1$and β_1 Approaching Infinity

As the precision of private information approaches infinity, the theory predicts that everyone should choose action A for all $\theta \leq \theta_1^*$, and no one should choose action A for all $\theta > \theta_1^*$, which is represented by the black line in Figure 1. However, for finite precision of private information, the theory predicts that $A_1(\theta)$ is monotonically decreasing in θ, resulting in the grey line in Figure 1.

The following propositions summarize the first-period predictions of the theory that we would like to test.

Prediction 1. *There exists a unique x_1^*, such that in any equilibrium of the dynamic game, an agent chooses action A ("attack") in the first period if and only if $x_1 < x_1^*$. By* *implication, $A_1(\theta)$ is decreasing in θ, and there exists a unique θ_1^* such that the status quo is abandoned in the first period if and only if $\theta < \theta_1^*$.*

Prediction 2. *The thresholds θ_1^* and x_1^* are decreasing in c, the cost of choosing action A.*

Note that in a coordination game with complete information, such as Obstfeld (1996), the cost of attacking plays no role in equilibrium play: either everyone attacks the status quo or no one attacks the status quo, regardless of cost. Moreover, agents' strategies need not be monotonic in θ.

2.2 Second-Period Predictions

The game continues into the second period as long as the status quo is in place, and the game ends if the status quo is abandoned in the first period. We will consider two possibilities for the information structure in the second period. First, suppose that the agents receive no additional private signal. In this case, when agents arrive at the second period, they observe that the status quo must have survived the first-period attack. From the observation that the status quo is still in place, the agents learn that the state of the fundamentals is not too weak, because otherwise it would have collapsed under the first attack. In fact, they now know that θ must be above θ_1^*. The knowledge that $\theta > \theta_1^*$ causes a first-order-stochastic-dominance shift of beliefs upwards, causing agents' behavior to become less aggressive. It turns out that this effect is strong enough to imply that no agent is willing to take action A in the second period, delivering a unique equilibrium.[5]

Prediction 3. *If no new information arrives in the second period, then choosing action B ("not attack") for all x is the unique continuation equilibrium.*

Testing this prediction allows us to isolate the effects of endogenous learning on agents' behavior. However, we would also like to examine the effects of the interaction between endogenous and exogenous learning. This can be accomplished by changing the information structure in period two, such that agents receive an additional signal that is sufficiently precise. That is, $x_{i2} = \theta + \xi_{i2}$, where $\xi_{i2} \sim N(0, 1/\beta_2)$ and β_2 is sufficiently high.

The size of the attack in period two is given by $A_2(\theta) = \Pr(x_2 \leq x_2^*|\theta)$, which is decreasing in θ, and the probability of regime change for an agent with signal x_2 is $\Pr(R_2 = 1|x_2, R_1 = 0) = \Pr(\theta \leq \theta_2^*|x_2, \theta > \theta_1^*)$, which is decreasing in x_2 if $\theta_2^* > \theta_1^*$.

[5] See Lemma 2 in Angeletos, Hellwig, and Pavan (2006) for proof of uniqueness in any monotone equilibrium. Overall uniqueness can be shown by iterated deletion of strictly dominated strategies. See Appendix A for the iterated dominance argument.

Therefore, in any equilibrium in which an attack occurs in the second period, θ_2^* and x_2^* solve

$$\theta_2^* = \Phi(\sqrt{\beta_2}(x_2^* - \theta_2^*)) \tag{3}$$

$$1 - \frac{\Phi(\sqrt{\beta_2+\alpha}(\frac{\beta_2}{\beta_2+\alpha}x_2^* + \frac{\alpha}{\beta_2+\alpha}z - \theta_2^*))}{\Phi(\sqrt{\beta_2+\alpha}(\frac{\beta_2}{\beta_2+\alpha}x_2^* + \frac{\alpha}{\beta_2+\alpha}z - \theta_1^*))} = c/y. \tag{4}$$

Equations (3) and (4) are the second-period equivalents of equations (1) and (2) in the dynamic setting. We can solve them for θ_2^* and x_2^*, which will tell us under which conditions we can have an attack in the second period. In particular, the public information revealed to all subjects before the first period must be such that z is sufficiently high.[6] Intuitively, when z is high ("lenient prior"), arrival of new more precise private information makes the marginal agent more aggressive and may eventually offset the incentive not to choose action A induced by the knowledge that the regime survived past attacks. Indeed, if z is sufficiently high, so that $\theta_1^* < \theta_\infty$, then a second attack necessarily becomes possible once β_2 is sufficiently high (i.e., the second signal is sufficiently precise). Note that θ_∞ is the limit of the equilibrium threshold of the static game as the precision of private information becomes infinite (in particular, $\theta_\infty = 1 - c/y$). In this case, the theory predicts that , in addition to the no-attack equilibrium (which is *always* an equilibrium), there can also be an attack equilibrium (for proofs and derivations see Appendix A). The effect brought about by the introduction of exogenous information, making the marginal agent more aggressive, counteracts the effect of endogenous learning that makes the agent relatively less aggressive. In fact, the effect of more precise information can offset the incentive not to choose action A induced by the knowledge that the regime has survived past attacks, thus making new attacks possible.

Note that, when z is low ("aggressive prior"), an increase in the precision of private information makes the marginal agent less aggressive (that is, θ_t^* decreases with β_t). The knowledge that the regime survived an attack in the first period then only reinforces this effect. Therefore, for sufficiently low z, such that $\theta_1^* > \theta_\infty$, there exists a unique equilibrium, such that no agent ever chooses action A after the first period. However, in this experiment, we explore the scenario where a new attack becomes possible, and therefore only test the following prediction:

Prediction 4. *There is always an equilibrium in which agents choose action B ("not attack") for all x in the second period. If the agents receive a new private signal in the*

[6] For a special case where the cost of attacking equals 1/2, "high z" means $z > 1/2$. Under this condition, $\theta_1^*(z) < 1/2$, because $\theta_1^*(z)$ is monotonically decreasing in z and $\theta_1^*(1/2) = 1/2$. (See Appendix A for proof.)

second period such that β_2 is sufficiently large and if $\theta_1^* < \theta_\infty$, then, in addition to the no-attack equilibrium, there also exists an x_2^* and an equilibrium in which an agent chooses action A ("attack") if and only if $x_2 < x_2^*$.

In the following sections, we proceed to test the above predictions in a laboratory experiment.

3 Overview of the Experiment

3.1 Procedures

We conducted six sessions of the experiment at the experimental laboratory at the Institute for Empirical Research in Economics at the University of Zurich, with the first four sessions held in June of 2006 and the following two sessions held in October of 2006. The subjects were all students at the University of Zurich. The procedure was kept the same throughout all six sessions, except that the order of the treatments was reversed to test whether the order of the treatments mattered for the results. All sessions were computerized using the program z-Tree (Fischbacher, 1999). The subjects were first asked to read through and sign informed consent forms for non-biomedical research.[7] Paper copies of the instructions[8] were distibuted to the participants prior to the beginning of the experiment. The subjects were asked to answer several control questions that tested their understanding of statistics, as well as the experimental procedures. Questions were answered in private. The subjects could not see or communicate with one another. At the end of the experiment, each participant filled out a computerized questionnaire.[9] The questionnaire asked the subjects about their strategies, as well as their understanding of statistics and probability. At the very end, each subject was paid in cash a show-up fee equal to 15 Swiss Francs (CHF) and his or her earnings over the course of the session. Final income of each subject was first given in points and then converted to Swiss Francs at the rate of 10 points = 50 centimes for Sessions 1-2 and at the rate 10 points = 25 centimes for Sessions 3-6. Average income (including the show-up fee) was 83.8 CHF, 45.5 CHF, and 25.5 CHF for Sessions 1-2, 3-4, and 5-6, respectively.[10]

Each of the six experimental sessions had 30 participants divided randomly into two groups of fifteen people. Each session consisted of 40 independent rounds of play, with

[7] Copies of the informed consent forms in German or English are available upon request.

[8] Full copies of the instructions in German or English are available upon request for all treatments.

[9] Copies of the questionnaire questions in German or English are available upon request.

[10] Note that the differences in average incomes across sessions 3-4 and 5-6 arise due to the differences in the cost of choosing action A.

each round corresponding to a new random number θ drawn from a normal distribution $N(z, 1/\alpha)$[11]. Thus, one can interpret each round as a new economy parametrized by the state of fundamentals, θ. The 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 private signal ("hint number" x_1) about the random number θ. The subjects were given information about the distribution of this hint number in the instructions.

Each round consisted of one or two periods ("stages") of decision-making. In stage 1 of each round, each subject had to decide between actions A or B as described in section 2. Once all the subjects chose their actions in each stage of every round, they were asked a follow-up question, namely: "How many other members of your group do you think chose action A?" Next, each subject received the following information: if the game ended after stage one, he or she found out that action A was successful, learned the value of the unknown number, how many other subjects chose action A, and his or her payoff in the round. If the game continued into the second stage, two scenarios were possible. In the treatment without new information, the subject did not get an additional "hint number" after the first stage, but only got a reminder of his or her original private signal and received notification that action A was not successful. In the treatment with new information, the subjects received a new more precise signal if the game continued into stage 2.

We ran 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 I.

[11]This experiment relies on the use of the normal distribution, as opposed to the uniform distribution used by Heinemann, Nagel, and Ockenfels (2004). The ideas are essentially identical. The benefit of running the experiment using the normal is the tractability of the analysis. The theory involved in the dynamic case is significantly more complicated as compared to the static benchmark, which necessitates the use of the normal. The normal distribution is also relatively simple to grasp for the test subjects, since it is fully parametrized by the mean and the precision. To ensure the subjects' full understanding during the experiment, we provided them with several examples and conducted a quiz to familiarize them with the normal distribution.

Table I.

Treatment Summary.

Session	Cost of Action A: First 20 Rounds	Cost of Action A: Last 20 Rounds	Information in Stage 2: First 20 Rounds	Information in Stage 2: Last 20 Rounds
1,2	**20**	**50**	No	No
3,4	**50**	**20**	No	No
5	60	60	**No**	**Yes**
6	60	60	**Yes**	**No**

3.2 Parameterization

We re-scaled all numbers by a factor of 100, so that the subjects did not have to deal with fractions. We chose the gross payoff, y, of a successful attack to be 100 and the gross payoff of an unsuccessful attack to be 0. This payoff scheme was chosen for its simplicity for the theory, as well as for the experiment participants.

Table II records the remaining parameters by session.

Table II.

Parameterization.

Session	z, $1/\alpha$	$1/\beta_1$	$1/\beta_2$
1-4	65, 50	7	N/A
5-6	75, 55	10	1

Note that the mean of the normal distribution was chosen to be greater than 50, in order to ensure that we can test the scenario where a new attack becomes possible with the arrival of new information in the second stage. In order to get a reasonable number of random draws within the interval of $[0, 100]$, we kept the mean"not too high." The standard deviation was chosen based on satisfying the criterion for stage-one uniqueness, which necessitates that the precision of public information must be sufficiently lower than the precision of private information. We also, however, needed to keep the precision from being too low, again in order to get enough random draws within the interval $[0, 100]$.

4 Data Analysis

4.1 Variables and Summary Statistics

In our analysis, the main dependent variables are the size of the attack (measured as the fraction of subjects choosing action A) on the aggregate level, and the action (a binary choice variable, with 1 representing action A and 0 representing action B) on the individual level.

The main explanatory variables are the random number, θ, on the aggregate level and the subject-specific hint number, x, on the individual level. We also look at several other variables which can have an effect on outcomes. One of these controls is the *cost* of action A (attacking), which is equal to 20 or 50, depending on the treatment, in Sessions 1-4 and is always equal to 60 in Sessions 5-6. When we explore the impact of endogenous learning, we look at the effect of *stage* on actions, where stage takes on values of 1 or 2. In order to understand the effects of exogenous learning, we introduce a new-information dummy (*NI dummy*), which takes on a value of 1 in the treatment where subjects receive a more precise private signal in the second stage and a value of 0 otherwise. Finally, we look at subjects' expectations about the size of the attack by creating a *belief* variable.

Table AI in the Appendix provides descriptive statistics for the experiment.

4.2 First Period Predictions

In the first stage of each independent round, the subjects chose between actions A and B. Both the aggregate and the individual-level data confirm that observed behavior is consistent with monotone strategies. In order to see this on the aggregate level, we plot the size of the attack (i.e., the total number of players out of 15 choosing action A) against θ for the cost-50 treatment and find that A is strictly decreasing in θ, just as the theory predicts (see Figure 1). Moreover, the figure below shows that, for low states, almost everyone always chose action A, while for high states, almost everyone always chose action B. There is an intermediate range of fundamentals for which the size of the attack is decreasing in θ. We test this nonparametrically by carrying out a locally weighted regression of the number of attackers on the value of θ which is represented by the black monotonically decreasing line in Figure 2.[12] The monotonicity of the fitted line confirms the hypothesis.[13]

[12] The fact that we have very few negative draws of θ is the reason behind the slightly positive slope of the fitted line in that range. That is, the positive slope is driven by "mistakes." The fitted line is otherwise monotonic over the critical range of fundamentals.

[13] See Figures B1 and B2 in Appendix B for similar plots for the other two cost treatment conditions.

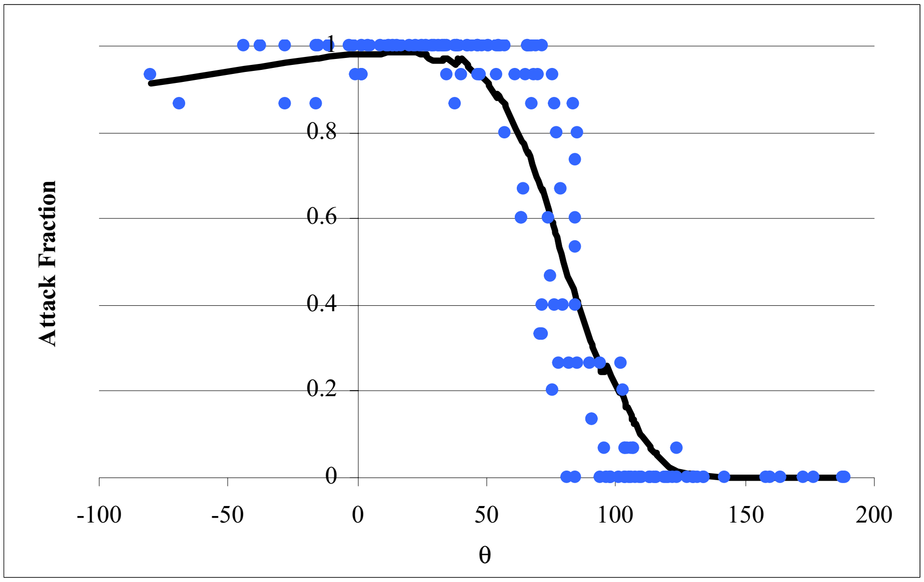

Figure 2: Kernel Regression (Cost-50 Treatment)

The analysis of the individual-level behavior in stage 1 also supports the theoretical predictions. Figure 3 was constructed by creating discrete bins for x and calculating the probability of choosing action A for each bin. The figure demonstrates that the probability of choosing action A is decreasing in the private signal over almost the entire range of x.[14]

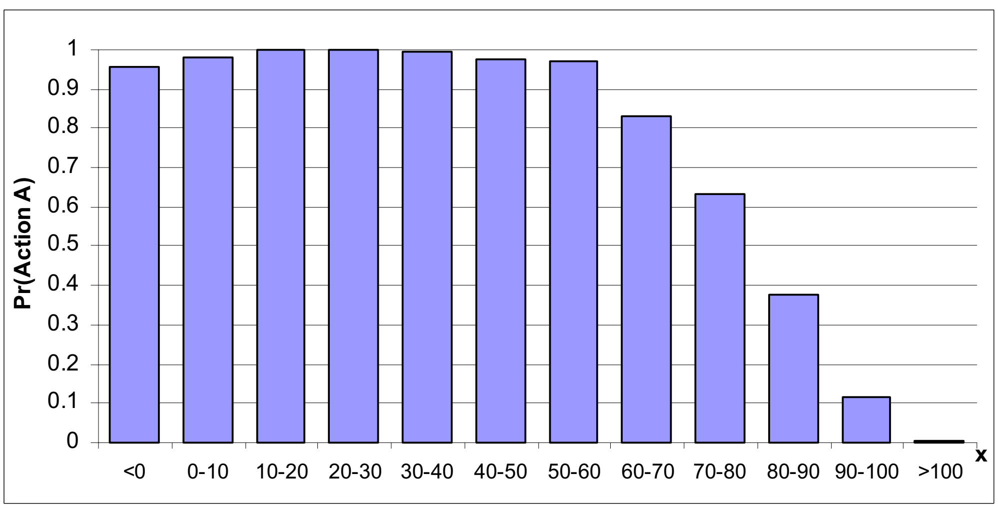

Figure 3: Probability of Action A vs. x (Cost-50 Treatment)

We also ran subject-level regressions of each individual's action on her private signal, x, and various controls. We use ordinary least squares to estimate the effects.[15] The

[14] See Figures B3 and B4 in Appendix B for similar plots for the other two cost treatment conditions.

[15] While a logit regression would be more appropriate given the binary nature of the dependent variable, we use a linear probability model since the logistic approach would result in biased estimates from a fixed-

results are reported in Table III.

Table III.
Stage 1 Individual Level Regressions (Pooled Data for All Sessions)

	Dependent Variable: Action		
	1	2	3
Private signal, x	-0.0065***	-0.0066***	-0.0012***
	(0.0001)	(0.0001)	(0.0001)
Cost of action A		-0.0026***	-0.0005**
		(0.0003)	(0.0002)
Belief			0.0646***
			(0.0011)
R^2	0.57	0.58	0.83
No. of observations	6000	6000	6000

Note: Robust standard errors in parentheses. Regressions include group and round fixed effects. For sessions 5 and 6, only the no-new-information treatment data are used. Significance levels: ** 5%, *** 1%.

According to Column 1, the effect of x on the choice of action is negative, as the theory predicts, and statistically significant at the 1 percent confidence level. This effect remains strong after controlling for other factors, such as the cost of action A and the beliefs about the size of the attack.

effects regression due to the panel structure of the dataset. Let p_{it} represent the probability that subject i attacks in period t, then $E[a_{it}] = 1\cdot\, p_{it} + 0 \cdot (1 - p_{it}) = p_{it}$. This is modeled as $p_{it} = \Pr[a_{it} = 1] = F(x'_{it}\beta{+}\mu_i)$. For a linear probability model, $F(x'_{it}\beta) = x'_{it}\beta{+}\mu_i$, and the usual panel data methods apply. That is, β can be consistently estimated by eliminating μ_i using the within transformation (demeaning the data). This is possible because the MLE of μ_i and β are asymptotically independent. The only issue is that $\widehat{a}_{it}$ is not guaranteed to lie in the unit interval.

In order to use the logistic approach, we need to define a threshold a^*_{it} such that

$$\begin{aligned} a_{it} &= 1 \text{ if } a^*_{it} > 0 \\ &= 0 \text{ if } a^*_{it} \leq 0, \end{aligned}$$

where $a^*_{it} = x'_{it}\beta + \mu_i + \nu_{it}$ with $Pr[a_{it} = 1] = Pr[a^*_{it} > 0] = Pr[\nu_{it} > -x'_{it}\beta - \mu_i] = F(x'_{it}\beta + \mu_i)$, where the last inequality holds as long as the density function for F is symmetric around zero. In this case, μ_i and β are unknown parameters and as $N \to \infty$, for a fixed T, the number of parameters μ_i increases with N. This means that μ_i cannot be consistently estimated for a fixed T. This is known as the incidental parameters problem.

The MLE of μ_i and β are no longer asymptotically independent with a qualitative limited dependent variable model with fixed T (like logit) as demonstrated by Chamberlain (1980). In his paper, Chamberlain proposes using the conditional logit approach to correct for this problem. For estimates using the conditional logit technique, see Table B-II in Appendix B.

Note also that the specifications in Table III include group and round fixed effects. We have also run regressions that include group, round, and subject fixed effects as a robustness check. The results are very close and are reported in Table B-III of Appendix B.

While subjects seem to be following threshold strategies, which supports the theoretical prediction of a unique equilibrium in the first period, we can learn more about the behavior of agents from estimating the magnitude of these thresholds and comparing them to the theoretical predictions. For every round of every session of the experiment, we run logit regressions of action on the private signal x, calculate predicted values from these regressions, and find the particular private signal, $\widehat{x}_{rs}$, for which the 45-degree line intersects the predicted values, where subscript r denotes round and subscript s denotes session. Finally, we average across rounds to get $\widehat{x}_s$, the estimated individual threshold for every session. We then use equation (1), derived in Section 2, to get the estimated aggregate threshold, $\widehat{\theta}_s$, for every session. Table IV reports the stage-one threshold estimates and compares them to the theoretical thresholds θ^* and x^*. Column 5 of the table reports the percentage of subjects who wrote that they followed a threshold strategy in their post-experiment questionnaire.

Table IV.
Threshold Summary.

Session	Cost	θ^*	$\widehat{\theta}$	x^*	$\widehat{x}$	Percent Using Thresholds
1, 2	20	81.5	78.6	87.8	84.3	93%
1,2	50	48.1	77.2	47.8	82.5	93%
3,4	20	81.5	79.2	87.8	85.3	95%
3,4	50	48.1	72.7	47.8	77.1	95%
5,6	60	34.8	69.2	30.9	72.9	92%

While the theoretical thresholds should be highly sensitive to the cost of attacking, the estimated thresholds decrease only slightly with cost. Even though we ensure that the subjects are well aware of the change in costs from treatment to treatment, they seem to not be very responsive to this change. While higher cost of attack does have a negative effect on the decision to attack on the individual level, as is evident from the negative and statistically significant coefficient on cost in Column 2 of Table III, the magnitude of this effect is clearly smaller than the theory predicts.

Most importantly, the estimated thresholds are considerably higher than the theoretical thresholds for the cost-50 and cost-60 treatments. This suggests that subjects are much more aggressive in their behavior than the theory would predict for the high-cost scenarios, but not for the low-cost ones.

We next seek to understand the trends in and the reasons behind the relative aggressiveness of subjects. Do thresholds change over rounds? Does the excess aggressiveness relative to theory arise due to "mistakes" in choices given beliefs, or due to "mistakes" in beliefs? To address the first question, we plot the estimated thresholds, $\widehat{\theta}$, across rounds.[16]

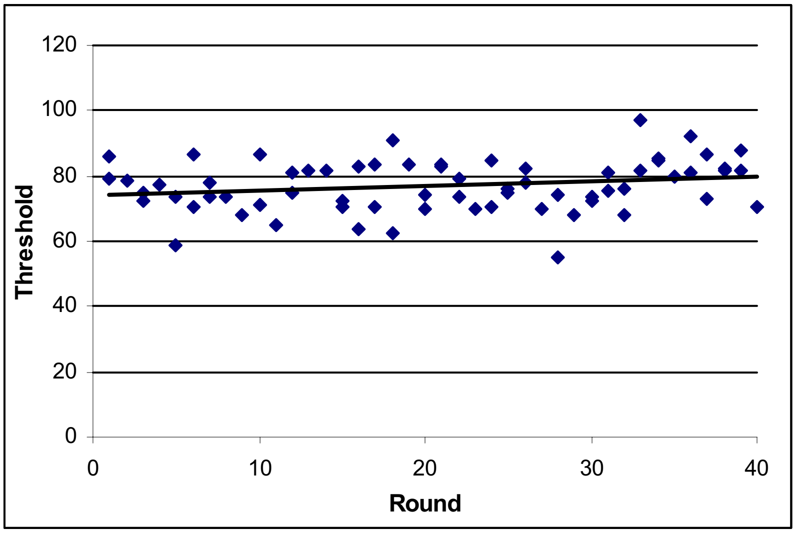

Figure 4: Estimated thresholds, $\widehat{\theta}$, vs. rounds (pooled data for sessions 1-4)

The estimated aggregate thresholds do not vary significantly with round, since the fitted black line exhibits only a slight upward trend. We can also see that the variance of thresholds does not decrease with rounds, which implies that subjects are not learning from round to round.

Regarding the second question, we can calculate the number of subjects who did not follow a best-response strategy given their private signal. In order to determine each subject's best response, we assume that each individual knows that others are using an average threshold $\widehat{x} > x^*$ in each round of the experiment. Assuming that each subject expects that others' follow the threshold, $\widehat{x}$, we next ask whether individual behavior is consistent with the theoretical best response strategy to this $\widehat{x}$. In order to answer this question, we compute a threshold $\widehat{\widehat{x}}$ that is the best response to $\widehat{x}$, which is simply the individual threshold that sets the posterior probability of regime change $Pr(\theta \leq \widehat{\theta}|\widehat{\widehat{x}})$ equal to the cost of attacking, c. Given this best-response threshold, $\widehat{\widehat{x}}$, a rational agent will attack if and only if her private signal, x, is below $\widehat{\widehat{x}}$ and refrain from attacking otherwise. We call strategies that do the opposite "mistakes." That is, a "mistake" would be an instance when a subject attacks when $x > \widehat{\widehat{x}}$ or does not attack when $x < \widehat{\widehat{x}}$.

We find that, on average, in approximately 91 percent of cases subjects followed a strategy that was a best response to the estimated threshold, $\widehat{x}$. Figures 5 and 6 show the number of "mistakes" by round for sessions 1-2 and 3-4, respectively.

[16] Note that the threshold $\widehat{\theta}$ is found using an equation equivalent to (1), $\Phi(\sqrt{\alpha}(\widehat{x} - \widehat{\theta})) = \widehat{\theta}$.

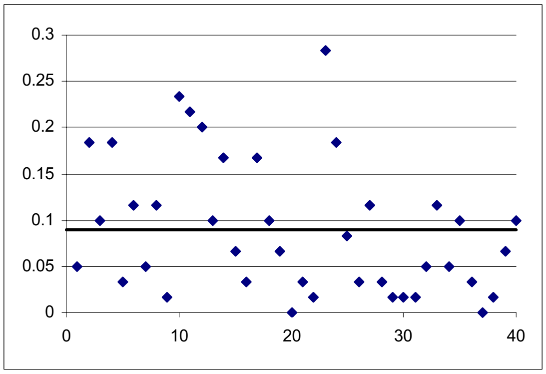

Figure 5: Proportion of "mistakes" vs. rounds (Sessions 1-2)

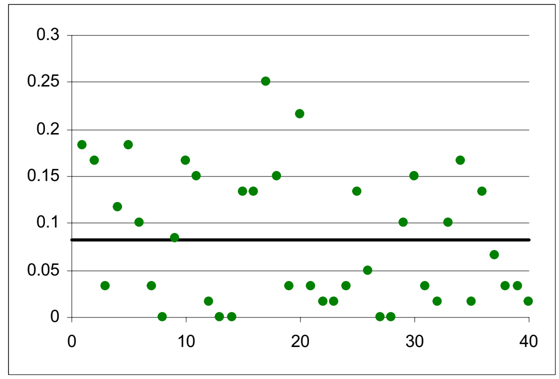

Figure 6: Proportion of "mistakes" vs. rounds (Sessions 3-4)

The figures above show that the number of "mistakes" relative to best-response does not change significantly with round. The dispersion also does not seem to decrease across rounds.[17] The solid black line represents the fraction of "mistakes" averaged across rounds.

4.3 Dynamic Predictions

4.3.1 Endogenous Learning

So far, we have analyzed the data from the first stage of the experiment. In this section, we will discuss the effects of learning across stages under the condition that subjects do not receive an additional signal in the second stage.

We find that the knowledge that the experiment has not ended in the first stage bears a strong effect on subjects' behavior in the second stage. In particular, we find that the average probability of attack in the second stage of the no-new-information treatments is much lower than the average probability of attack in stage one, as is shown in Figure 7.

[17]Using standard OLS techniques to fit a regression line through these data produces slope coefficients that are no statistically significantly different from zero. See Table B-VII in Appendix B for regression results.

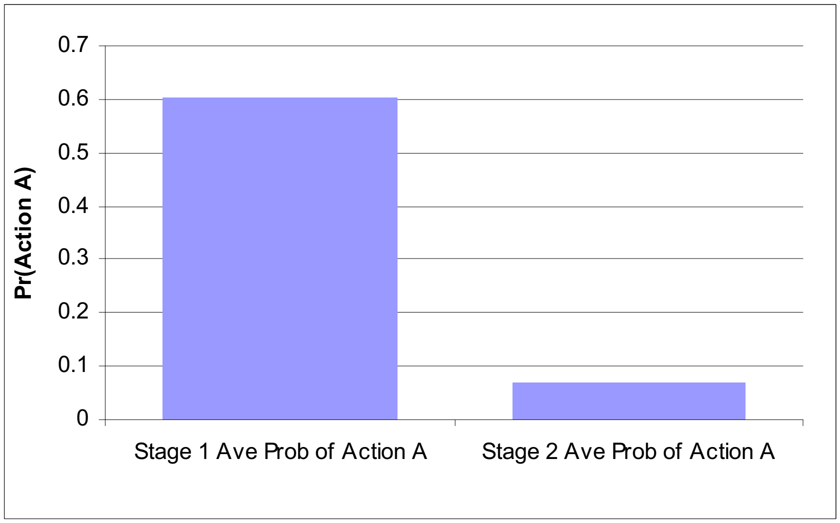

Figure 7: Average Probability of Action A for the No-New-Information Treatments

The same conclusion can be drawn from looking at the individual-level regressions of action on the private signal, x, that now include pooled data for stages one and two and add stage as a control variable. Table V reports that stage has a negative and highly statistically significant effect on action, which means that the probability of subjects choosing action A is reduced as we go from stage one to stage two.[18]

[18] Again, see Table B-IV in Appendix B for estimation using conditional logit.
Note also that the specifications in Table V include group and round fixed effects. We have also run regressions that include group, round, and subject fixed effects as a robustness check. The results are very close and are reported in Table B-V of Appendix B.

Table V.

Individual Level Regressions (Pooled Data for All Sessions)

	Dependent Variable: Action		
	1	2	3
Private signal, x	-0.006***	-0.006***	-0.001***
	(0.0001)	(0.0001)	(0.0001)
Cost of action A		-0.0026***	-0.0006**
		(0.0003)	(0.0002)
Belief			0.0632***
			(0.0011)
Stage	-0.2308***	-0.2272***	-0.0466***
	(0.0073)	(0.0073)	(0.0057)
R^2	0.61	0.62	0.84
No. of observations	8820	8820	8820

Note: Robust standard errors in parentheses. Regressions include group and round fixed effects. For sessions 5 and 6, only the no-new-information treatment data are used. Significance levels: ** 5%, *** 1%.

Figures 8 and 9 show the reduction in aggressiveness in the strategy space of the agents for cost-50 and cost-20 treatments, respectively.

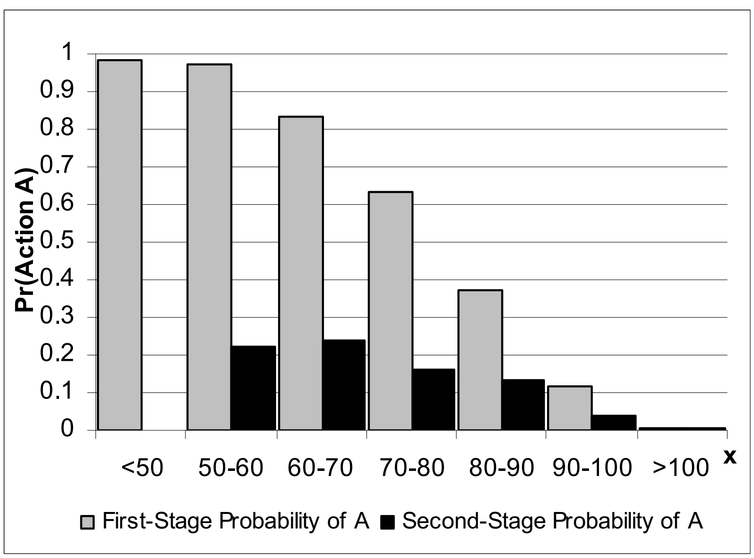

Figure 8: Probability of Attack vs. x by Stage for Cost-50 Treatments

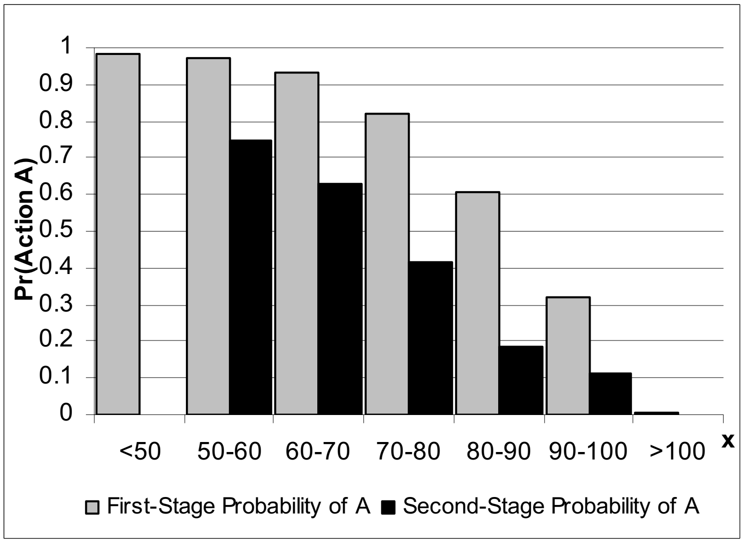

Figure 9: Probability of Attack vs. x by Stage for Cost-20 Treatments

In both figures, the probability of action A is lower in the second stage relative to the first stage. This shows that subjects exhibit some endogenous learning. However, the learning effect is not as strong as the model predicts, which suggests that agents continue to act overly aggressively in the second stage of the experiment. Note also that the reduction is substantially larger in the high-cost treatment, suggesting that the cost of action A has an important effect on aggressiveness in stage two.

4.3.2 New Information

In this section, we turn to the experimental treatments that explore the effects of providing the subjects with a new more precise private signal in the second stage. Here, the theory predicts that an attack becomes possible in stage two, given that the parameters have been chosen appropriately. Figure 10 contrasts the average probability of action A in the two stages of the experiment for the no-new-information (NNI) and the new-information (NI) treatments. Note that the figure was constructed using only the rounds that continued into the second stage and for which the random number drawn was below 100. This allows us to make the clearest possible comparison between treatments.

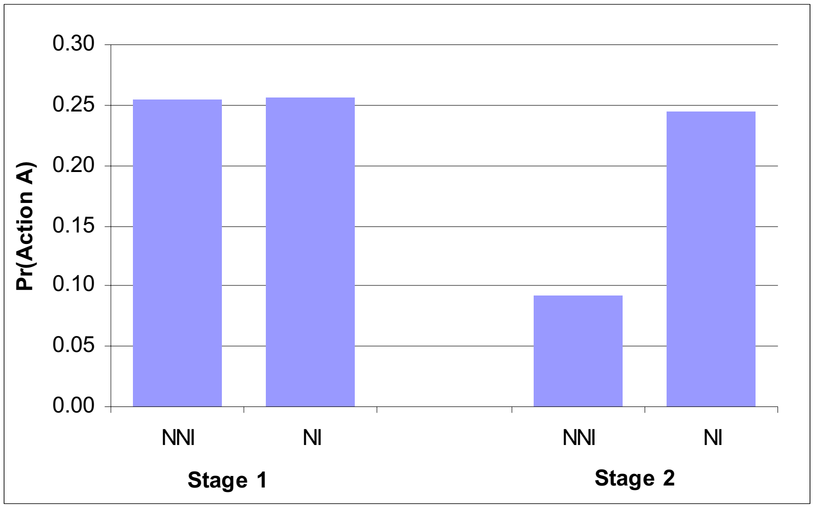

Figure 10: Average Probability of Action A for the NNI and the NI Treatments

First, consider the average probability of action A for the two treatment conditions in the first stage. The probabilities are very close in magnitude: 0.25 for the treatment with no new information in stage two and 0.26 for the one with new information in stage two. We can therefore conclude that a subject in the new-information treatment, who knows before the experiment begins that she will be receiving a more precise private signal in the second stage, does not wait to take action until the second stage, but rather behaves in a similar fashion to a subject in the no-new-information treatment. Secondly, note that the probability of action A is reduced dramatically in the NNI treatment, as the experiment continues into the second stage. However, we find that the average probability of action A in the second stage of the NI treatment is not statistically significantly different from the probability of action A in the first stage.

We confirm this result by running an individual-level regression of action in the second stage on the private signal, x, and the new-information treatment dummy. The results of this regression are reported in Table VI.[19] The statistically significant coefficient on the NI dummy tells us that subjects are more likely to choose action A in stage two of the NI treatment than in stage two of the NNI treatment, which is consistent with the theoretical prediction.

[19] For estimates using the conditional logit technique, see Table B-VI in Appendix B.

Table VI.
Stage 2 Individual Level Regressions (Sessions 5 and 6)

	Dependent Variable: Action
Private signal, x	-0.0021^{***}
	(0.0002)
NI dummy	0.0842^{***}
	(0.0145)
R^2	0.20
No. of observations	1395

Note: Robust standard errors in parentheses. Regressions include group, subject, and round fixed effects. Significance level: *** 1%.

Figure 11 shows the probability of action A in the strategy space of the agents for the NNI and NI treatments with cost of action A of 60. For most bins, the probability of action A in the second stage for the NI treatment exceeds the probability of action A for the NNI treatment. Moreover, for x above 80, the probability of action A in the second stage in the NI treatment exceeds the probability of action A in the first stage.

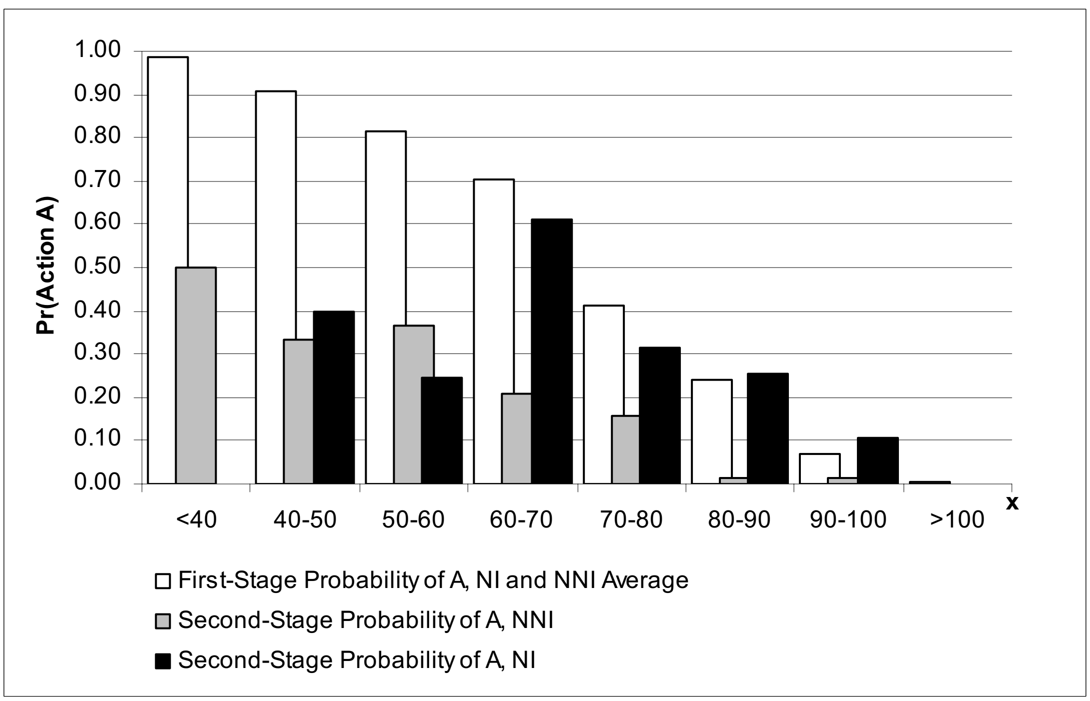

Figure 11: Probability of Action A vs. x by Stage for the NNI and NI Treatments

Recall that the theory predicts that, in addition to a no-attack equilibrium, there is a possibility of new attacks in the second stage if $\theta_1^* < \theta^\infty$. In order to test this prediction, we selected the parameters that would allow us to satisfy $\theta_1^* < \theta^\infty$, relying on the theoretical prediction for θ_1^* ($\theta_1^* = 34.8$ for $z = 75$ and $\theta^\infty = 40$). Given $\theta_1^* = 34.8$, our findings support the theory. However, given that the estimated threshold, $\widehat{\theta}_1$, is more aggressive than the theory predicts, this condition is not satisfied in the data ($69.2 > 40$). Conditional on the actual first-period threshold being above θ^∞, the theory predicts that no attacking should remain the unique equilibrium.[20]

In summary, the evidence suggests that the quantitative predictions about the thresholds are rejected in both stages of the experiment due to the subjects' aggressive actions. However, the evidence supports the qualitative theoretical predictions, namely, that the mass of agents choosing action A is monotonically decreasing in θ and that endogenous learning reduces the probability of choosing action A in the second stage, while new information (exogenous learning) increases this probability.

4.4 Rationality and Consistency of Beliefs

4.4.1 Rationality

So far, we have shown that the data seem to support the qualitative predictions of the theory; however, we see that the behavior of experimental subjects is more aggressive than predicted in both the static and dynamic frameworks. In this section, we seek to understand whether the aggressiveness results from behavior that is irrational given the agents' beliefs, and whether the agents' actions are consistent with their expectations of the actions of others.

First, we address the question of rationality by testing two hypotheses. The first hypothesis postulates that agents make "mistakes" or act "too aggressively," given their beliefs of what other agents will do. This pattern of behavior can be explained as a "rule of thumb" strategy, where subjects attack all the time simply because attacking is perceived to be more "fun" for some reason. This should be especially the case whenever the cost of action A is relatively small. An alternative hypothesis is that agents do not make "mistakes" given their beliefs of what others will do, and it is their beliefs that are more aggressive that the theory predicts.

Recall that before revealing the outcome of the experiment at a particular stage, we ask each subject about his or her belief as to the number of other group members choosing action A. Figure 12 plots these beliefs about the fraction of agents choosing action A (the attack fraction) and the theoretical benchmark for these beliefs over x_i in the range

[20] If, however, the theory were to be extended to capture first-stage aggressiveness, the predictions for the second stage would change.

[0, 100] for Cost-20, Cost-50, and Cost-60 treatments. The theoretical benchmark is the function $g(x_i, x^*) = E\left[\Phi\left(\frac{x^*-\theta}{\sigma_x}\right)|x_i\right]$, where x^* is the theoretical threshold derived in section 2.1. That is, $g(x_i, x^*)$ gives the expectation of the mass of agents choosing action A held by an agent with signal x_i who expects all other agents to use the threshold x^*. The actual reported beliefs are a function $\widehat{g}(x_i)$, which we estimate from the actual data using a kernel regression approach.

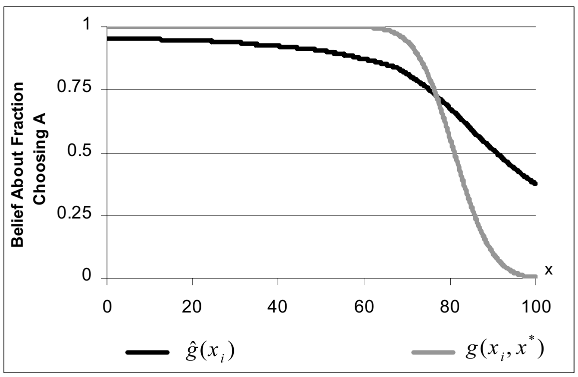

Cost 20

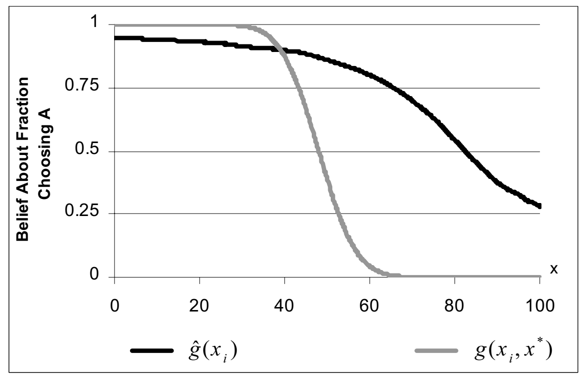

Cost 50

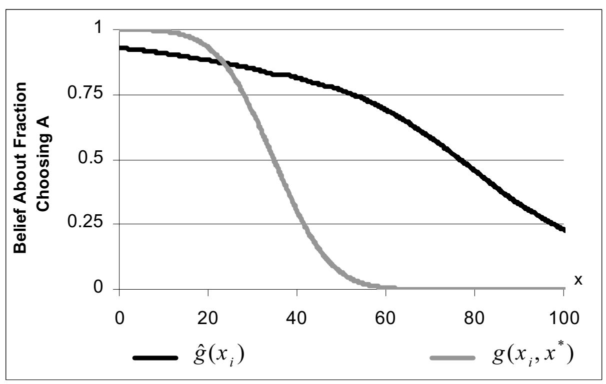

Cost 60

Figure 12: Beliefs about Fraction of Agents Choosing Action A vs. Theory Prediction

In sections 4.2 and 4.3, we found that the agents' actions are more aggressive than the theory predicts. Figure 12 shows that the subjects' *beliefs* are more aggressive than predicted by the theory in the cost-50 and cost-60 treatments. (The agents' expectations of the size of the attack lie above/to the right of the theoretical expectation for higher values of x_i.) These are exactly the cost treatments for which we find more aggressive actions relative to theory. However, the beliefs in the cost-20 treatment do not seem to be much more aggressive than the theoretical prediction. Looking back at Table 4, agents' actions are not more aggressive than the theoretical prediction in the cost-20 treatment. Note also that the plots in Figure 12 demonstrate that the observed beliefs are less sensitive to the individual signal than the theory predicts, since $\widehat{g}(x_i)$ is flatter than $g(x_i, x^*)$. This feature of the observed beliefs is present in all cost treatments.

Figure 12 provides visual support for the view that the subjects of these experiments do not act aggressively merely due to irrationality, but rather due to a rational response to their overly aggressive expectations. However, we need to provide still more concrete evidence. In particular, we can construct a measure of rationality based on our data of the subjects' beliefs. A rational agent will choose action A if and only if the expected payoff from choosing it is greater than the cost. However, for the sake of this comparison, we would actually need to know the subject's belief about the probability of a successful attack, not his or her belief about the size of the attack.[21] This, however, is not a question a typical subject is capable of answering.

Thus, we compute $\overline{x}_i$ such that $\widehat{g}(x_i) = g(x_i; \overline{x}_i)$. In other words, for every subject i with signal x_i, we look for a value $\overline{x}_i$ that minimizes the distance between the two curves in Figure 12 for each cost treatment. The threshold $\overline{x}_i$ is the value of x that rationalizes subject i's belief about the size of the attack. These values are graphed across x_i for different costs of attacking in Figure 13.

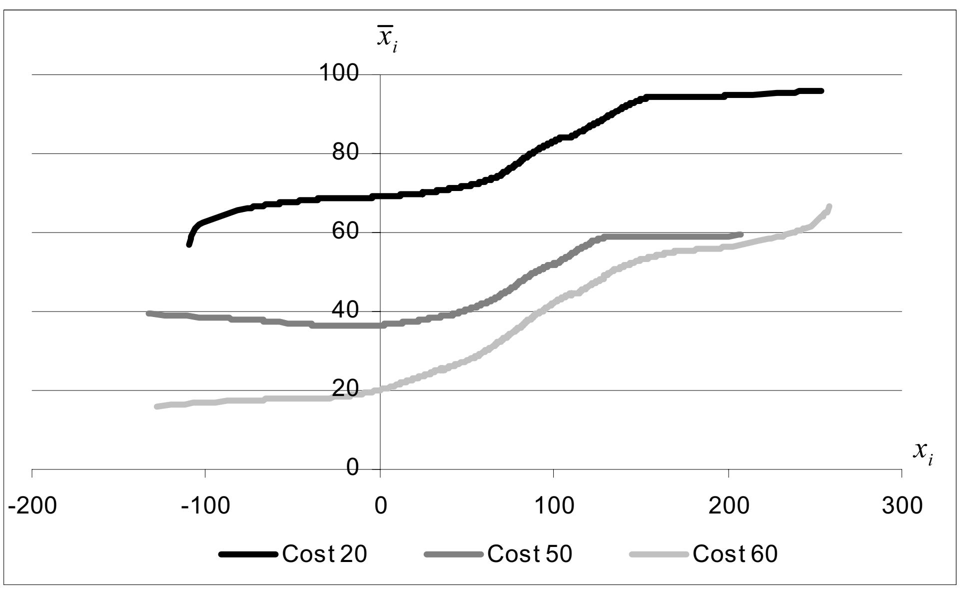

Figure 13: Thresholds $\overline{x}_i$ for different cost treatments

The values of $\overline{x}_i$ decrease with the cost of attacking. Within each cost treatment, $\overline{x}_i$ increases with the individual private signal, x_i. This suggests that there is a positive

[21] That is, the relevant condition for rationality is to choose action A if and only if $E[A(\theta)|x_i] - (\frac{\alpha}{\alpha+\beta}x_i + \frac{\beta}{\alpha+\beta}z) > c$. Rather, instead of $E[A(\theta)|x_i]$ in expression, the proper term is the expectation of $I_{A(\theta)>\theta}$, where $I_{A(\theta)>\theta} = I_{\theta<\theta^*} = \left\{\begin{smallmatrix} 1 \text{ if } \theta<\theta^* \\ 0 \text{ if } \theta>\theta^* \end{smallmatrix}\right.$.

relationship between the individual agent's error in beliefs about the actions of others and the agent's expectation about θ. Agents with a high private signal x_i expect others to use less aggressive strategies (i.e., have a higher $\overline{x}_i$).

Given this $\overline{x}_i$, we compute the corresponding value of $\overline{\theta}_i$, using equation (1). Finally, in order to see which decisions were rational, we use the fact that a rational agent would always choose action A if and only if $\Phi\left(\sqrt{\beta_1+\alpha}(\overline{\theta}_i - \frac{\beta_1}{\beta_1+\alpha}\overline{x}_i - \frac{\alpha}{\beta_1+\alpha}z)\right) > c$ and choose action B otherwise. Using this criterion, we find that subjects are rational in 76.98 percent of cases in the cost-20 treatment, 90.79 percent of cases in the cost-50 treatment, and 89.44 percent of cases in the cost-60 treatment.

The above evidence leads us to reject the hypothesis that subjects are making "mistakes" or acting "too aggressively" given their beliefs. The data suggest that the subjects' beliefs are overly aggressive. However, given these beliefs, agents do not make "mistakes" (i.e., act "rationally") in 86 percent of cases, on average.

4.4.2 Consistency

Secondly, we address the issue of consistency by asking whether the subjects' actions are consistent with their beliefs about the size of the attack. Ideally, we would like to compare each individual's expectation of the fraction of agents attacking given his or her individual signal, x, to the actual fraction of agents attacking given x. However, the latter metric is not available (that is, we can only plot the aggregate attack against θ, not x). However, we can employ the law of iterated expectations, in order to provide a measure of consistency. In particular, the theory tells us that the following equality must hold by the law of iterated expectations:

$$E[A(\theta)] = E\left[E[A(\theta)|x]\right] \tag{5}$$

Thus, we can test whether actions are consistent with individual beliefs by estimating the right- and the left-hand sides of equation (5). The left-hand side of equation (5) is the expectation of the actual fraction of agents choosing action A, which we compute by integrating numerically the size of the attack that we get from the data, weighted by the distribution of θ. On the right-hand side of (5), $E[A(\theta)|x]$ is the subject's belief about the fraction of agents choosing action A. To compute $E\left[E[A(\theta)|x]\right]$, we average beliefs across x. That is, we integrate numerically the function $\widehat{g}(x_i)$, weighted by the distribution of x. Table VII presents the results from this analysis.

Table VII.
Test of Consistency in Stage 1.

Cost Treatment	Average Realized Attack Fraction	Beliefs Averaged Across x
20	0.5857	0.5354
50	0.5799	0.5347
60	0.5071	0.4615

To construct the above table, we use stage-one data for the three cost treatments. Note that the values reported in Columns 2 and 3 of Table VII are close to one another, which serves as evidence that actual outcomes are consistent with beliefs.

Next, we perform a consistency test for the second stage of the experiment (Table VIII).

Table VIII.
Test of Consistency in Stage 2.

Cost & Info Treatment	Average Realized Attack Fraction	Beliefs Averaged Across x
20 NNI	0.0227	0.0610
50 NNI	0.0769	0.2101
60 NNI	0.0461	0.1505
60 NI	0.1248	0.1484

Note that there is a reversal in aggressiveness: in the first stage, agents' beliefs are slightly less aggressive than their actions, while in the second stage, agents' actions are less aggressive relative to their beliefs, especially in the high-cost treatments with no new information. The subjects seem to think that others will be daring and attempt another attack, while they themselves behave cautiously. There is also a difference between the new-information and the no-new-information treatments within the cost-60 treatment. The subjects beliefs are approximately the same, but actions are vastly different. Agents act much more aggressively upon arrival of new information than without it, just as the theory predicts.

5 Conclusion

This paper uses a laboratory experiment to examine how learning affects the dynamics of coordination in a global-game environment. The main goal of this study is to understand the interaction between the dynamics of expectations about the state of the fundamentals and the dynamics of coordination. In particular, we focus on the effects of "endogenous learning" (learning over time from past regime outcomes) and "exogenous learning" (learning by receiving private information from outside sources)

The experimental evidence supports the theoretical predictions of our model. In the first stage of the experiment, we find that the size of the attack is monotonically decreasing in the strength of the economic fundamentals, θ, and that the agents' strategies are decreasing in their individual private signals. In the second stage, we find that the probability of attack is significantly higher in treatments where subjects receive an additional private signal (treatments with "endogenous" and "exogenous learning") than in treatments where agents do not receive an additional private signal (treatments with "endogenous learning" only).

However, we also observe that subjects display excess aggressiveness in their actions in both stages. In the first stage, the estimated thresholds followed by the subjects are higher than the theoretical thresholds in the high-cost treatments. Moreover, while the cost of attacking does decrease subjects' aggressiveness significantly, the reduction is not as large in magnitude as the theory predicts. In the second stage, the learning effect is not as strong as the model predicts.

In order to understand the source of this excessively aggressive behavior, we test two hypotheses. The first posits that the subjects make "mistakes" or act too aggressively, given their beliefs about the actions of others. The second states that agents do not make "mistakes" or act too aggressively, given their beliefs about the actions of others. This test of the theory of global games proves to be useful, since in these models actions are driven by the formation of expectations about other players' actions. We find that agents do not seem to make "mistakes" or act too aggressively, given their beliefs about the actions of other subjects.

This study can be extended along several dimensions. Experimentally, the framework can be extended to have more than just two periods, in order to shed some light on the timing of crises. As another extension of this experimental framework, it would be interesting to see if communication plays a significant role for equilibrium selection in the dynamic global game setting, and multiplicity detection in particular.

Finally, we find evidence that agents' expectations about the actions of others are overly aggressive relative to theory. Thus, a natural extension of the theory seems to involve exploring theoretical reasons for this excess aggressiveness and incorporating them into our models. Extending the theoretical literature in this direction may therefore prove to be a fruitful line of research.

References

[1] Angeletos, George-Marios, Hellwig, Christian, and Alessandro Pavan (2006). "Dynamic Global Games of Regime Change: Learning, Multiplicity, and Timing of Attacks." *Econometrica, forthcoming.*

[2] Angeletos, George-Marios and Ivan Werning. (2006) "Crises and Prices: Information Aggregation, Multiplicity, and Volatility." *American Economic Review* 96 (5) (December):1720-1737.

[3] Atkeson, Andrew. (2000). "Discussion of Morris and Shin's 'Rethinking Multiple Equilibria in Macroeconomic Modelling'." National Bureau of Economic Research *Macroeconomics Annual*.

[4] Baltagi, Badi H. (1995) *Econometric Analysis of Panel Data*. John Wiley & Sons.

[5] Boston, Robert. (2003). "Contagion Without Borders: The Lessons Taught by the 'Tequila Effect' or the 'Brazilian Sneeze'." 2003 Global Investment Conference.

[6] Brunnermeier, Markus K. and John Morgan (2004). "Clock Games: Theory and Experiments." manuscript, Haas School, UC Berkeley.

[7] Cabrales, Antonio, Nagel, Rosemarie, and Roc Armenter (2003). "Equilibrium Selection Though Incomplete Information in Coordination Games: An Experimental Study." Discussion Paper No. 601, Universitat Pompeu Fabra, Barcelona.

[8] Carlsson, Hans and Eric van Damme (1993a). "Global Games and Equilibrium Selection." *Econometrica* 61, 989-1018.

[9] Carlsson, Hans and Eric van Damme (1993b). "Equilibrium Selection in Stag Hunt Games" in *Frontiers of Game Theory*, ed. by K. Binmore, A. Kirman, and P. Tani. Cambridge, MA: MIT Press, 237-253.

[10] Chamberlain, Gary. (1980) "Analysis of Covariance with Qualitative Data." *Review of Economic Studies* 47, 225-238.

[11] Chamley, Christophe. (1999). "Coordinating Regime Switches." *The Quarterly Journal of Economics*, 114(3), 869-905.

[12] Cheung, Yin-Wong and Daniel Friedman (2006). "Speculative Attacks: A Laboratory Study in Continuous Time." Working paper, Dept. of Economics, University of California, Santa Cruz.

[13] Cooper, Russell W., DeJong, Douglas V., Forsythe, Robert, and Ross, Thomas W. (1990). "Selection Criteria in Coordination Games: Some Experimental Results." *American Economic Review* 80 (1) (March): 218 – 33.

[14] Cooper, Russell W., DeJong, Douglas V., Forsythe, Robert, and Thomas W. Ross. (1992). "Communication in Coordination Games." *Quarterly Journal of Economics* 107 (2) (May): 739-71.

[15] Corsetti, Giancarlo, Guimaraes, Bernardo, and Nouriel Roubini. (2003). "International Lending of Last Resort and Moral Hazard: A Model of IMF's Catalytic Finance." NBER Working Papers 10125, National Bureau of Economic Research, Inc.

[16] Kristin Forbes. (2000). "The Asian Flu and Russian Virus: Firm-level Evidence on How Crises are Transmitted Internationally." NBER Working Papers 7807, National Bureau of Economic Research, Inc.

[17] Freedman, Amy L. (2004). "Economic crises and political change: Indonesia, South Korea, and Malaysia." Heldref Publications and Gale Group, http://www.findarticles.com/p/articles/mi_m2393/is_4_166/ai_114631686/

[18] Goldstein, Itay, and Ady Pauzner. (2001). "Contagion of Self-Fulfilling Financial Crises due to Diversification of Investment Portfolios." working paper, Duke University and Tel Aviv University.

[19] Edmond, Chris. (2005). "Information and the Limits to Autocracy." working paper, NYU Stern.

[20] Heinemann, Frank, Nagel, Rosemarie, and Peter Ockenfels (2004). "The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information." *Econometrica* 72 (5) (September): 1583-1599.

[21] International Monetary Fund. (2000). Recovery from the Asian Crisis and the Role of the IMF.

[22] Krugman, Paul R. (1979). "A Model of Balance of Payments Crises." *Journal of Money, Credit, and Banking* 11 (August): 311 – 25.

[23] Morris, Stephen and Hyun Song Shin. (1998). "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review* 88 (3) (June): 587 – 97.

[24] Morris, Stephen and Hyun Song Shin. (2004). "Coordination Risk and the Price of Debt." *European Economic Review* 48, 133-153.

[25] Obstfeld, Maurice. (1996). "Models of Currency Crises with Self-Fulfilling Features." *European Economic Review* 40 (3 – 5) (April): 1037 – 47.

[26] Rochet, Jean-Charles and Xavier Vives. (2004). "Coordination Failures and the Lender of Last Resort: Was Bagehot Right After All?" working paper, University of Toulouse.

[27] Van Huyck, John B., Battalio, Raymond C., and Richard O. Beil. (1990). "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure." *American Economic Review* 80 (1) (March): 234-48.

6 Appendix A: Derivations

Let us solve for some of these equilibrium values and do some comparative statics that will help with the experiment. First, we solve for the equilibrium threshold in period 1, $\theta_1^*(z)$, for the special case where cost of attacking, c, equals $\frac{1}{2}$ (as was the case in some of the experimental treatments)

$$\theta_1^* = \Phi(\sqrt{\beta_1}(x_1^* - \theta_1^*)). \tag{6}$$

Solving (6) for x_1^*, we get

$$x_1^* = \theta_1^* + \beta_1^{-1/2}\Phi^{-1}(\theta_1^*).$$

Now we substitute x_1^* into the other equilibrium equation, namely

$$\Phi\left(\sqrt{\beta_1+\alpha}(\theta_1^* - \frac{\beta_1}{\beta_1+\alpha}x_1^* - \frac{\alpha}{\beta_1+\alpha}z)\right) = 1/2$$

$$\Phi\left(\sqrt{\beta_1+\alpha}(\theta_1^* - \frac{\beta_1}{\beta_1+\alpha}\left[\theta_1^* + \beta_1^{-1/2}\Phi^{-1}(\theta_1^*)\right] - \frac{\alpha}{\beta_1+\alpha}z)\right) = 1/2.$$

Putting the terms that contain θ_1^* on the left-hand side of the expression, we get

$$\frac{\alpha}{\beta_1+\alpha}\theta_1^* - \frac{\beta_1^{1/2}\Phi^{-1}(\theta_1^*)}{\beta_1+\alpha} = \frac{\Phi^{-1}(1/2)}{\sqrt{\beta_1+\alpha}} + \frac{\alpha}{\beta_1+\alpha}z. \tag{7}$$

We can simplify this expression further, using the fact that $\Phi^{-1}(1/2) = 0$ and multiplying both sides by $\frac{\beta_1+\alpha}{\alpha}$

$$\theta_1^* - \frac{\beta_1^{1/2}}{\alpha}\Phi^{-1}(\theta_1^*) = z. \tag{8}$$

Solving for θ_1^* gives us the equilibrium threshold in the first round, $\theta_1^*(z)$, such that the regime will collapse in the first round ($R_1 = 1$) if and only if $\theta \leq \theta_1^*$. Such a solution exists and is unique if and only if the following relationship holds for the precisions: $\beta_1 \geq \alpha^2/(2\pi)$. To see this, we define

$$G(\theta_1^*(z), z) \equiv z - \theta_1^* + \frac{\beta_1^{1/2}}{\alpha}\Phi^{-1}(\theta_1^*) = 0.$$

Note that $G(\theta_1^*(z), \cdot)$ is continuous and differentiable in $\theta \in (0,1)$, and that $G(0,z) = -\infty$ and $G(1,z) = \infty$, which implies that there necessarily exists a solution and any solution satisfies $\theta_1^*(z) \in (0,1)$. This establishes existence. To prove uniqueness, note that

$$\frac{\partial G(\theta_1^*(z),z)}{\partial \theta_1^*} = \frac{\beta_1^{1/2}}{\alpha}(\Phi^{-1})'(\theta_1^*) - 1.$$

We can re-write this using the formula for the derivative of an inverse function:

$$\left(\Phi^{-1}\right)'(\theta_1^*) = \frac{1}{\Phi'(\Phi^{-1}(\theta_1^*))} = \frac{1}{\phi(\Phi^{-1}(\theta_1^*))},$$

where $\phi(\cdot)$ is the p.d.f. of the standard normal distribution and is bounded by $\frac{1}{\sqrt{2\pi}}$ (i.e. $\max_{\omega \in \mathbb{R}} \phi(\omega) = \frac{1}{\sqrt{2\pi}}$). Therefore,

$$\frac{\partial G(\theta_1^*(z),z)}{\partial \theta_1^*} = \frac{\beta_1^{1/2}}{\alpha}\frac{1}{\phi(\Phi^{-1}(\theta_1^*))} - 1 > \frac{\beta_1^{1/2}}{\alpha}\sqrt{2\pi} - 1.$$

Then if $\frac{\beta_1^{1/2}}{\alpha}\sqrt{2\pi} - 1 \geq 0$, or if $\frac{\beta_1}{\alpha^2} \geq \frac{1}{2\pi}$, the function G is strictly increasing in θ_1^* ($\frac{\partial G(\theta_1^*(z),z)}{\partial \theta_1^*} > 0$), which implies a unique solution to (8).

We can use the Implicit Function Theorem to demonstrate that the threshold $\theta_1^*(z)$ is monotonically decreasing in z. Let $F(\theta_1^*(z), z)$ be defined as

$$F(\theta_1^*(z),z) \equiv \frac{\alpha}{\beta_1+\alpha}\theta_1^* - \frac{\beta_1^{1/2}\Phi^{-1}(\theta_1^*)}{\beta_1+\alpha} - \frac{\alpha}{\beta_1+\alpha}z = 0$$

$$\frac{\partial \theta_1^*}{\partial z} = -\frac{\partial F/\partial z}{\partial F/\partial \theta_1^*} = -\frac{-\frac{\alpha}{\beta_1+\alpha}}{\frac{\alpha}{\beta_1+\alpha} - \frac{\beta_1^{1/2}}{\beta_1+\alpha}(\Phi^{-1})'(\theta_1^*)} = \frac{1}{1-\frac{\beta_1^{1/2}}{\alpha}(\Phi^{-1})'(\theta_1^*)}. \tag{9}$$

As we have shown above, the derivative of the inverse of the c.d.f. of the standard normal is positive and reaches its minimum at $\sqrt{2\pi}$. We also know that the relationship between the precisions is $\frac{\beta_1^{1/2}}{\alpha} \geq \frac{1}{\sqrt{2\pi}}$. Thus, the whole fraction in (9) is negative (i.e., $\frac{\partial \theta_1^*}{\partial z} \leq 0$). Intuitively, θ_1^* is decreasing in z because when the public signal (z) has a high mean, the fundamentals are relatively good. So, the region where the attack will be successful in the first period is relatively small. Thus, the threshold theta is low. In other words, when the mean of the prior is high, the agents are initially pessimistic about their ability to overthrow the regime. So, in the first period, the size of the attack is relatively small. Then, in the second period, if the agents get a sufficiently precise private signal, an attack becomes possible. (That is, agents can become optimistic about their ability to change the status quo.) This is why this scenario can lead to multiplicity.

We can also verify that $\theta_1^*(1/2) = 1/2$ (but only if the public signal is completely uninformative relative to the private signal, that is if $\alpha/\beta_1 \to 0$). Let us substitute 1/2 into equation (7):

$$\frac{\alpha}{\beta_1+\alpha}\left(\frac{1}{2}\right) - \frac{\beta_1^{1/2}\Phi^{-1}(1/2)}{\beta_1+\alpha} = \frac{\Phi^{-1}(1/2)}{\sqrt{\beta_1+\alpha}} + \frac{\alpha}{\beta_1+\alpha}\left(\frac{1}{2}\right)$$

$$-\frac{\beta_1^{1/2}}{\beta_1+\alpha} = \frac{1}{\sqrt{\beta_1+\alpha}}.$$

Squaring both sides

$$\frac{\beta_1}{(\beta_1+\beta)^2} = \frac{1}{\beta_1+\beta}$$

$$\frac{1}{1+\beta/\beta_1} = 1.$$

which is true for all $\beta_1$if $\alpha/\beta_1 \to 0$ or equivalently if $\beta_1/\alpha \to \infty$. To put this differently, we just found the Morris-Shin limit threshold. The Morris-Shin limit is the limit as the ratio of precisions of private and public information approaches infinity, or in other words private information becomes infinitely precise relative to public information. It is

$$\lim_{\frac{\beta_1}{\alpha}\to\infty} \theta_1^*(z) = \frac{1}{2} = 1 - c \equiv \theta_\infty.$$

Finally, we employ the Implicit Function Theorem again to show that $\theta_1^*(z)$ is monotonic in β_1. Define

$$H(\theta_1^*(z), z) \equiv \theta_1^* - \frac{\beta_1^{1/2}}{\alpha}\Phi^{-1}(\theta_1^*) - z = 0$$

$$\frac{\partial\theta_1^*}{\partial\beta_1} = -\frac{\partial H/\partial\beta_1}{\partial H/\partial\theta_1^*} = \frac{\frac{1}{2\alpha}\beta_1^{-1/2}\Phi^{-1}(\theta_1^*)}{1 - \frac{\beta_1^{1/2}}{\alpha}(\Phi^{-1})'(\theta_1^*)}. \tag{10}$$

We already know that the denominator of (10) is always negative. Let us focus on the numerator. The inverse c.d.f. of the normal has the following property:

$$\begin{aligned} \Phi^{-1}(\theta_1^*) &< 0 \text{ if } \theta_1^* < 1/2 \\ \Phi^{-1}(\theta_1^*) &> 0 \text{ if } \theta_1^* > 1/2. \end{aligned}$$

Therefore, to sign this fraction, we need to consider our two cases. In Case 2′, z is low ($z < 1/2$), which implies that $\theta_1^*(z) > 1/2$ because we have proved above that $\theta_1^*(z)$ is monotonically decreasing in z and that $\theta_1^*(1/2) = 1/2$. This implies that $\Phi^{-1}(\theta_1^*) > 0$, and the numerator of (10) is positive (i.e., $\theta_1^*(z)$ is decreasing in β_1). In Case 2$''$, z is high, that is $z > 1/2$, so we know that $\theta_1^*(z) < 1/2$. In this case, $\Phi^{-1}(\theta_1^*) < 0$, the numerator of (10) is negative (i.e., $\theta_1^*(z)$ is increasing in β_1). This shows that $\theta_1^*(z)$ is monotonic in β_1.

Iterated Dominance Argument.

We have established there exists a unique monotone equilibrium in the first stage whenever the noise in private information is small enough. This result, however, leaves open the possibility that there are other non-monotone equilibria. We now show that there is no other equilibrium and, what is more, that the equilibrium is dominance solvable.

For simplicity, consider the special case where $\alpha = 0$. This allows us to eliminate the dependence on z and denote the strategy by $a(x_1)$ and the aggregate size of the attack by $A(\theta)$.

For any $\widetilde{x}_1 \in [-\infty, +\infty]$, let $A_{\widetilde{x}_1}(\theta)$ denote the mass of agents attacking (choosing action A) when (almost every) agent chooses action A if and only if $x_1 \leq \widetilde{x}_1$. Next, we define the function

$$V(x_1, \widetilde{x}_1) = E[U(1, A_{\widetilde{x}_1}(\theta), \theta) - U(0, A_{\widetilde{x}_1}(\theta), \theta)|x_1],$$

which represents the utility difference between choosing action A and choosing action B for an agent who has a private signal x_1 and expects the other agents to attack if and only if their signals fall below $\widetilde{x}_1$. From the model,

$$A_{\widetilde{x}_1}(\theta) = \Phi(\sqrt{\beta_1}(\widetilde{x}_1 - \theta))$$

and

$$V(x_1, \widetilde{x}_1) = y - y\Phi(\sqrt{\beta_1}(x_1 - \widetilde{\theta}_1)) - c,$$

where $\theta = \widetilde{\theta}_1(\widetilde{x}_1)$ is the unique solution to $A_{\widetilde{x}_1}(\widetilde{\theta}_1) = \widetilde{\theta}_1$, or equivalently the inverse of

$$\widetilde{x}_1 = \widetilde{\theta}_1 + \beta_1^{-1/2}\Phi^{-1}(\widetilde{\theta}_1).$$

Note that $\widetilde{\theta}_1$ is increasing in $\widetilde{x}_1$, which implies that $V(x_1, \widetilde{x}_1)$ is increasing in $\widetilde{x}_{1.}$ That is, the more aggressive the other agents are, the higher the payoff from choosing action A. Furthermore, $V(x_1, \widetilde{x}_1)$ is decreasing in x_1: the higher the private signal, the lower the expected payoff from choosing action A.

Next, note that $V(x_1, \widetilde{x}_1)$ is continuous in x_1 and satisfies $V(x_1, \widetilde{x}_1) \to y - c. > 0$ as $x_1 \to -\infty$ and $V(x_1, \widetilde{x}_1) \to -c < 0$ as $x_1 \to +\infty$. We can therefore define a function

$h(\cdot)$ such that $x_1 = h(\widetilde{x}_1)$ is the unique solution to $V(x_1, \widetilde{x}_1) = 0$ with respect to x_1. That is, when agents $j \neq i$ choose action A if and only if $x_{j1} \leq \widetilde{x}_1$, agent i finds it optimal to choose action A if and only if $x_{i1} = h(\widetilde{x}_1)$. Since $V(x_1, \widetilde{x}_1)$ is continuous in both arguments, decreasing in x_1 and increasing in $\widetilde{x}_1$, the function $h(\widetilde{x}_1)$ is continuous and increasing in $\widetilde{x}_1$. Finally, note that $h(\widetilde{x}_1)$ has a unique fixed point $x_1^* = h(x_1^*)$ and this fixed point is indeed the threshold x_1^* of the unique monotone equilibrium that we constructed in section 2.1.

Now, construct a sequence $\{\underline{x}_{1,k}\}_{k=0}^{\infty}$ with $\underline{x}_{1,0} = -\infty$ and $\underline{x}_{1,k} = h(\underline{x}_{1,k-1})$ for all $k \geq 1$. In particular, letting $\underline{\theta}_{1,k-1}$ be the solution to

$$\underline{x}_{1,k-1} = \underline{\theta}_{1,k-1} + \beta_1^{-1/2}\Phi^{-1}(\underline{\theta}_{1,k-1}),$$

we have

$$V(x_1, \underline{x}_{1,k-1}) = y - y\Phi(\sqrt{\beta_1}(x_1 - \underline{\theta}_{1,k-1})) - c,$$

and thus

$$\underline{x}_{1,k} = \underline{\theta}_{1,k-1} + \beta_1^{-1/2}\Phi^{-1}(1 - c/y).$$

Thus, with $\underline{x}_{1,0} = -\infty$, we have $\underline{\theta}_{1,0} = 0$, $\underline{x}_{1,1} = \beta_1^{-1/2}\Phi^{-1}(1 - c/y)$, and so on. Clearly, the sequence $\{\underline{x}_{1,k}\}_{k=0}^{\infty}$ is increasing and is bounded above by x_1^*. Hence, it converges to some $\underline{x}_1$. By continuity of $h(\cdot)$, the limit $\underline{x}_1$ must be a fixed point of h. But we have already proved that $h(\cdot)$ has a unique fixed point, and therefore $\underline{x}_1 = x_1^*$.

Next, construct a sequence $\{\overline{x}_{1,k}\}_{k=0}^{\infty}$ with $\overline{x}_{1,0} = -\infty$ and $\overline{x}_{1,k} = h(\overline{x}_{1,k-1})$ for all $k \geq 1$. Note that the sequence is decreasing and is bounded below by x_1^*. Hence, sequence $\{\overline{x}_{1,k}\}_{k=0}^{\infty}$ converges to some $\overline{x}_1$. By continuity of $h(\cdot)$, the limit $\overline{x}_1$ must be a fixed point of h. But we have already proved that $h(\cdot)$ has a unique fixed point, and therefore $\overline{x}_1 = x_1^*$.■

7 Appendix B: Tables and Figures

Table B-I.

Descriptive Statistics.

	Sessions 1-4 Cost 20	Sessions 1-4 Cost 50	Sessions 5-6 No New Info	Sessions 5-6 New Info
Mean θ	73.48	64.40	77.78	74.10
Min θ	-73.83	-92.73	-107.48	-49.06
Max θ	208.46	188.30	256.30	256.68
Mean x (Stage 1)	73.37	64.45	77.80	74.25
Min x (Stage 1)	-82.55	-106.92	-128.23	-71.84
Max x (Stage 1)	220.85	201.84	277.64	281.59
Mean # Attackers Stage 1	9.16	8.89	6.88	7.18
Mean # Attackers Stage 2	1.28	0.89	0.50	1.94
Mean Belief (Stage 1)	8.57	8.35	6.62	7.03
Mean Belief (Stage 2)	1.76	1.58	1.08	2.71
% Successful Attacks (Stage 1)	0.57	0.54	0.43	0.41
% Successful Attacks (Stage 2)	0.019	0.00	0.00	0.021
Median Comfort Level with Stats	4	4	4	4
Number of Subjects	120	120	60	60

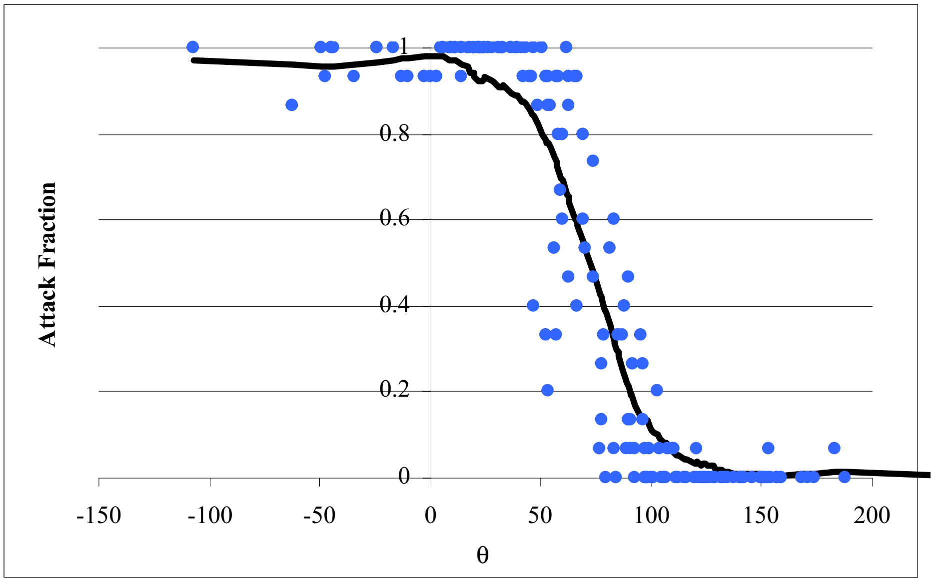

Figure B1: Kernel Regression (Cost-60 Treatment)

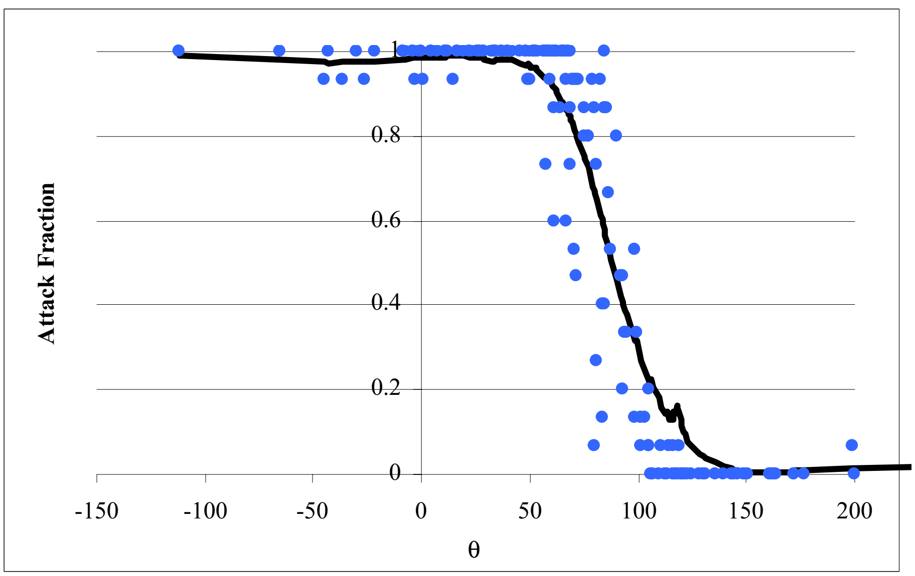

Figure B2.: Kernel Regression (Cost-20 Treatment)

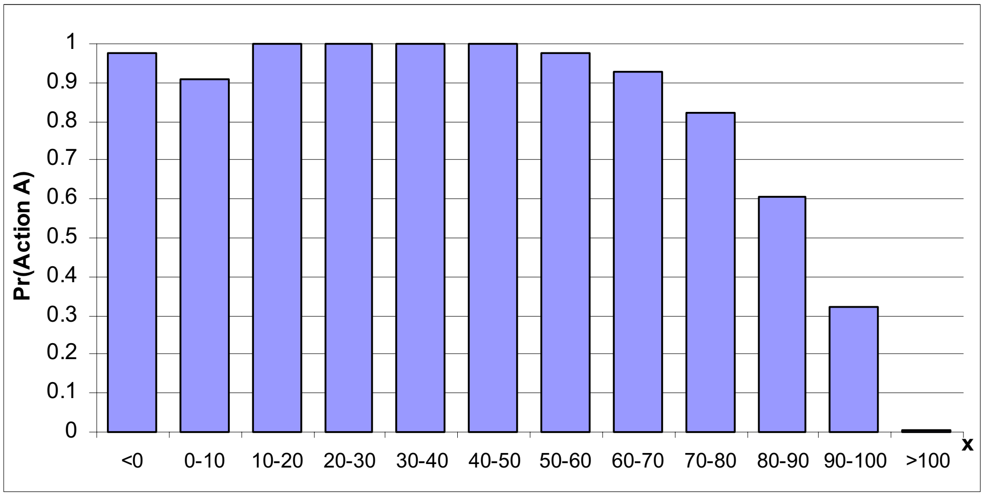

Figure B3: Probability of Action A vs. x (Cost-20 Treatment)

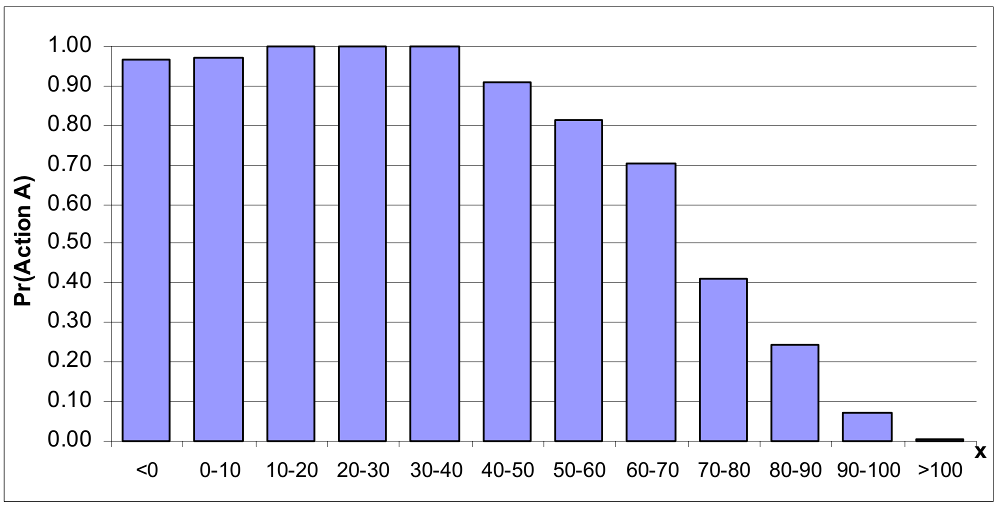

Figure B4: Probability of Action A vs. x (Cost-60 Treatment)

Table B-II.

Stage 1 Individual Level CLogit Regressions (Pooled Data for All Sessions)

	Dependent Variable: Action		
	1	2	3
Private signal, *x*	-0.1046***	-0.1064***	-0.0603***
	(0.0100)	(0.0104)	(0.0091)
Cost of action A		-0.0306***	-0.0115*
		(0.0046)	(0.0065)
Belief			0.6842***
			(0.0511)
Pseudo R^2	0.74	0.75	0.91
No. of observations	6000	6000	6000

Note: Standard errors in parentheses. Regressions include group and round fixed effects. For sessions 5 and 6, only the no-new-information treatment data are used. Significance levels: * 10%, ** 5%, *** 1%.

Table B-III.

Stage 1 Individual Level Regressions (Pooled Data for All Sessions; Subject Fixed Effects)

	Dependent Variable: Action		
	1	2	3
Private signal, *x*	-0.0065***	-0.0066***	-0.001***
	(0.0001)	(0.0001)	(0.0001)
Cost of action A		-0.0026***	-0.0004**
		(0.0003)	(0.0002)
Belief			0.0668***
			(0.0010)
R^2	0.58	0.59	0.84
No. of observations	6000	6000	6000

Note: Robust standard errors in parentheses. Regressions include group, subject, and round fixed effects. For sessions 5 and 6, only the no-new-information treatment data are used. Significance levels: ** 5%, *** 1%.

Table B-IV.

Individual Level CLogit Regressions (Pooled Data for All Sessions)

	Dependent Variable: Action		
	1	2	3
Private signal, x	-0.1052***	-0.1078***	-0.0614***
	(0.0090)	(0.0096)	(0.0097)
Cost of action A		-0.0339***	-0.0148**
		(0.0044)	(0.0064)
Belief			0.7079***
			(0.0485)
Stage	-1.6341***	-1.5964***	-0.8440***
	(0.1157)	(0.1145)	(0.1785)
R^2	0.75	0.76	0.92
No. of observations	8820	8820	8820

Note: Standard errors in parentheses. Regressions include group and round fixed effects. For sessions 5 and 6, only the no-new-information treatment data are used. Significance levels: ** 5%, *** 1%.

Table B-V.

Individual Level Regressions (Pooled Data for All Sessions; Subject Fixed Effects)

	Dependent Variable: Action		
	1	2	3
Private signal, x	-0.006***	-0.006***	-0.0008***
	(0.0001)	(0.0001)	(0.0001)
Cost of action A		-0.0026***	-0.0005**
		(0.0003)	(0.0002)
Belief			0.0655***
			(0.0008)
Stage	-0.2307***	-0.2272***	-0.0399***
	(0.0073)	(0.0073)	(0.0054)
R^2	0.62	0.62	0.85
No. of observations	8820	8820	8820

Note: Robust standard errors in parentheses. Regressions include group, subject, and round fixed effects. For sessions 5 and 6, only the no-new-information treatment data are used. Significance levels: ** 5%, *** 1%.

Table B-VI.

Stage 2 Individual Level CLogit Regressions Using (Sessions 5 and 6)

	Dependent Variable: Action
Private signal, *x*	-0.1428***
	(0.0199)
NI dummy	0.7783*
	(0.4518)
R^2	0.67
No. of observations	1395

Note: Standard errors in parentheses.
Regressions include group, subject, and round fixed effects. Significance level: * 10%, *** 1%.

Table B-VII.

Aggregate-Level OLS Regressions of Average Fraction of Mistakes on Rounds

	Dependent Variable: Average Fraction of Mistakes	
	Sessions 1-2	Sessions 3-4
Round*Cost	-0.0024	-0.0011
	(0.0019)	(0.0021)
Constant	0.1148**	-0.071**
	(0.0261)	(0.0044)
R^2	0.04	0.01
No. of observations	40	40

Note: Robust standard errors in parentheses.
Significance levels: ** 5%, *** 1%.