Speculative Attacks:
A Laboratory Study in Continuous Time

Yin-Wong Cheung
and
Daniel Friedman

Department of Economics,
University of California, Santa Cruz

August 14, 2006

Acknowledgement
The work was funded under NSF grant IIS-9986651. The authors are also grateful to programmers Nitai Farmer and Adam Freidin and to research assistants Lidan Wang, Paul Viotti, and especially Kai Pommerenke. We are greatly indebted to Joshua Aizenman and Federico Ravenna, who made substantial contributions to early versions of the paper. We also benefited from helpful comments offered by Jim Costain, Joseph Fung, Frank Heinemann, Hans Genberg, Antonio Guarino, Ali Hortascu, Steffen Huck, Ken Kletzer, Brian Kluger, Steve Morris, Peter Norman Sørensen, Jean-Robert Tyran, Eric Wincoop, Donald Wittman, Matthew Yiu, and seminar participants at University College London, The Hong Kong Institute for Monetary Research, and the 2006 CESifo Conference on Macro, Money & International Finance.

Speculative Attacks:
A Laboratory Study in Continuous Time

Abstract

We examine speculative attacks in a controlled laboratory environment featuring continuous time, size asymmetries, and varying amounts of public information. Attacks succeeded in 233 of 344 possible trading periods. When speculators have symmetric size and access to information: (a) weaker (or more rapidly deteriorating) fundamentals increase the likelihood of successful speculative attacks and hasten their onset, and (b) contrary to some theory, 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, and experience with large speculators increases small speculators' response to the public information. However, giving the large speculator increased size or better information does not significantly strengthen his impact.

Keywords: Currency Crisis, Speculative Attack, Laboratory Experiment, Coordination Game, Pre-emption, Large Player

JEL Codes: F30, C73, C92

Speculative Attacks:
A Laboratory Study in Continuous Time

1. Introduction

For more than 30 years, recurrent currency crises have confounded policy makers, and have challenged economic theorists to find explanations. The "First Generation" models, (e.g., Krugman, 1979, and Flood and Garber, 1984) gave insight into the Latin American and other crises of the 1970s and early 1980s, but had difficulty accounting for later events such as the 1992 British Pound crisis. Theorists responded with "Second Generation" models of self-fulfilling crises (e.g., Obstfeld 1995, 1996), leading to quite different policy implications. The 1997-8 crisis in East Asia and other recent events inspired "Third Generation" models that focus on financial market imperfections. Most Second and Third Generation models feature strategic interactions among speculators and government agencies, and many have multiple equilibria.

International finance economists now have an embarrassment of riches. With so many models to choose from, it is hard to trust the policy implications of any one of them; indeed, the next currency crisis might inspire Fourth Generation models with still different implications. Empirical work is essential to understand the range of applicability of current models. Traditional econometric methods help, but are limited by a mismatch between data and theory. Key variables, such as speculators' information and expectations, are not observable in the field so inferences are limited. The historical data, of course, can exhibit at most one equilibrium at a time, so the existence of multiple equilibria remains a matter of conjecture.

The laboratory offers an underexploited source of data for examining the robustness of theoretical predictions. In the lab, one can observe and even control key variables including the information available to players, can replicate a given scenario many times, and can make rigorous causal inferences. Of course, laboratory markets are much smaller scale than in the field, and one must investigate external validity. Due to their very different strengths, laboratory and field data are complements, not substitutes.

The present paper describes one of the first laboratory investigations of speculative attacks. The laboratory game is not closely tailored to any particular theoretical model. It is intended to provide evidence not otherwise obtainable that can help sharpen the issues, inform econometric analysis of field data, and refine existing models.

Unlike most current models and experiments, time is continuous in our laboratory game. During each trial, the strength of fundamentals can deteriorate (or improve or remain constant) moment by moment, and the speculators can switch back and forth between passive mode and attack mode. Such asynchronous decisions by human subjects illuminate aspects of the coordination problem that are invisible in discrete time models, but that may be crucial in the decentralized 24 hour global foreign exchange market.

The experiment focuses on a question central to many of the Second and Third Generation models: when can speculators coordinate on their preferred equilibrium, an attack that forces devaluation? The question is particularly acute when the speculators are uncertain about the strength of fundamentals, or are uncertain about what the other speculators are doing. We examine whether public information about the fundamentals and about other players' choices affects coordination and, hence, the probability of a successful attack.

In some trials we include one speculator able to take a larger position and sometimes with better information on the fundamentals. It is widely believed that a single large player, George Soros, was essential in coordinating the successful 1992 attack on British pound, and a branch of the theoretical literature justifies that belief. The experiment also examines some predictions from that theoretical literature.

Section 2 briefly surveys literature on some relevant theoretical models, some field evidence, and the handful of relevant laboratory studies. Section 3 offers some simple theoretical perspectives and summarizes existing theoretical predictions that are testable in the laboratory. Section 4 describes our laboratory environment and treatments.

Section 5 presents the results. When feasible, attacks succeed more often than not. With speculators of symmetric size and access to information, speculative attacks are more often successful and occur sooner when fundamentals are weaker. Contrary to some theory, 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, and experience with large speculators increases small speculators' response to the public information. However, giving the large speculator increased size or better information does not significantly strengthen his impact.

Section 6 puts the findings into broader perspective. Appendix A is a copy of instructions to laboratory participants.

2. Some Related Literature

As just noted, there are already several generations of models designed to explain foreign currency crises. Henderson and Salant (1978) is an important precursor, dealing with the interaction between the speculators and government policies. The First Generation models, formulated by Krugman (1979) and Flood and Garber (1984), focused on economies where the fundamentals are inconsistent with preserving a fixed exchange rate. This is the case when a sizable fiscal deficit is monetized and the money supply increases faster than money demand under a fixed exchange rate, inducing a crisis. In the absence of uncertainty, the crisis hits and the currency devalues when the shadow exchange rate [i.e., the exchange rate that would prevail if the exchange rate were allowed to float] is equal to the official exchange rate.

The First Generation models provide a useful interpretation of crises triggered by fiscal weaknesses, as in Latin America in the 1970s-80s. However, the financial instability in Europe in the early 1990s, including the 1992 attack on the British Pound, shifted attention to the fragility of equilibrium under a fixed exchange rate. The role of financier George Soros in the 1992 attack also raised questions regarding speculators' behavior when some of them have market power. "Second Generation" models, formulated by Obstfeld (1995, 1996), feature self-fulfilling crises and multiple equilibria. Obstfeld (1995), for instance, extends the Barro-Gordon framework and models a policy maker who faces a trade-off between credibility and flexibility when devaluation entails a fixed cost. The potential for self-fulfilling crises stems from circularity: the behavior of the policy maker depends on prior private sector expectations of inflation and depreciation. These expectations, however, depend on market perceptions of the policy maker's behavior. In such circumstances, an increase in the cost of abandoning the peg may increase the likelihood of a crisis [Flood and Marion (1999)], and toughness can be counterproductive [Drazen and Masson (1994)].

Later events, especially the 1997-8 crisis in East Asia, led to "Third Generation" models. Their focus is the role of balance sheet factors and financial sector weaknesses, as well as the possibility of bailouts by international financial institutions, central banks, and governments. Leading examples include Krugman (1999), Corsetti, Pesenti and Roubini (1999), Chang and Velasco (1999) and Dooley (1997).

The growing dynasty of models attests to the rich interplay between economic theory and economic events. The downside is that models seem to "chase" the last crisis in an attempt to

provide ex-post explanations. Of course, the research has also identified some ex-ante vulnerabilities and suggested ways to prevent or mitigate crises; see e.g. Obstfeld (1998).

A common thread of the Second and Third Generation models is the role of public information and the existence of multiple equilibria. Morris and Shin (1998) resolve multiplicity by introducing private information into the global games approach of Carlsson and Van Damme (1993). When privately informed speculators make simultaneous choices, one obtains a unique equilibrium for almost every value of the fundamentals. Much of the subsequent theoretical literature investigates when private information does or does not eliminate multiple equilibria (e.g., Angeletos and Werning, 2004).

Several later authors, surveyed in Allegret and Cornand (2005), extend the static Morris and Shin model to examine policy issues. There has been a widespread consensus in the policy community, bolstered by theoretical models such as Heinemann and Illing (2002) and by some evidence such as Bannier (2004), that greater transparency (i.e., better information on fundamentals and perhaps on speculative behavior) will discourage speculative attacks. Using a slightly different model, Bannier and Heinemann (2005) confirm the consensus except in some circumstances when prior beliefs about fundamentals are optimistic, in which case a lesser degree of transparency will minimize speculative attacks. The survey finds other exceptions and concludes that central bank transparency has an ambiguous effect theoretically (and empirically).

Noting that "the coordination problem among investors is at the heart of most second- and third-generation models of currency crises," Corsetti *et al.* (2004) work out the theoretical implications of introducing a single large player into the global games approach. Their model predicts that the mere presence of a large player makes other (small) players more willing to join a speculative attack. The effect is enhanced when the small players can observe the large player's position, but is not necessarily enhanced by increasing the size of the large player. In a more strategically symmetric model, Bannier (2005) concludes that attacks are more likely with a larger and less well informed large player when beliefs about fundamentals are pessimistic, but the reverse is true when beliefs are optimistic.

Guimaraes (2004) extends the uniqueness result to a continuous time setting in which fundamentals follow Brownian motion and a continuum of speculators chooses at Poisson times. Costain (2004) obtains bimodal outcomes (both successful and unsuccessful attacks) over a wide range of fundamentals when speculators choose sequentially and observe some earlier choices.

Angeletos et al (2006) reach a similar conclusion in a discrete time dynamic extension of global games where players can learn from earlier periods. For our purposes, their most interesting conclusion is that "equilibrium dynamics can alternate between phases of distress and phases of tranquility, even without changes in the fundamentals."

Evidence on the role of large players in contemporary foreign exchange markets includes the Federal Reserve Bank of New York (1998) report that the top five trading firms account for 31% of market share. Recent surveys of dealers (Cheung and Wong, 2000; Cheung and Chinn, 2001; Cheung, Chinn and Marsh, 2004) indicate that trading in some currencies is dominated by a few big players. Respondents attribute deviation of FX rates from fundamental values mainly to "excess speculation" and "institutional customer or hedge fund manipulation." Many observers agree that George Soros' hedge fund, followed by other investors, precipitated devaluation of the British pound and Italian lira in the 1992 EMS crisis. Some argue that the hedge funds led the way in the rapid devaluation of the Japanese yen in late 1998. It would seem that their high degree of leverage allows hedge funds to influence the price of thinly traded currencies (*Economist*, October 10, 1998; Sesit and Pacelle, 1998). On the other hand, Eichengreen *et al.* (1998) and others argue that hedge funds did not exacerbate the 1997-8 Asian currency crises.

We are aware of only a few relevant laboratory experiments. The first, reported in Heinemann, Nagel and Ockenfels (2004, henceforth HNO) examines the static symmetric model of Morris and Shin (1998). Subjects play the role of speculators in a series of simultaneous move coordination games, with public or private information about the strength of fundamentals. Contrary to the key global games prediction that private information resolves multiplicity, the speculators coordinate more successfully and achieve higher payoffs with public information. Consistent with the basic comparative static predictions, speculative attacks succeed more often when fundamentals are weaker and when speculation costs are lower.

A second study by Costain, Heinemann and Ockenfels (2005, henceforth CHO) reports bimodal outcomes when traders move sequentially, as predicted in the Costain (2004) model. Cornand (2005) extends the static HNO setup and finds that speculative attacks succeed more often with two public signals than with one public and one private signal. Arikawa et al. (2006) also extend HNO and find, consistent with Corsetti *et al.* (2004), that small players are more likely to attack when a large player is present.

Do similar results hold in continuous time? Huberman and Glance (1993) argue that continuous time can make coordination more difficult because choices are asynchronous—it's not only what you do but when you do it. In the only other continuous coordination laboratory game we know, Brunnermeier and Morgan (2004, henceforth BM) [1] 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.

3. Theoretical Considerations

This section spotlights the main issues by describing speculators' activities and writing down a simple stylized model.

When speculators think that monetary authorities may not be able to maintain a fixed exchange rate for some specific currency (against, say, the US Dollar), they can "attack" by selling the currency short, hoping to profit from the difference between the exchange rate before and after devaluation. Devaluation occurs when in aggregate speculators' net short position exceeds some critical level that represents the monetary authorities' willingness and ability to defend the currency (the so-called "fundamentals"). While waiting for devaluation, the speculators incur ongoing costs. The main explicit cost is the interest rate differential: authorities often defend a currency by raising the interest rate hundreds or even thousands of basis points above the US dollar rate. The main implicit cost is the risk that the authorities might impose capital controls, or might otherwise penalize speculators taking large short positions.

Many of the strategic issues are captured in the following continuous time coordination game. At each moment t in [0, 1], each player i with mass (or position size) m_i can switch between passive mode, denoted $a_i(t) = 0$, and attack mode $a_i(t) = 1$.[2] Let $B(t) = \sum_i a_i(t)\, m_i$ be the total mass of players in attack mode at time t, and let $T(t)$ represent the threshold mass, i.e., the exogenous strength of fundamentals. Normalize payoff flows so that players earn 100 per unit time in passive mode (an opportunity cost that represents the ongoing costs noted above) and,

[1] We believe that all the laboratory studies mentioned here were independently initiated at about the same time as (or after) our own except for HNO, which came first.

[2] A careful formal model would impose the assumption that the a_i functions are right-continuous, or at least measurable.

when $B(t) < T(t)$, earn 0 in attack mode. Let $t^* = \min\{1, \inf\{t: B(t) \geq T(t)\}\}$ represent the moment that the speculative attack succeeds, or the end of the period if the attack never succeeds. The attackers receive a lump-sum payoff of $L>0$ when the attack succeeds. After that point, all players earn the opportunity cost. Set the discount factor to 1 for simplicity.

Thus, using the Heaviside function $\theta(x) = 1$ if $x > 0$ and $\theta(x) = 0$ otherwise, the payoff for player i can be written as $100(1-\int_0^{t^*} a_i(t)dt) + \theta(1-t^*)a_i(t^*)L$. In the simplest variant of the game, players at time t observe the history $T(s)$ and $B(s)$ for all $s<t$, and all have equal mass normalized to unity, $m_i = 1$. More complicated variants allow for noisy observations of $T(s)$ and $B(s)$ and for size asymmetries.

Some caveats are in order before discussing the equilibria of these games and the testable implications. The threshold $T(t)$ is exogenous, i.e., the monetary authority is not a strategic participant, in order to keep a sharp focus on coordination and information issues. The games suppress two margins that sometimes play a minor role in the literature: the size of the devaluation is fixed at L, and each player either attacks fully or not at all.

At any fixed time t, the (static) equilibrium correspondence depends on the strength of the fundamentals $T(t)$ relative to the mass $m = \max_i m_i$ of the largest individual player, and to the total mass $M = \sum_i m_i$ of all players. When $T(t) \leq m$, there are only devaluation equilibria in which a sufficiently large player unilaterally forces devaluation. When $T(t) > M$, there is only the trivial equilibrium with everyone passive. Our focus is the intermediate case. When $m < T(t) \leq M$, there are devaluation equilibria as well as the trivial equilibrium. The question then is, when will the players be able to coordinate an attack to achieve the higher payoffs in a devaluation equilibrium?

Continuous time introduces some interesting complications. Player i would find it in her immediate interest to join a speculative attack whenever she is pivotal, i.e., at any time t such that $B(t) < T(t) \leq B^{+i}(t) \equiv m_i + \sum_{j \neq i} a_j(t)\, m_j$. Typically several passive players would find themselves as pivotal at the same time, so they play a preemption game: as with patrons of a burning nightclub or depositors in a bank run, the reward goes to those who move first. The incentives are in the opposite direction when a player is not pivotal. It is in her interest to wait for the other players to attack (and let them forego the opportunity cost of 100 per unit time) until she becomes pivotal. Thus the direct incentive switches abruptly from free riding (waiting) to

preemption (rushing) when a player becomes pivotal.[3]

The efficient equilibrium in the continuous time game is for all players is to switch simultaneously into attack mode at the first time t^* that $T(t^*) \leq M$. Of course, there are many other equilibria, e.g., "sunspot" equilibria in which a sufficient mass of players attacks at an arbitrary t^* such that $T(t^*) \leq M$.

In practice, one wonders how the players might learn to coordinate on any equilibrium. Some of our evidence suggests a home-grown signaling process. When $B(s)$ is publicly observable, a player switching into attack mode signals to the other players currently in passive mode that they are now closer to being pivotal, which may encourage them also to switch. This signal is low cost if the switch is brief. Such signaling episodes would be observed as pulses in which B increases rapidly and either reaches the threshold T so the attack succeeds, or else tops out below T and then rapidly declines.

The interesting questions now may be posed as follows.

1. Fundamentals. Is the probability of a successful attack a decreasing function of $T(t)$? The static Nash equilibrium correspondence allows for that possibility but does not require it. Recall that the global games approach predicts a sharp transition from probability 0 to probability 1. To expand slightly on the comments at the end of the previous section, HNO obtain a more gradual transition using different fixed levels of T in different trials, while CHO confirm an almost linear transition with simultaneous choice but persistent bimodality with sequential choice. By contrast, our laboratory design will allow us to evaluate the effects of continuously changing fundamentals.

2. Information. Does the probability of a successful attack increase or decrease when players have better public information about B and T? The discussion above (regarding signaling and being pivotal) suggests an increase. The policy consensus, together with some theory and field evidence, predicts a decrease. Bannier and Heinemann (2005) predict that it can go either way depending on market sentiment and other factors.

3. Asymmetries. Does the presence of a player with larger mass or better information increase the

[3] Similar timing issues arise in several other contexts besides speculative attacks. These range from trying to be fashionably late to a dinner party, to introducing a new high technology product or harvesting a seasonal crop, to selling off a bubble asset. Park and Smith (2003) consider such games with no information on other players' choices and no option to switch back to non-speculative mode. They derive symmetric mixed Nash equilibria characterized by a trickle of attackers until a crucial moment (determined by the known future path $T(t+s)$) when all remaining players attack.

probability of a successful attack? Corsetti *et al.* (2004) predict affirmative answers that are consistent with the laboratory results of Arikawa et al. (2006). Bannier (2005) has more nuanced predictions. Our design will provide evidence in a continuous time setting.

4. Coordination. In terms of individual choice, do players respond differently when they are nearly pivotal than when they are far from pivotal? More precisely, is a player in passive (or active) mode more (or less) likely to switch when the shortfall $T(t)$ - $B^{+i}(t)$ is smaller?

4. The Experiment

To answer such questions, we conducted an experiment at UCSC's LEEPS lab using a custom-built software platform called EvTech that supports asynchronous (i.e., continuous time) binary choice games. A typical session lasted 90 minutes and involved 8 inexperienced human subjects recruited by email from a campus-wide pool of undergraduate volunteers. Subjects silently read the instructions, reproduced in Appendix A, and then listened to an oral summary by the conductor. After a quiz and a couple of practice rounds, they played about 30 periods. Each subject received in cash the profits earned in every period, typically between $15 and $25.

During each period, each subject privately viewed a player screen display similar to the one shown in Figure 1. Except for one condition explained below, each subject knew that other subjects saw the same information on their screens. Each subject had two buttons, labeled Red (corresponding to the passive mode) and Blue (the attack mode), and switched mode by clicking the button not grayed out. The per-period payoff rates were exactly as in the theoretical discussion above: 0 for Blue prior to success, and 100 for Red (the opportunity cost) and for either mode after a successful attack. Players in Blue mode at the time that an attack succeeded received a known bonus, typically L=100 points. At the end of each period, the subject could see her payoff performance and that of all players, as in the right side panels in Figure 1.

Continuously updated panels as on the left in Figure 1 allow the player to observe the mass $B(s)$ of players in attack mode, the threshold $T(s)$ for success, the payoffs actually earned and payoffs foregone, for all times s from the beginning of the period up to the current time t. More precisely, the program refreshes the displays every "tick", set at half a second (500 msec).

Table 1 below lists 12 experimental sessions with a total of 344 relevant periods of 90 seconds each. A typical session consisted of several blocks of four periods. Each block held constant all treatments except the time path of fundamentals; the other treatments varied across

blocks and across sessions. Key treatments include the following.

- Number of players and experience. The number of human subjects varied across sessions from 4 to 11, and was known to everyone. The last two sessions used only subjects with experience in a prior session of the experiment; the other sessions used only inexperienced subjects.
- Fundamentals. The threshold $T(t)$ always started out at 110% of total player mass M at t=0, fell linearly to a predetermined final value at a predetermined time, and remained at that final value for the rest of the period. Final values and times ranged from 50% of M at 15 seconds to 105% at 60 seconds. Of course, attacks can not succeed in the 16 periods (4.4% of the 360 potentially relevant periods) for which $T(t)$ remains above 100%, and those periods are excluded from the analysis.[4] The vast majority of fundamentals paths lie between a decrease to 90% of M within 45 seconds, and a decrease to 60% within 30 seconds. The path changed each period (to discourage uninteresting focal point equilibria), and varied across blocks orthogonally to the other treatment variables (to sharpen statistical inferences).
- Information. The threshold line $T(t)$ and blue activity line $B(t)$ were displayed sharply in baseline periods, as in Figure 1. In some other periods, participants saw a thick line or band encompassing (not centered on) the actual value; its width allowed observational errors of up to two players. In some periods we completely suppressed either the T line or the B line. Thus each source of information had three possible levels, yielding 3*3=9 different information treatments. Preliminary analysis showed that performance in the thick line treatments was not significantly different from the sharp line treatments, so the main data analysis involves only 2*2 = 4 distinct information treatments.
- Symmetry. In a bit over half of the reported periods, all players had the same screen displays and the same weight. In three periods of most blocks of some sessions, one player (called Soros below) had weight 3 or 5 times the weight of the other players. Soros observed the same B line as everyone else (either perfect or none), but in some blocks when other players did not observe T, Soros observed it perfectly, and this possibility was known to all players. The Soros role rotated among the players across blocks. Each player knew when Soros was

[4] Each block of four periods was separated by a single period using a different binary choice game, such as a congestion game. Of course, we also exclude these separation periods from the present paper, as well as periods from early pilot sessions (in which we often varied the number of active players within a session and introduced some automated agents). No other periods are excluded.

present, could easily infer Soros' size from periods when the *B* line appeared, and knew when she had the Soros role.

Our laboratory environment is constructed to address the questions posed at the end of the previous section, regarding the impact of fundamentals, information, player asymmetries and pivotality. It does not address some other questions covered in the HMO experiment, such as the comparative statics of devaluation size and opportunity cost, or the general impact of private information (we examine only Soros' private information). The sequential environment of CHO and the continuous time environment of BM allow players only to switch once from passive to attack mode; ours allows them to switch freely back and forth. Thus our asynchronous environment provides a new and stringent test of coordination.

5. Results

Figure 2 summarizes behavior in period 19 from the last session of the experiment. In that block (periods 16-20), all players observed neither the blue line $B(t)$, the weight of players in attack mode, nor the gray line $T(t)$, the threshold weight or strength of fundamentals. Player #4 had the Soros role, with weight 5. In period 19, the threshold dropped from 110% of total mass M (= 5+8*1 = 13) to 50% of M at t=45 seconds. In the first 20 seconds or so players remained in passive mode except for brief solo attacks by player #5 and player #4 (Soros). Then several small players successively switched into attack mode and remained there for several seconds, but the attack reached a plateau and began to subside by t=27 sec.[5] At t=41, with much weaker fundamentals, the attack was renewed and it succeeded at t^*=45.5 when Soros finally switched to attack mode. Soros and the other attackers at t^* earned about 190 points each in period 19, while passive players earned 100 points. Player #9 abandoned the attack just before it succeeded and earned only 89 points.

The other 343 periods differed in many respects. Some sessions saw many periods with several unsuccessful attack pulses, while other sessions (including the last) usually saw attacks succeed on the first pulse. Some sessions had relatively few successful attacks, but in one session attacks usually succeeded almost as soon as they became feasible. Summaries for all periods in Figure 2 format cannot be included here but are available on request.

[5] Such a pulse may be reminiscent of Angeletos et al. (2006).

Figure 3 provides an overview of speculative success for the 296 symmetric periods, by information condition. The success rate for speculative attacks is, of course, zero when the fundamentals are too strong ($T > M$). A simultaneous attack by all speculators is required when $T = 0.95M$, and the figure shows that the success rate in this case is under 10% in all information treatments. When speculators can see both $B(t)$ and $T(t)$, the success rate climbs steeply as fundamentals weaken, e.g., it is over 70% when $T = 0.75M$, i.e., ¾ of the speculators are required. By contrast, the success rate at the same threshold is only about 20% when speculators see neither $B(t)$ nor $T(t)$. Players achieve intermediate success rates when they can see one but not both of the lines.

The figure suggests affirmative answers to the first two research questions: the probability of a successful attack does seem to be a decreasing function of the threshold, and seems enhanced by more public information. But the figure doesn't indicate whether the differences are significant, it conflates rapidly moving with slowly moving fundamentals, and it omits the asymmetric periods.[6] More quantitative tests are in order.

Our quantitative data analysis first treats each period as a single observation, and examines four measures of performance.

1. Success: this dummy variable =1 iff there is a successful attack that period.
2. Time to success, $\tau = t^* - t_o$: seconds elapsed from the time t_o when the attack first became feasible (so $T(t_o) = M$) until the time t^* when the attack succeeded. Since time-dependent variables are recorded by tick (half-seconds), we can also measure time as the integer number of ticks instead of the decimal number τ of seconds. The τ and tick variables are right-censored by the end of the period when $t^*=90$, i.e., when no attack succeeds that period.
3. Minimum gap, Mingap = $\min\{T(t) - B(t): 0 \leq t \leq 90\}/M$, the minimum shortfall of the mass attackers from the threshold within the period, expressed as a fraction of player mass M. Mingap measures continuously how close attacks come to success. It equals zero and is left-censored when Success = 1.

[6] After making conventions on how to classify periods in which Soros sees T but other players do not, etc., one can create versions of Figure 3 that includes the asymmetric periods. The results are roughly similar to Figure 3 except that the DB=1, DT=0 line is closer to the full information line DT=DB=1, suggesting that information on speculative activity is more useful when Soros is present.

4. Bluearea = $\int_0^{t^*} B(t)dt$: the area under the blue activity line $B(t)$ before the attack succeeds. It measures speculators' foregone earnings due to attacks before achieving success. It is right-censored when Success = 0.

We regress these performance measures on a set of explanatory variables defined in Table 1. To reiterate:

1. Threshold information dummy: DT=1 if a threshold line is displayed (exact or with noise) and =0 if it is not displayed. (We found no qualitative difference between an exact DT dummy and a noisy DT dummy in the symmetric sessions, so we combine them.)
2. Blue information dummy: DB=1 if the blue line is displayed (exact or with noise) and =0 if it is not displayed. (Again, for simplicity we combined the exact and noisy DB dummies after finding no qualitative difference between them.)
3. Steep dummy: Steep=1 if $T(t)$ fell to 0.5 within 50 seconds, or to 0.6 within 30 seconds. A value of 1 indicates rapid deterioration of the fundamentals.

5.1. Symmetric Periods

We begin the regression analysis with the 193 periods with symmetric players. To explain the first performance measure, success, we estimate the Logit model

$$P(y_i = 1) = F(x_i'\beta),\ i = 1, \ldots, N,$$

where y_i is the period i success dummy taking the value of 1 or 0, while x_i is a vector of explanatory variables introduced earlier, and β is the corresponding vector of coefficients. $F(w) = e^w/(1+e^w)$ is the logistic distribution function, and N=193 is the number of periods included in the regression analysis.

Panel A of Table 2 reports the coefficient estimates. The probability of a successful attack increases when the threshold and attack mass variables are observable and when the fundamentals deteriorate rapidly. Indeed, a steep decline in the fundamentals increases the log odds of success by almost three. We also considered the effect of the interaction terms DT*Steep and DT*DB, but omit reporting them because neither is highly significant in any specification or even marginally significant in most specifications.

Due to their censoring structure, the performance variables τ, Mingap, and Bluearea are

estimated using the Tobit specification

$$y_i^* = x_i'\beta + \varepsilon_i,\ i = 1, \ldots, N$$

where the observed variable y_i is equal to its latent counterpart y_i^* if there is no censoring and is equal to 1 (0) if y_i^* is right- (left-) censored. (For τ, the censor-value of 1 period length corresponds to 90 sec or 180 ticks.) The coefficients are estimated under the standard Type I extreme-value distribution assumption.[7]

The estimation results are summarized in columns 3, 4, and 5 in Table 2A. The Tobit results reinforce the implications of the Logit estimates. For instance, the time to success (τ) comes significantly sooner when players see the T and the B lines and when T drops steeply. Each dummy variable reduces the time to success by about 15% (e.g., 0.88 /5.87 ≈ 0.15 for DB). Again, the estimates indicate that the 'Steep' dummy has an even larger impact.

So far the regression results indicate that information about the fundamentals and speculative activity hastens the occurrence of a successful attack. What happens when no successful attack is observed? The Mingap Tobit regression again indicates that all three explanatory variables foster speculation and help narrow the gap to success. The Bluearea Tobit regression likewise shows that the same variables reduce the time spent in fruitless speculation.

To check robustness and to offer a different perspective on the data, we re-examine the time to success τ using the Cox proportional hazard model (PHM)

$$h_i(t) = h(t,x_i) = h_0(t)\exp(x_i'\beta),\ \ i = 1, \ldots, N.$$

The variable $h_i(t)$ is the hazard rate (or, in the present context, the success rate), the probability density that a speculative attack is successful at time t conditional on the explanatory variables in the vector x_i (and conditional on no earlier success, so we can use t and τ interchangeably). The baseline hazard rate $h_0(t)$ is $h_i(t)$ when $x_i = 0$. One advantage of PHM is that it is semiparametric and imposes no functional form on the baseline function $h_0(t)$.

The results, given in the sixth column of Table 2, strengthen the earlier conclusions. Other things equal, displaying the strength of fundamentals (DT=1) increases the probability of a successful attack per unit time by an estimated factor of $e^{1.08} \approx 2.94$, or 194%. Likewise, displaying the overall speculative position (DB=1) more than triples the probability of success ($e^{1.29} \approx 3.65$) and steeply declining fundamentals (Steep=1) increases it sixfold ($e^{1.85} \approx 6.36$). The

[7] The Type I extreme-value distribution function, also known as the log Weibull distribution function, is given by $F(w) = \exp[-\exp(-w)]$.

coefficient signs are reversed relative to the Tobit model because the PHM estimates the success rate instead of the time to success.

The PMH also allows us to examine directly the effect of the time-varying threshold, as reported in the last column of Table 2. Of course, the Steep dummy is now redundant and is omitted and, due to the way the data are recorded, the time variable is expressed in ticks rather than in seconds. Both information dummies DT and DB maintain their significance and approximate magnitudes in the presence of $T(t)$, although their relative size changes. The coefficient estimate -7.04 implies that deterioration of fundamentals causes a dramatic increase in the probability of a successful attack. For example, the success probability doubles when T decreases by just 10% of speculators' total capacity M (since $e^{(-0.10)*(-7.04)} \approx 2.02$).

Players in a given session might be especially good or bad at coordinating, and this could affect the results. Therefore in Panel B we report estimates using session fixed effects. None of the estimated effects changes sign or loses significance, and several estimated effects become stronger. A few of the session fixed effects are notable, e.g., session exp101 has a very large (but statistically insignificant) fixed effect for Success, suggesting excellent coordination.

5.2. Soros data

Recall from Table 1 that the Soros player has mass three or five times the mass of the other players and sometimes also has an informational advantage. Additional explanatory variables include the following dummies:

Soros	equals one if there was a player with weight>1
BigSoros	equals one if his weight=5; =0 otherwise
PT	equals one if the threshold line {T(s): s<t} is displayed for the Public
OST	equals one if only Soros sees the threshold line, but not the Public

Table 3 reports regressions for the asymmetric sessions, which include 46 symmetric periods as well as 151 periods with Soros. The specifications are similar to those in Table 2, augmented by the 'Soros' variable and the more significant interactions. The variable 'BigSoros' is insignificant in all specifications and so is dropped from the reported results. One possible interpretation, consistent with the theoretical results of Corsetti *et al.* (2004), is that what matters is Soros' mere presence, not his exact trading capacity. The 'OST' variable is also insignificant

in all specifications and omitted from the table.

The results in Table 3 confirm that Soros has a real impact. The "Soros" dummy variable is significant at least at the 10% level in each regression, and always indicates that speculative attacks are more likely to succeed and come sooner when Soros is present. For example, the third column of Panel A indicates that Soros speeds the time to success by almost 30% (from 1.94/6.76 ≈ 0.287). Even when the attack is not successful, the Mingap and pressure estimates in columns 4 and 5 indicate that the presence of Soros significantly narrows the gap from success and adds to speculative pressure.

The impact of fundamentals $T(t)$ is less in Table 3 than in Table 2 but is still quite strong. For example, the devaluation probability increase for a 10% deterioration in fundamentals is 65% (based on the -5.01 coefficient estimate in the last column of Table 3), about two-thirds as strong as in Table 2. Steeply declining fundamentals seldom appear in the asymmetric data, so the steep dummy has more erratic coefficients and usually lower significance than in Table 2.

Table 3, like Table 2, reports that both informational dummy variables significantly increase the probability of an attack, hasten its onset, increase speculative pressure, and reduce the shortfall of unsuccessful attacks. However, in Table 3 most of the direct effects are stronger and they often interfere with each other. The Tobit time to success estimates, for example, indicate separate reductions of about 18% and 36% respectively for PT (the counterpart of DT from Table 2) and DB, but an offset of 12% in the combination. Thus when speculators observe both speculative activity and fundamentals, the average time to success declines by about [1.24 + 2.42 – 0.82]/6.76 ≈ 18% + 36% - 12% = 42%.

The interaction term DB*Soros is insignificant in many of the regressions, but comes up marginally significant in two of the three timing regressions, with a sign again indicating an offset of the level effects. Looking once more at the Tobit time to success estimates, we note that the separate effects are 36% and 29% reductions respectively for observing speculative activity and for Soros' presence, but the coefficient estimate for DB*Soros implies an offset of 1.55/6.76 ≈ 23% and thus a combined effect of about 36% + 29% - 23% = 42%. In comparing these estimates to those from Table 2, the main puzzle is the large impact for DB (estimated at 36%) when Soros is absent. Perhaps players somehow become more attuned to signaling in sessions where Soros is usually present.

The estimates reported in Panel B show once again that allowing for fixed effects

generally strengthens the inferences.

5.3 Individual Player Decisions

The last column of Table 3 investigates a finer-grained question: how do individual speculators decide when to attack? We fit the proportional hazard model to the 2448 ticks at which a speculator first switched from passive mode into attack mode. The estimates indicate that such switches are much more likely when speculators can see the strength of fundamentals and somewhat more likely when they can see the amount of speculation. Stronger fundamentals (when observable) deter attacks and the presence of Soros encourages attacks.

Table 4 investigates whether Soros exploits mainly his greater opportunity to be pivotal and so tends to attack last, or exploits mainly his signaling strength and so tends to attack early. The most direct evidence comes from defining the variable Shortfall from Pivotality, $SP(i,t) = T(t) - B(t) + (a_i(t)-1)m_i$, for each player i at each time t. Thus SP is the number of additional regular (mass=1) players it would take for the attack to succeed. Table 4 reports mean values of SP of 5.25 for regular players and 3.70 for Soros players when Soros has no informational advantage, and reports similar averages (if anything, a bit further apart) when Soros knows the strength of fundamentals but small speculators do not. Thus, contrary to a prediction of the Corsetti *et al.* (2004) model, our data indicate that Soros attacks later than ordinary speculators.

Table 5 looks at each player's choices every half-second tick prior to a successful attack in all symmetric periods. For players in passive mode (about 109 thousand observations), it uses a Logit model to fit the cases (about 3 thousand) where the player switched to attack mode on the next tick. For players already in attack mode (about 27 thousand observations) it fits the choice of staying in attack mode (about 25 thousand cases) or switching back to passive mode. The second column reports direct effect estimates. The presence of either line increases the switching probability significantly, and a greater shortfall from pivotality decreases it. The more refined quantitative estimate in the fourth column allows for individual subject fixed effects and interactions. The specification assumes that the log odds of switching on the next tick are linear and, when both lines are visible, the estimates give slope -.34+.09+.13-.29 = -0.41 and intercept -2.80-.20-.18+1.18 = -2.00. The corresponding switch probability is exp (-2.00 - 0.41SP) /(1+ exp(-2.00 - 0.41SP)), about 12% when the player is pivotal (SP=0, essentially a measure of reaction time) and about 8% when only one more speculator is required for a successful attack

(SP=1). The estimated probability falls below 1% when at least 7 more speculators are needed.

These coefficient estimates indicate that seeing both lines together has a much stronger effect than seeing either of them by itself: for the intercept, the direct effects (-.20, -.18) are at best marginally significant while the interaction (+1.18) is very significant, and for the slope estimates the interaction term (-0.29) again is the most significant and strongest term.

The last two columns of Table 5 obtain parallel results for players already in attack mode. The direct effect estimates in the third column indicate that players are more likely to remain in attack mode when they have more information and when the attack is closer to success. The more refined specification in the last column obtains stronger direct effect estimates but relatively weak interaction estimates.

Subsets of the data provide evidence on the robustness of two conditional predictions regarding transparency. Recall that Bannier and Heinemann (2005) conclude that greater transparency may encourage speculation in some circumstances when speculators believe that fundamentals are strong. Their model, an extension of global games, doesn't translate precisely into our laboratory setup, but if their conclusion is very robust it would suggest that the positive DB and DT coefficient estimates in the second column of Table 5 are due to observations when $T(t)$ is high. Figure 3 points to $T(t)$=0.8 as a natural breakpoint, since the success rate of speculative attacks in the full information condition first exceeds 50% and climbs sharply at this threshold. Restricting to $T(t) \geq 0.8$ data, we obtain coefficient estimates -0.06 and 0.23**, smaller than the unrestricted estimates of 0.15** and 0.24** reported in the table. Thus the first conditional prediction is not supported in our data.

Recall that Bannier (2005) concludes that attacks are more likely with a larger and less well informed large player when beliefs about fundamentals are pessimistic, but the reverse holds when beliefs are optimistic. A rough translation into our data predicts higher attack rates by ordinary players in condition 5N-p (Soros has weight 5 and neither Soros nor regular players see the threshold line, and fundamentals are "pessimistic," $T(t) < 0.8$) than in 3A-p (Soros has weight 3 and is the only one who sees the threshold line, and $T(t) < 0.8$), and higher attack rates in 3A-o than in 5N-o (defined as before except with "optimistic" fundamentals $T(t) \geq 0.8$). The raw attack rates in our data are 3.2% in 5N-p versus 7.0% in 3A-p (contrary to the first part of the prediction) and are 2.9% in 3A-o versus 2.7% in 5N-o (consistent with the prediction but negligible). To control for information conditions and SP, one compares intercept coefficients

the basic Switch to Attack Logit. The first comparison becomes negligible (log-odds of -1.70 in 5N-p versus -1.76 in 3A-p) while the second becomes more favorable (-1.95 in 3A-o versus -2.57 in 5N-o). Thus the evidence on the second conditional prediction is mixed.

6. Discussion

Second and third generation models raise as many questions as they answer. What governs the timing of speculative attacks? What factors make attacks more or less frequent, and more or less likely to succeed? The models, and the available field evidence, are either silent or inconsistent on these important matters. The few existing laboratory studies use a simplified one-time, simultaneous-move setting that doesn't capture the dynamic nature of currency speculation.

Our laboratory experiment allows speculators to switch in and out of attack mode and lets the strength of fundamentals change in real time. It controls explanatory variables and so their impact on performance has a causal interpretation. Thus the experiment provides direct evidence on the models' predictive power in a simple dynamic setting, and offers a fresh perspective on field evidence. The main results can be summarized as follows.

1. Fundamentals. The probability of a successful attack indeed is a decreasing function of the strength of the fundamentals. The function appears to be continuous, not a jump from 0 to 1 as in some models. For example, a deterioration of the fundamentals from 90 to 80 (percent of potential speculation required to force a devaluation) doubles the probability of a successful attack in our data.

2. Information. Successful attacks are more likely and come sooner when speculators have information about the strength of fundamentals and about the activities of other speculators. In our symmetric data, for example, providing precise public information about net speculative positions shortens the average time to devaluation by about 15%, and providing precise public information about the fundamentals has a very similar effect. The effects persist even with weaker fundamentals and with less precise information.

3. Asymmetries. The mere presence of a speculator with larger mass ("Soros") increases the probability of a successful attack and hastens its onset; e.g., the time to success decreases by almost 30%. Consistent with some theory, giving the large speculator increased size or better information does not significantly strengthen his impact in our data. Contrary to some theory, our results indicate that large speculators tend to attack later than regular ones.

4. Coordination. Speculators are more likely to switch into (and less likely to switch out of) attack mode when they are nearer to being pivotal. The effect is enhanced by information about fundamentals and speculative activity, and it persists whether fundamentals are weak or strong.

Of course, the laboratory results should not be translated directly into policy recommendations, but they should sharpen subsequent theoretical and empirical research and thus improve policy indirectly. For example, part of result 3. seems inconsistent with surveys of foreign exchange dealers, who say that the two most important advantages possessed by large traders are a "large customer base" and "better information" about the market (Cheung and Wong, 2000; Cheung and Chinn, 2001; and Cheung, Chinn and Marsh, 2004).[8] Future empirical work, both laboratory and field, should resolve the inconsistency and point to appropriate policy.

Another example is the second result that better information about fundamentals encourages speculative attacks and makes success more likely. As noted, this conclusion seems to contradict conventional wisdom, some theoretical models, and one piece of field data. In our experiment, the strength of fundamentals is, by design, independent of what information is available. The two variables might be correlated in field data. Indeed, the central bank might strategically increase or decrease the availability of information, and from a decrease speculators might infer that fundamentals are weak. If this strategic effect is sufficiently strong, it could account for the conventional result and its very different policy implications. One way to pursue the question is to conduct laboratory experiments with information controlled by a player in the role of a central bank. The baseline results of our completed experiment would allow clear inferences about the strategic effect.

To conclude, experimental methods have made useful contributions to many theoretical and applied fields of economics in the last ten or twenty years, but have hardly touched international finance. We believe that the work reported here offers evidence that is complementary to both field studies and theoretical models, and we hope that it encourages greater cross fertilization.

[8] Wei and Kim (1997) find a very weak association between big players' position and the subsequent currency movement, and infer that large players do not generally have better information.

Bibliography

Arikawa, Yasuhiro, Kumi Suzuki-Loffelhotz, and Kenshi Taketa, 2006, "Experimental Analysis on the Role of a large Speculator in Currency Crises," Bank of Japan, Institute for Monetary and Economic Studies Discussion Paper No. 2006-E-4.

Angeletos, George-Marios, Christian Hellwig and Alessandro Pavan, 2006, "Learning in Global Games of Regime Change and the Dynamics of Crises," Manuscript, UCLA Economics Department, February.

Angeletos, George-Marios and Iván Werning, 2004, Crises and Prices: Information Aggregation, Multiplicity and Volatility, NBER Working Paper 11015.

Allegret, Jean-Pierre and Camille Cornand, 2005, "The Pros and Cons of Higher Transparency: The Case of Speculative Attacks," GATE WP 05-02.

Bannier, Christina E [nee Metz], 2004, "The Role of Information Disparity in the 1994/95 Mexican Peso Crisis: Empirical Evidence," manuscript, Finance Department, Goethe University Frankfurt, March.

Christina E. Bannier, 2005, "Big Elephants in Small Ponds: Do Large Traders Make Financial Markets More Aggressive?" *Journal of Monetary Economics* 52, 1517–1531

Bannier, Christina E. and Frank Heinemann, 2005, "Optimal Transparency and Risk - Taking to Avoid Currency Crises," *Journal of Institutional and Theoretical Economics* 161, 374 – 391.

Brunnermeier, M. K. and John Morgan, 2004, "Clock Games: Theory and Experiments," manuscript, Haas School, UC Berkeley, September.

Carlsson, H., and E. Van Damme, 1993, "Global Games and Equilibrium Selection," *Econometrica*, 61, 989–1018.

Chang, Roberto and Andres Velasco (1999), "Financial Fragility and the Exchange Rate Regime," NBER Working Paper 7272.

Cheung, Yin-Wong and Clement Yuk-Pang Wong, 2000, "A Survey of Market Practitioners' Views on Exchange Rate Dynamics," *Journal of International Economics* 51, 401-419.

Cheung, Yin-Wong and Menzie Chinn, 2001, "Currency Traders and Exchange Rate Dynamics: A Survey of the U.S. Market," *Journal of International Money and Finance* 20, 439-471.

Cheung, Yin-Wong, Menzie Chinn and Ian Marsh, 2004, "How Do UK Foreign Exchange Dealers Think their Market Operates," *International Journal of Finance and Economics* 9, 289-306.

Cornand, Camille, forthcoming, Speculative Attacks and Informational Structure: An Experimental Study, *Review Of International Economics*.

Corsetti, Giancarlo, Amil Dasgupta, Stephen Morris, and Hyun Song Shin, 2004, "Does One Soros Make a Difference? A Theory of Currency Crisis with Large and Small Traders," *Review of Economic Studies*, 71, 87-113.

Corsetti, Giancarlo, Paolo Pesenti, Nouriel Roubini (1999), "Paper Tigers? A Model of the Asian Crisis," *European Economic Review*, Vol. 43: 1211-1236.

Costain, James (2004), A Herding Perspective on Global Games and Multiplicity. U. Carlos III Economics Working Paper 03-29 (08), May 2003; revised May 2004.

Costain, James, Frank Heinemann, and Peter Ockenfels (2005). "Multiple Outcomes of Speculative Behavior in Theory and in the Laboratory," Department of Economics manuscript, Universidad Carlos III Madrid. (April).

Dooley Michael P. (1997). "A Model of Crises in Emerging Markets," NBER Working Paper # 6300.

Drazen Allan and Paul R. Masson (1994), "Credibility of Policies versus Credibility of Policymakers," *Quarterly Journal of Economics* 104 (August), 735-54.

Economist, 1998. Finance and economics: Freefall. Economist 349(8089),78-79.

Eichengreen, B., Mathieson, D., with Chadha, B., Jansen, A., Kodres, L., Sharma, S., 1998. Hedge funds and financial market dynamics. Occasional Paper 166. International Monetary Fund, Washington, D.C.

Federal Reserve Bank of New York, 1998, *Foreign Exchange and Interest Rate Derivatives Markets Survey Turnover in the United States* (New York: Federal Reserve Bank of New York, September 29th).

Flood, Robert P. and Peter M. Garber (1984), "Collapsing Exchange Rate Regimes: Some Linear Examples," Journal of International Economics 17 (August), 1-13.

Flood, Robert P. and Nancy P. Marion (1999), "Perspectives on the Recent Currency Crisis Literature," International Journal of Finance and Economics, 4 (January), 1-26.

Guimaraes, Bernardo (2004), "Dynamics of currency crises with asset market frictions," Yale

Economics mimeo (May).

Heinemann, Frank, and Gerhard Illing, “Speculative Attacks: Unique Equilibrium and Transparency,“ *Journal of International Economics* 58 (2), 2002, pp. 429-450..

Heinemann, Frank, Rosemarie Nagel, and Peter Ockenfels, 2004, “The Theory Of Global Games On Test: Experimental Analysis Of Coordination Games with Public And Private Information,” *Econometrica*, Vol. 72, No. 5 (September, 2004), 1583–1599.

Henderson, Dale and Steven Salant, "Market Anticipations of Government Policies and the Price of Gold", *Journal of Political Economy* 86 (August 1978), pp. 627-48

Huberman, Bernardo, and Natalie Glance (1993), “Evolutionary Games and Computer Simulations,” *Proceedings of the National Academy of Science* 90 (August) 7716-7718.

Krugman, Paul (1979), “A Model of Balance of Payment Crises,” Journal of Money, Credit and Banking, 11 (August), 311-25.

Krugman, Paul (1999), “Balance Sheets, the Transfer Problem, and Financial Crises,” in International Finance and International Crises, ed. By Peter Isard, Assaf Razin and Andrew K. Rose, International Monetary Fund, Washington DC.

Morris, S., and H. S. Shin, 1998, “Unique Equilibrium In A Model Of Self-Fulfilling Currency Attacks,” *American Economic Review*, 88, 587–597.

Obstfeld, Maurice (1995), “The Logic of Currency Crises,” in *Monetary and Fiscal Policy in an Integrated Europe*, eds. Barry Eichengreen, Jeffrey Frieden, and Jurgen von Hagen, Springer Verlag, London.

Obstfeld, Maurice, 1996, “Models Of Currency Crisis With Self-Fulfilling Features,” *European Economic Review*, 40, 1037–1047.

Obstfeld, Maurice, (1998), “EMU: Ready or Not?” Princeton Essays in International Finance, no. 209, July.

Park, A. and L. Smith, 2003, “Caller Number Five: Timing Games that Morph from One Form to Another,” University of Michigan manuscript, April.

Sesit, M.R., Pacelle, M., 1998. Selling blitz hits dollar 7.4% on Tokyo’s move: hedge funds scramble in bid to cut losses. Wall Street Journal, October 8, C1.

Wei, Shang-Jin and Jungshik Kim, 1997, The Big Players in the Foreign Exchange Market: Do They Trade on Information or Noise?, NBER Working Paper 6256.

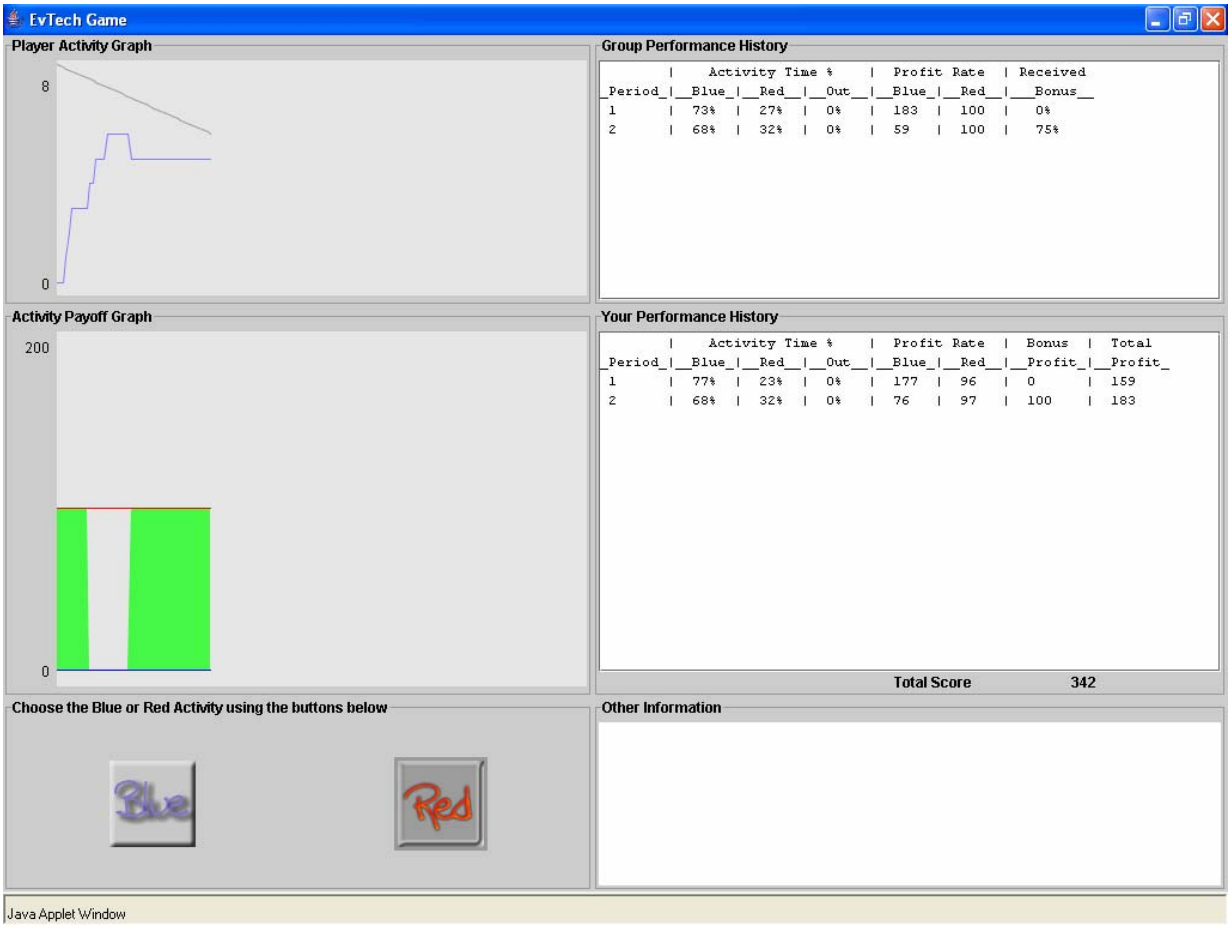

Figure 1. Example of speculator's screen at mid-period.

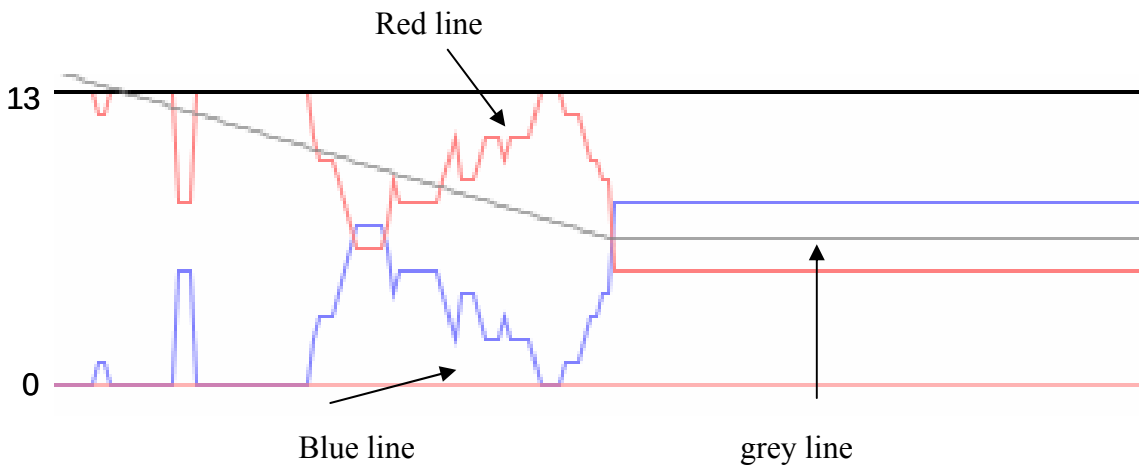

Figure 2. Player activity graph from period 19 on Feb. 8, 2005. The threshold (grey line) falls to 0.5 (6.5) after 45 seconds. There were 9 human players, but player 4 had a weight of 5 in this period. No player saw any of the lines in this period.

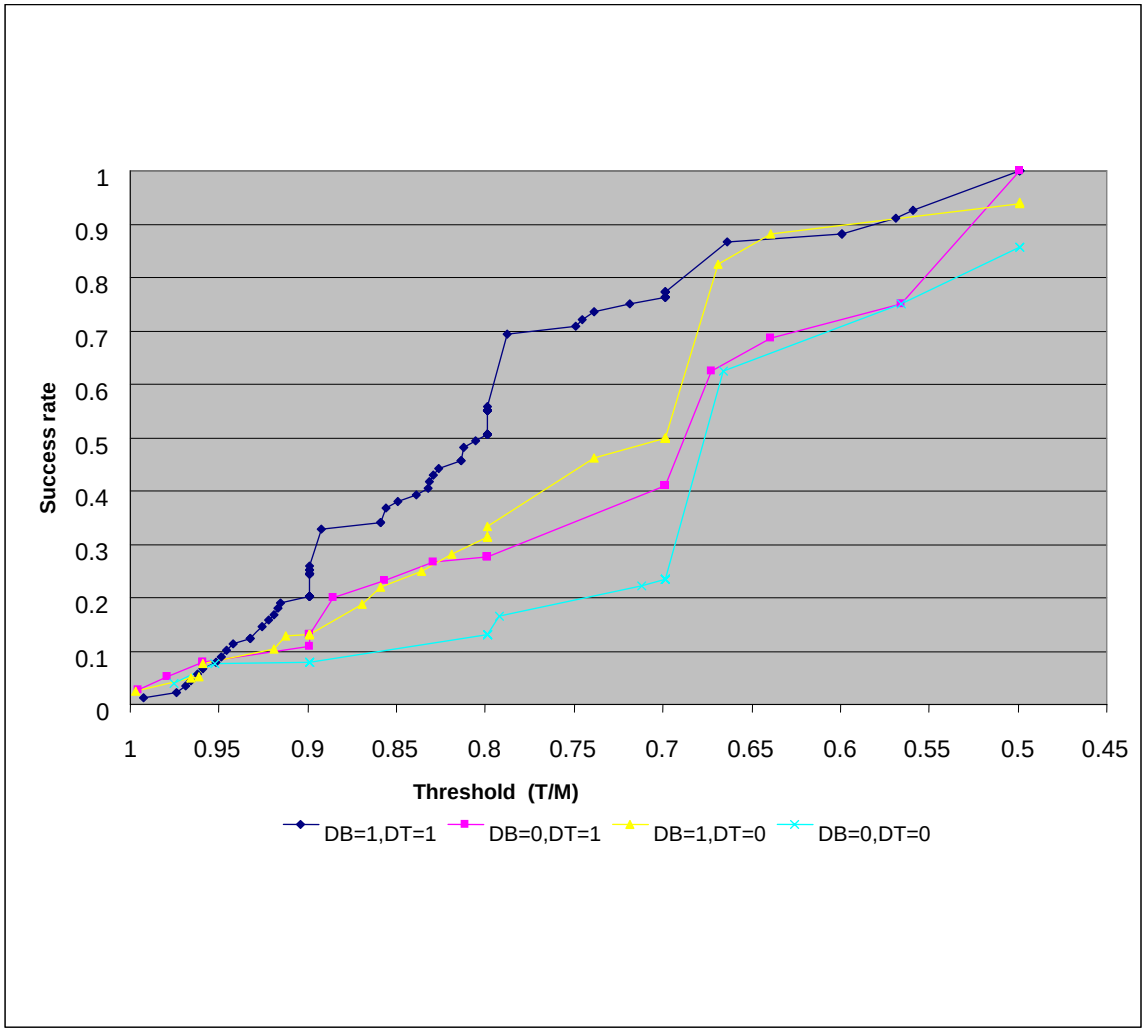

Figure 3: Success Rates for Speculative Attacks by the Strength of Fundamentals in Symmetric Periods.

The success rate at threshold x is defined as S(x)/[S(x)+U(x)], where S(x) = number of periods where an attack succeeded at a threshold at or above x, and U(x) = number of periods where the threshold declined to x without a successful attack. The horizontal axis is the threshold T as a fraction of the mass M of players.

Table 1: Experimental Design

Session number	Session date	Number of periods with max player weight[a]			Number of players	Number of periods with displayed lines[b]			Number of periods with asymmetric information[c]			Number of periods with steep[d] fundamental		Players experi-enced?
		=1	=3	=5		T	B	TB	A	D	N	yes	no	
85	7/21/04	15	0	0	6	10	11	8	0	0	0	2	13	No
86	8/16/04	36	0	0	4	24	24	16	0	0	0	8	28	No
87	8/18/04	23	0	0	6	19	16	12	0	0	0	5	18	No
88	9/28/04	36	0	0	7	24	24	16	0	0	0	8	28	No
89	10/6/04	18	0	0	5	10	10	6	0	0	0	4	14	No
91	10/13/04	18	0	0	8	10	10	6	0	0	0	4	14	No
94	11/19/04	0	31	0	6	18	31	18	6	18	7	4	27	No
100	1/25/05	11	10	12	7	13	33	13	8	6	8	0	33	No
101	1/31/05	8	9	11	11	13	28	13	7	7	6	0	28	No
102	2/1/05	9	12	14	8	15	19	9	9	9	8	5	30	No
103	2/7/05	7	12	8	7	11	18	8	7	7	6	3	24	Yes
104	2/8/05	12	14	18	9	17	25	10	11	10	11	5	39	Yes
Sum		193	88	63		184	249	135	48	57	46	48	296	

Note: The table and subsequent analysis excludes practice periods, the first 5 periods in each session, periods not related to speculative attacks, and 16 periods in which the attack was not feasible. The sessions after 10/13/04 include only the extreme information treatments (sharp lines or none) regarding fundamentals and speculative activity, while the prior sessions also include the intermediate (thick line) treatments.

[a]Symmetric periods are indicated by max player weight =1, regular Soros periods by =3 and BigSoros periods by =5.

[b]T indicates that the threshold line is displayed; B indicates that the blue (attack) line is displayed; and TB indicates that both lines are displayed.

[c]A indicates that only Soros sees the threshold line; D indicates that all players see the threshold line; and N indicates that no players see the threshold line.

[d]The fundamental path is classified as steep if $T(t)$ fell to $0.5M$ within 50 sec., or to $0.6M$ within 30 seconds.

Table 2: Coefficient Estimates for Symmetric Periods

Panel A. Pooled Data

Dependent Var.	Success	τ in sec	Mingap	Bluearea	τ in sec	Ticks
Model	Logit	Tobit	Tobit	Tobit	PHM	PHM
Intercept	-2.29 **	5.87 **	-.31 **	6.88 **		
DT	1.61 **	-.81 **	-.32 **	- .69 **	1.08 **	1.03 **
DB	1.59 **	-.88 **	-.32 **	- .87 **	1.29 **	.91 **
Steep	2.96 **	-1.33 **	-.71 **	-1.62 **	1.85 **	
T(t)						-7.04 **
No. of obs.	193	193	193	193	193	193
Censor if success		=0	=1	=0	=0	=0
Type of censoring		right	left	right	right	right
No. censored obs.		87	106	87	87	87

Panel B. Session dummies included

Dependent Var.	Success	τ in sec	Mingap	Bluearea	τ in sec	Ticks
Model	Logit	Tobit	Tobit	Tobit	PHM	PHM
Intercept	-3.16 *	5.58 **	-.25 +	6.53 **		
DT	2.87 **	-.85 **	-.48 **	-.88 **	1.47 **	1.45 **
DB	2.70**	-.92 **	-.43 **	-1.14 **	1.68 **	1.17 **
Steep	4.51 **	-1.28 **	-.72 **	-1.47 **	2.33 **	
T(t)						-8.75 **
exp103	.45	-.08	.17	.22	.09	.23
exp102	-.00	-.29	-.01	.49	.61	.55
exp101	13.25	-.40	-5.93	.12	.66	.87
exp100	-2.27 +	.71 +	.47 *	1.61 **	-1.14 +	-.76
exp91	-.43	.16	.28 +	.56	-.36	-.31
exp89	1.59	-.59 +	-.24	.39	1.19 +	1.28 +
exp88	-2.60 +	.78 *	.53**	1.31 **	-1.25 +	-1.22 +
exp87	-3.37 *	.97 *	.69**	.89 +	-1.70 *	-1.89 *
exp86	1.43	-.13	-.19	- .03	.22	.36
exp85	-2.57 +	.75 +	.56**	1.50 **	-1.25 +	-1.05 +
No. of obs.	193	193	193	193	193	193
Censor if success		=0	=1	=0	=0	=0
Type of censoring		right	left	right	right	right
No. censored obs.		87	106	87	87	87

Note: ** denotes p-values ≤ .001, * denotes p-values ≤ .01, + denotes p-values ≤ .1. Data are all symmetric periods reported in Table 1.

Table 3: Coefficient Estimates for Asymmetric Sessions

Panel A. Pooled Data

Dependent Var.	Success	τ in sec	Mingap	Bluearea	τ in sec	Ticks	Attack[a]
Model	Logit	Tobit	Tobit	Tobit	PHM	PHM	PHM
Intercept	-5.12 *	6.76 **	.25 **	9.31 **			
PT	3.68 *	-1.24 **	-.24 **	-1.66 **	1.87 **	1.16 +	2.48 **
DB	5.18 *	-2.42 *	-.28 **	-3.74 **	3.56 *	2.92 *	.31 **
PT*DB	-1.96	.82 *	-.02	1.26 *	-1.23 +	-.45	
Steep	14.52	-.82 **	-3.54	1.62 **	1.27 **		
T(t)						-5.01 **	-2.47 **
Soros	3.80 +	-1.94 +	-.22 **	-2.39+	2.93 *	2.42 +	.31 **
DB*Soros	-2.19	1.55 +	-.003	2.28 +	-2.33 +	-1.73	
No. of obs.	198	198	198	198	198	198	2448
Censor if success		=0	=1	=0	=0	=0	
Type of censoring		right	left	right	right	right	right
No. censored obs.		44	154	44	44	44	282

Panel B. Session dummies included

Dependent Var.	Success	τ in sec	Mingap	Bluearea	τ in sec	Ticks	Attack[a]
Model	Logit	Tobit	Tobit	Tobit	PHM	PHM	PHM
Intercept	-6.05 *	6.94 **	.26 **	8.39 **			
PT	3.90 *	-1.28 **	-.23 **	-1.22 **	2.31 **	1.59 **	3.23 **
DB	6.84 **	-2.78 **	-.39 **	-4.07 **	4.75 **	3.76 **	.17 *
PT*DB	-1.80	.97 **	-.05	.94 +	-1.72 **	-.94 +	
Steep	14.54	-.63 **	-3.29	-.93 **	1.19 **		
T(t) [a]						-5.24 **	-3.32 **
Soros	4.04 *	-1.90 **	-.20 **	-1.94 +	3.41 *	2.71 *	.21 *
DB*Soros	-2.02	1.60 *	-.08	1.91 +	-2.79 *	-2.01 +	
exp103	1.35	-.26 +	-.02	.47 +	.54 +	.61 +	
exp102	.89	-.60 **	-.06	.35	1.36 **	1.38 **	
exp101	11.90	-.01	-3.09	1.29 **	.06	.78 *	
exp100	-2.89 *	.65 **	.37 **	2.10 **	-.90 *	-.64 +	
exp94	-1.17	-.05	.20	.86 **	.27	.31	
No. of obs.	198	198	198	198	198	198	2448
Censor if success		=0	=1	=0	=0	=0	
Type of censoring		right	left	right	right	right	right
No. censored obs.		44	154	154	44	44	282

Note: ** denotes p-values ≤ .001, * denotes p-values ≤ .01, + denotes p-values ≤ .1. Data in all columns except the last are from all periods of sessions conducted after 11/1/04 as reported in Table 1. [a]The columns for Attack are based on all data from all sessions reported in Table 1, so Panel B doesn't have room to report the full set of dummy coefficients.

Table 4: Timing of Speculative Attacks

	No. of obs.	Mean SP	t Test	MW Test
Soros Knows Same			6.05**	4.44**
Regular Players	2033	5.25		
Soros Players	178	3.70		
Soros Knows More			5.78**	4.32**
Regular Players	664	5.36		
Soros Players	93	3.32		

Note: MW refers to the nonparametric Mann-Whitney (or Wilcoxon rank-sum) test, and t refers to the unequal variance version of Student's t test for equality of means. The variable Shortfall from Pivotality (SP) is defined in the text. Data are all periods from sessions conducted after 10/13/04 reported in Table 1.
** denotes p-values ≤ .001.

Table 5: Switching Behavior

Dependent Var.	Switch to Attack	Switch to Attack	Stay in Attack	Stay in Attack
Model	Logit	Logit + Subject fixed effects	Logit	Logit + Subject fixed effects
Intercept	-3.04 **	-2.80	2.37 **	2.77
DT	0.15 **	-0.20	0.41 **	0.64 **
DB	0.24 **	-0.18	0.21 **	0.92 **
DT*DB		1.18 **		-0.14
SP	-0.22 **	-0.34 **	-0.09 **	-0.14 **
SP*DT		0.09 +		0.05
SP*DB		0.13 **		-0.07 +
SP*DT*DB		-0.29 **		-0.08
No. of obs.	109051	109051	27061	27061

Note: ** denotes p-values ≤ .001, * denotes p-values ≤ .01, + denotes p-values ≤ .10. Data include all ticks prior to a successful attack in all symmetric periods and for all subjects reported in Table 1.