UNIVERSIDAD CARLOS III DE MADRID

TESIS DOCTORAL

Bubbles, Currency Speculation, and Technology Adoption

Autor:
Carlos José Pérez González

Director/es:
Manuel Santos Santos
Francisco Marhuenda Hurtado

DEPARTAMENTO DE ECONOMIA

Getafe, noviembre de 2010

A la Consuela

Contents

Preface (in Spanish)

Esta es una tesis de carácter teórico que propone y discute tres aplicaciones de la teoría de juegos al análisis económico. Consta de tres capítulos:

El primer capítulo se titula "Bubbles with Random Behavioral Trading". Este capítulo se enmarca dentro de la teoría de las finanzas conductuales. Es un análisis teórico de las condiciones necesarias para que aparezcan y se sostengan burbujas especulativas en los mercados de activos. Explica cómo la presencia de cierta masa de agentes irracionales en los mercados induce a los especuladores a elegir estrategias de inversión que favorecen una subida injustificada de los precios. La idea básica es que los especuladores entienden que, a menudo que los precios crecen durante la burbuja, cada vez más agentes irracionales entran en el mercado atraídos por la subida. Esto permite a los primeros vender a los segundos a precios aún más altos en el futuro. La principal aportación original de este capítulo es la consideración del carácter aleatorio del comportamiento de los agentes irracionales. El modelo está construido sobre la base del artículo "Bubbles and Crashes" de Abreu y Brunnermeier, que considera el comportamiento de los agentes irracionales como constante. Este capítulo demuestra que introducir aleatoriedad en ese punto permite sostener una burbuja, aún cuando se elimina el principal supuesto de Abreu y Brunnermeier: información asimétrica sobre el valor fundamental de los activos.

El segundo capítulo se titula "Currency Speculation in a Game-Theoretic Model of International Reserves" y está escrito en colaboración con Manuel S. Santos. Este capítulo propone un modelo de especulación en los mercados internacionales de divisas. Cuando un gobierno trata de mantener un tipo de cambio fijo que no se corresponde con el tipo de equilibrio, corre el riesgo de que su moneda sea objeto de un ataque especulativo. Si los especuladores consideran que el gobierno no tiene suficientes reservas para hacer frente al ataque, pueden coordinarse vendiendo masivamente la moneda doméstica —al gobierno— al tipo de cambio fijo, provocando así una devaluación, y recomprando posteriormente a un precio más ventajoso. El volumen

de reservas de que dispone el gobierno es clave para determinar si un ataque tendrá o no lugar y cómo de provechoso será este último para los especuladores que lo lleven a cabo; un ataque exitoso produce más beneficios para los atacantes cuantas más reservas hay en juego. La principal aportación original de este artículo es la incorporación de este papel fundamental de las reservas en un modelo de ataques especulativos con información asimétrica sobre los fundamentales de la economía. Extiende los resultados de existencia y unicidad de equilibrio establecidos por Morris y Shin (1998) en un modelo sin reservas, y demuestra que algunas de las conclusiones obtenidas por estos autores acerca del valor de la transparencia de la política monetaria no son robustas.

El tercer capítulo se titula "Technology Adoption with Switching Costs and Learning by Doing" y está escrito en colaboración con Carlos J. Ponce. En este capítulo se discute el problema de la adopción de nuevas tecnologías en un contexto dinámico de competencia en precios. En nuestro modelo, dos vendedores compiten en cada periodo por un comprador de vida corta y demanda unitaria. Uno de los vendedores tiene la opción de adoptar una nueva tecnología que presenta costes de ajuste y aprendizaje por experiencia. Mostramos que existen tecnologías eficientes que, sin embargo, no se adoptan en equilibrio. La razón de ello estriba en que el vendedor que no tiene la opción de adoptar tiene incentivos a bajar su precio de venta para evitar que el otro vendedor aprenda la nueva tecnología. Esto reduce los beneficios derivados de adoptar la nueva tecnología y, en algunos casos, llega a hacer que la adopción sea inviable. También se muestra que existe un sesgo endógeno a adoptar tecnologías que se aprenden más rápido, sean o no las más eficientes. Ambos resultados son originales.

Getafe, Madrid.
Noviembre 2010.

CHAPTER 1

Bubbles with Random Behavioral Trading

ABSTRACT. We present a model in which a bubble can persist despite that (a) it is common knowledge among arbitrageurs, and (b) they can make it burst.

1.1. Introduction

In its most inclusive form, the efficient markets hypothesis (EMH henceforth) states that, besides the fact that behavioral traders are present in the stock market and their trades do not tend to cancel each other, the trading activity of rational arbitrageurs keeps stock prices close to their fundamental values (Shleifer, 2000, page 2). The EMH thus presents arbitrageurs as a stabilizing force whose action neutralizes the distorting effect of behavioral traders. On the other hand, literary explanations of the emergence and persistence of bubbles (Kindleberger and Aliber, 2005) argue on the contrary; they say that bubbles rise with the complicity of arbitrageurs who buy overpriced stocks that they later sell to behavioral traders at an even higher price. Arbitrageurs represent here a destabilizing force which acts against informational efficiency.

Informal accounts of bubbles usually go as follows: After some particularly good news about the profitability of a certain investment, arbitrageurs buy stocks bidding up their price. The initial price increase calls the attention of behavioral traders who extrapolate the most recent trend and get into the market seeking for capital gains. As behavioral traders keep entering the market, the price grows even higher, at an unsustainable rate. Then, a spiral of speculation develops which allows arbitrageurs to leave the market before the inevitable burst of the bubble.

From this perspective, the strategic side of a bubble has a very simple logic. Arbitrageurs use the initial price increase as a bait to attract behavioral traders into the market. Since they know that behavioral traders buy stocks whenever they see signs of a price up-trend, they do their best by buying stocks, inducing an initial price increase, and selling them later to the excited behavioral traders. However, a mere
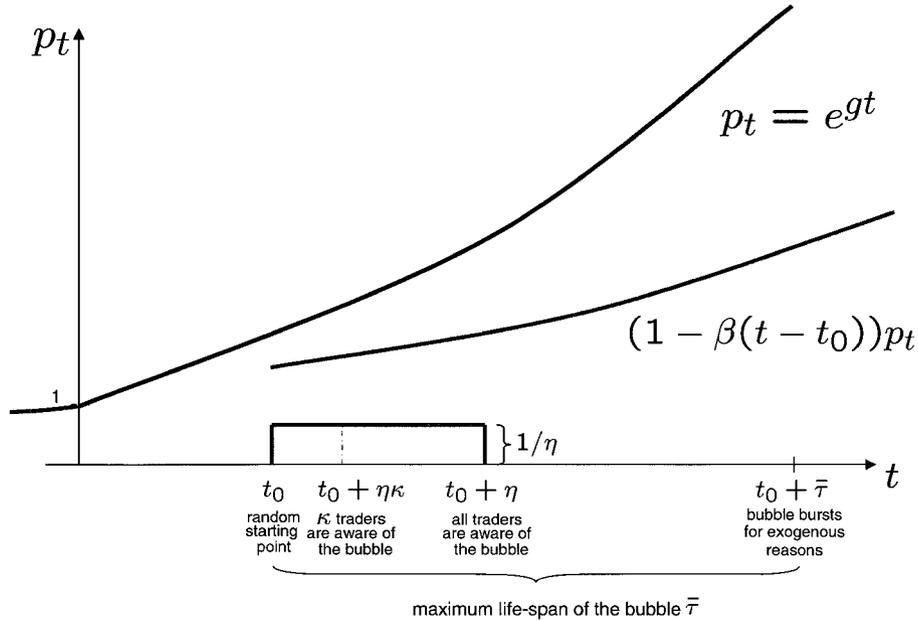
FIGURE 1. Illustration of price paths from Abreu and Brunnermeier (2003).

price up-trend cannot be a sufficient condition for the rise of a bubble, since otherwise we would observe bubbles occurring everywhere, all the time. Certain complementary conditions should be met. One of them is precisely the one that would justify the price increase; a series of good news about the stocks or, in the words of Hyman P. Minsky, a "displacement" which could lead behavioral traders to interpret the initial price increase as an evidence that something fundamentally good has happened.

1.1.1. The Model of Abreu and Brunnermeier (2003). Dilip Abreu and Markus K. Brunnermeier proposed a model which rationalizes the idea that speculation can be destabilizing. They explain why in the presence of "irrationally exuberant" behavioral traders, speculators may choose to ride the bubble rather than to attack it.

In the model of Abreu and Brunnermeier there is a mass of size one of rational arbitrageurs and a mass of behavioral traders. The price process is exogenously given as depicted in figure 1. The fundamental value of stocks is e^{gt} before the random date t_0 and drops to $(1 - \beta(t - t_0))e^{gt}$ at $t = t_0$. The price $p_t = e^{gt}$ equals the fundamental value until t_0, but keeps growing at rate g thereafter. That is, a fraction $\beta(\cdot)$ of the price is not justified by fundamentals from t_0 onwards.

The price is kept above its fundamental value by behavioral traders. As soon as the cumulative selling pressure by arbitrageurs exceeds κ, the absorption capacity of behavioral traders, the stock price drops back to its fundamental value. If this does not happen, the bubble is assumed to burst for exogenous reasons at $t_0 + \bar{\tau}$.

Arbitrageurs do not observe t_0. At each instant t from t_0 to $t_0 + \eta$ a new cohort of mass $1/\eta$ of arbitrageurs is informed of the fact that $t_0 \leq t$. As figure 1 illustrates, only after instant $t_0 + \eta$ is every arbitrageur aware of the misprizing. The key feature of this information structure is that no arbitrageur can tell how many arbitrageurs became aware before him.

Arbitrageurs wish to sell their stocks at the highest possible price (note, however, that because $\kappa < 1$ not all arbitrageurs can sell before the date of burst). This could be done by selling just before the crash, but its exact timing is not known to anyone. This is because the behavior of each arbitrageur—and thus the date of burst of the bubble—depends on the random date t_0 about which arbitrageurs are imperfectly and asymmetrically informed.

Abreu and Brunnermeier show that there is an equilibrium in which each arbitrageur rides the bubble for some time. In particular, they show that the misprizing lasts beyond instant $t_0 + \eta\kappa$—at which the mass of arbitrageurs who are aware of the bubble is sufficient to make it burst. Arbitrageurs wait to sell because they know that they are possibly among the first κ who became aware of the bubble.

1.1.2. Our proposal. In the model of Abreu and Brunnermeier (2003) the behavior of the behavioral traders is deterministic and common knowledge among arbitrageurs. On the other hand, they attribute the mayor role in explaining the rise and persistence of bubbles to the asymmetric information about the fundamental value of stocks.

We propose here a model which takes the opposite view. We argue that it is the behavior of the behavioral traders, by its essentially irrational nature, what puzzles arbitrageurs more. The basic idea is that the behavior of the crowd—the market sentiment—is hard to predict and subject to sudden changes. Of course, it exhibits certain regularities as, for example, trend-chasing behavior, but it is ultimately seen as a random phenomenon by rational investors.

On the other hand, we assume that the fundamental value of stocks is common knowledge among arbitrageurs. This is an exaggeration, but it helps to make our point clearer. Recall that it is the asymmetry and not the lack of perfect information about fundamentals what drives the results in Abreu and Brunnermeier (2003). Our model presumes

that the fundamental valuations that rational investors make of stocks are based mainly on publicly available "hard" information, following similar valuation procedures. This is something that can hardly be true for the assessment of market sentiment. We defer the analysis of asymmetric information about behavioral trading to a future paper.

1.2. The Model

We consider a market for stocks. There is a continuum of mass $0 < \mu < 1$ of arbitrageurs who seek to maximize the expected discounted value of their transactions. They can sell and buy back shares at any time at the discounted cost $0 < c < 1$ per transaction, but they are constrained on the maximum long and short positions they can take. In particular, the selling pressure exerted by each arbitrageur must lie within the unit interval at any time. There is also a continuum of mass 1 of behavioral traders whose trading behavior is exogenously given and summarized by their aggregate absorption capacity κ. The absorption capacity of behavioral traders is an unobservable stochastic process.

The fundamental value of the stock is e^{rt} for all $t \geq 0$. The *pre-crash* price of the stock is given by e^{gt}, with $g > r$. The pre-crash price is the market price of the stock as long as the aggregate selling pressure of arbitrageurs stays below the absorption capacity of behavioral traders. At the first instant at which this ceases to happen, the market price drops to its *post-crash* price e^{rt}.

The market price process should be interpreted as follows: A series of good news have happened before $t = 0$ which justified the higher price growth rate g. From then on, this higher rate is no longer justified by fundamentals, which now grow at rate r. A bubble starts at $t = 0$ which persists until the first instant at which the selling pressure of arbitrageurs equals (or exceeds) the absorption capacity of behavioral traders.

The absorption capacity depends on the realization of a random variable X with standard uniform distribution. It is a hidden state variable, but its probability law is common knowledge among arbitrageurs. Given a realization x of X, the absorption capacity κ has the form

$$(1.1) \qquad \kappa(x,t) := \begin{cases} x \sin\left(\frac{t}{x}\right) & \text{if } 0 \leq t < \pi x \\ 0 & \text{if } t \geq \pi x. \end{cases}$$

Various sample paths of κ are plotted in figure 2. It is a continuous function which takes only positive values, non-decreasing in x for t fixed and uni-modal in t for fixed x. The shape of the sample paths of κ reflects the idea that, as long as the bubble persists, more and
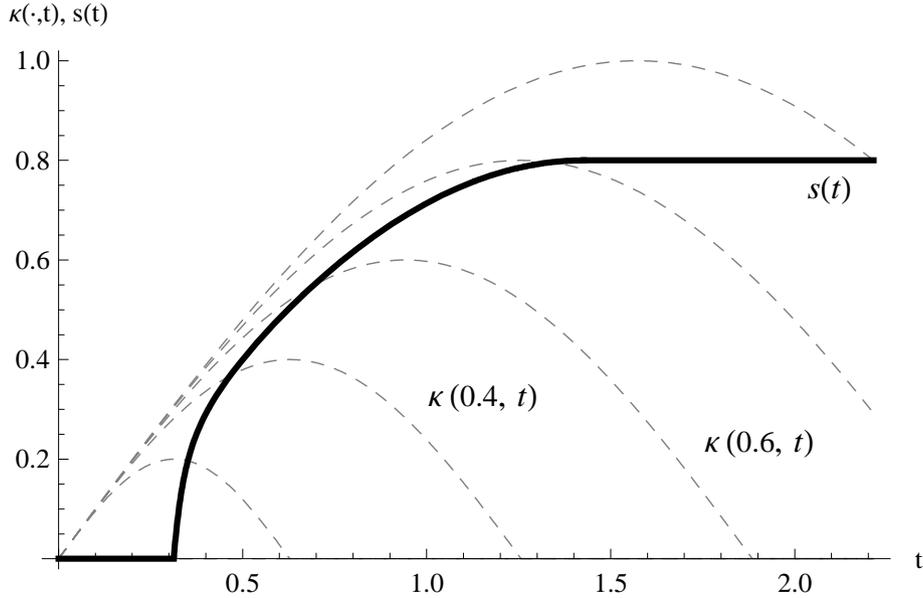
FIGURE 2. Various sample paths of κ (dashed) and the equilibrium aggregate selling pressure for the parameter values $\mu = 0.8$ and $g - r = 0.1$.

more money from the behavioral traders enters the market until the maximum established by the state variable X is reached; at this point, the process is reversed. Larger states correspond unambiguously to more aggressive behavior from the part of behavioral traders; for larger values of x two things happen: (a) more money flows into the market, and (b) it stays in for longer. This allows us to interpret the state variable X as an index of their confidence on the persistence of the bubble.

For every state $x > \mu$, there is an interval of time in which the absorption capacity exceeds μ, the maximum feasible aggregate selling pressure. We call such an interval a *mania*. A mania is, thus, a period in which not even a coordinated attack from the part of arbitrageurs can burst the bubble; it is a paradise for speculation. We will see below that if a mania could not happen, we would only have equilibria in which the bubble collapses at $t = 0$ for every state. Its relevance comes form the fact that, whereas the price growth gives incentives to stay in the market, the possibility of a mania allows an arbitrageur to stay in for longer than others.

As will become clear below, our choice of κ is to a great extent arbitrary. In the first place, we want a process whose sample paths

start and end below μ—and for which a mania takes place with positive probability—to represent the behavioral traders's confidence. On the other hand, our proofs make use of the continuity and the following monotonicity property of κ: if $x_1 > x_0$, then $\kappa(x_1, t) > \kappa(x_0, t)$ whenever $\kappa(x_0, t) > 0$. Both properties seem quite natural in our setting.

The possibility of occurrence of a mania is the cornerstone of our main result. This means that its empirical relevance rests on the plausibility of an assumption about the relative market power of the real-world investors represented by our two sets of traders. The customary association of arbitrageurs with institutional investors and of behavioral traders with individual investors seems overly naive to us. What should count for somebody to qualify as an arbitrageur or as a behavioral trader is the way in which he perceives and analyzes the situation rather than his main occupation. People who think strategically, who understand the logic of the situation at hand and act accordingly, should qualify as arbitrageurs; those who exhibit trend-chasing behavior, who form extrapolative expectations from the most recent data alone, should not.

It is simply impossible to estimate the proportion of people that should be assigned to each class. However, we argue that the amount of people who may qualify as behavioral traders tends to be larger in both groups, of institutional and of individual investors, during a bubble. If this is true, that is, if most of the irrational money which enters the stock market during a bubble is not in the market under normal circumstances, we would have also an explanation for why behavioral traders do not necessarily disappear in the long run as suggested by Friedman (1953); behavioral traders would save in quiet times the money that they lose in the bubble. Regarding individual investors, we refer once more to the work of Kindleberger and Aliber (2005) who comment on how people who hardly ever bought stocks before entered the market, even getting themselves into debt, at the peak of the great bubbles. On the other hand, Galbraith (1994) and Shiller (2005) expound that many institutional investors, besides their presumed superior training, are subject to the same behavioral biases as the inexperienced individual investors. As an example, Greenwood and Nagel (2009) have shown that young fund managers betted disproportionately on technology stocks and exhibited trend-chasing behavior during the technology bubble of the late 1990's. They argue, thoroughly, that young managers formed adaptive expectations which put more weight on the most recent trend than those who were more experienced. But it is not just that their behavior resembled that of behavioral traders, Greenwood and Nagel also found that such behavior was rewarded by substantial

inflows of new capital which enhanced their market power along the bubble. Again, the question is how much wealth is in the hands of more and less sophisticated investors—a question which can hardly be given a serious answer. We do not pretend to have a solid quantitative argument supporting our assumption; we just say that, once one abandons the individual versus institutional investors dichotomy, it becomes a possibility whose consequences are worth exploring.

Arbitrageurs receive no signal before the crash. On the other hand, the best they can do after the crash is to quit and never re-enter the market—because all transactions made after the crash are costly and worthless. There are no contingencies to which they may have to adapt besides the crash itself, and the adaptation to this event is trivial. This means that a strategy for an arbitrageur only has to specify a set of orders that will be placed sequentially as long as the bubble persists. Because transactions are costly, each arbitrageur plans, at most, a finite amount of them.

A pure strategy profile is a measurable function $\sigma : [0, \mu] \times \mathbb{R}_+ \to [0, 1]$ which specifies the selling pressure $\sigma(i, t)$ exerted by each arbitrageur $i \in [0, \mu]$ at each instant $t \in \mathbb{R}_+$. Without loss of generality, we assume that all arbitrageurs start at their maximum long position, i.e.: $\sigma(i, 0) = 0$ for all i. Their aggregate selling pressure is defined as

$$s(t) := \int_0^\mu \sigma(i, t) \, di.$$

A *trigger-strategy* for an arbitrageur specifies a unique transaction at which he completely sells out. The set of such trigger-strategies is indexed by the time at which the sales occur, which allows us to give a name to each one of them. For example, if each arbitrageur i plays some trigger-strategy t_i, we obtain the profile $\sigma(i, t) = \mathbf{1}_{[t_i, +\infty)}(t)$. A mixed trigger-strategy is a mixed strategy which only contains trigger-strategies in its support. If each arbitrageur independently draws a trigger-strategy from the same distribution F, we have that the corresponding aggregate selling pressure is $s(t) = \mu F(t)$ almost surely for all t.

Once we have defined the absorption capacity and the aggregate selling pressure, we can define the *date of burst* of the bubble.

DEFINITION 1.1. The date of burst of the bubble is the random variable defined by the function

(1.2) $$T(x) := \inf \{t : s(t) \geq \kappa(x, t), t > 0\}.$$

The date of burst determines the path of the market price. As we have said before, the market price jumps from the pre-crash price to

the post-crash price precisely at the date of burst, i.e.:

$$p(x,t) := \begin{cases} e^{gt} & \text{if } t < T(x) \\ e^{rt} & \text{if } t \geq T(x). \end{cases}$$

All transactions take place at the market price. This assumption is questionable if there are states for which the limit from the left of the aggregate selling pressure is strictly smaller than the absorption capacity at the date of burst. It would be more natural to assume that some of the orders placed at the date of burst, up to the limit imposed by the outstanding absorption capacity at that moment, are executed at the pre-crash price. Our assumption simplifies the analysis and does not change the results.

1.3. Symmetric Equilibria in Trigger-Strategies

Arbitrageurs try to do their best against other arbitrageurs. They form a rational conjecture about the date of burst and choose their most preferred selling strategy accordingly. The second step is straightforward, but to form a rational conjecture about the date of burst requires a good deal of strategic thinking.

Our aim is to show that there exist symmetric equilibria in mixed trigger-strategies. These equilibria are characterized by a mixed trigger-strategy F such that, when all the rest are playing according to it, each arbitrageur finds it optimal to play any strategy in its support. The mixed trigger-strategy F determines the aggregate selling pressure which, trough (1.2), defines the date of burst. Since there is a continuum of arbitrageurs, no-one of them can affect the date of burst, which is the only way in which the pay-off of an arbitrageur is affected by the choices of other arbitrageurs. Hence, the problem of an arbitrageur is to choose a best response given the distribution of the date of burst induced by F. A symmetric equilibrium is found when every strategy in the support of F is indeed a best response.

The pay-off from the trigger-strategy t is given by

$$v(t) := E\left[e^{-rt}p(X,t) - c\right]$$

(1.3)
$$= e^{(g-r)t}[1 - G(t)] + G(t) - c,$$

where G is the distribution function of the date of burst. An expression for the pay-off from an arbitrary strategy is given in the proof of Proposition 1.6 in Appendix A.1. The formula (1.3) neatly expresses the arbitrageur's trade-off between staying in and quitting the market. On one side, the discounted pre-crash price grows with time; on the other side, the probability of survival of the bubble decreases.

1.3.1. Pure strategies. Our first result is rather trivial. When all other arbitrageurs play the trigger-strategy $t = 0$, the bubble bursts immediately. Any transaction takes place at the post-crash price and, hence, yields, at most, $1 - c$. This shows that the trigger-strategy $t = 0$ is a best response, which characterizes an equilibrium.

PROPOSITION 1.2. *There is a unique symmetric equilibrium in pure trigger-strategies. In this equilibrium, each arbitrageur plays the trigger-strategy $t = 0$.*

It is easy to see that there is no other symmetric equilibrium in pure trigger-strategies. No such equilibrium exists for $t \geq \pi$ because the bubble never bursts after π. When everyone else sells at the same date $0 < t < \pi$, there is an upward jump in the probability distribution of the date of burst at t. An arbitrageur who deviates and sells a bit earlier sacrifices an infinitesimal reduction in the pre-crash price in return for a discrete decrease in the probability of burst. This shows that selling at the same date as others do cannot be a best response. The usual argument which rules out bubbles in models in which all traders are rational goes along this line.

Note that there are uncountably many equilibria in which the bubble bursts at $t = 0$. For example, there are uncountably many symmetric equilibria in which every arbitrageur plays a mixed trigger-strategy F which puts strictly positive mass on the trigger-strategy $t = 0$. From now on, we restrict ourselves to equilibria in which the bubble has some chance of survival, that is, to equilibria in which $G(0) < 1$.

We see that our model keeps the standard (EMH) solution: despite the potential speculative profits, there are equilibria in which competition among arbitrageurs causes an early burst and no-one benefits from them. As we have written down the model, the bubble is very weak in its infancy. The survival of the bubble requires coordination among arbitrageurs. They have the opportunity to feed the bubble, perhaps all the way to a mania, but the fear that other arbitrageurs may not concur can ruin it all. Their problem is that they cannot observe the realization of X; things would be much easier otherwise.

1.3.2. Non-degenerate mixed strategies. Do arbitrageurs have to content themselves with this solution? Our answer is no. It is possible to reconcile the individual incentive that arbitrageurs have to time the market with their collective interest in feeding the bubble. The ingredient that makes it possible is the possibility of occurrence of a mania. If a mania could not happen, an arbitrageur would stay in the market only if he believed that a sufficiently big mass of other

arbitrageurs will stay in as well. However, arbitrageurs cannot hold these beliefs because they are inconsistent: all arbitrageurs stay in the market for a finite period and there has to be one who stays in the longest. But if there is the possibility of a mania taking place, some arbitrageurs may find it optimal to stay in for long, knowing that they will succeed only if a mania develops, as long as they believe that the other arbitrageurs will not kill the bubble too soon. The reason for this is that arbitrageurs are no longer competitors during a mania; they can all quit the market at a profit.

But, how can those arbitrageurs be sure that the rest will not kill the bubble too soon? We show below how this coordination problem is resolved in equilibrium. Arbitrageurs will leave the market orderly, in such a way that their aggregate selling pressure does not exceed the absorption capacity too soon; they will adapt their behavior to the shape of κ. But this behavior must also be compatible with individual rationality to constitute an equilibrium; we prove that it is optimal for arbitrageurs to behave in a way that allows the outbreak of manias.

What is optimal for an arbitrageur is determined by the shape of the distribution function G. We know that for a mixed trigger-strategy F to characterize an equilibrium, all strategies in the interior of its support must be optimal. This imposes restrictions on G that translate into restrictions on F trough the function T. The following result states two properties of the function T that will be used to explain how the conditions of optimality restrict the shape of any equilibrium F.

LEMMA 1.3. *Suppose that there is a symmetric equilibrium in mixed trigger-strategies which fulfills $G(0) < 1$. Then, the function T is strictly increasing and continuous.*

The previous lemma tells us how we can obtain the distribution function of the date of burst from the distribution function of the state variable:

$$G(t) = \mathbf{P}(T(X) \le t) = \mathbf{P}(X \le T^{-1}(t)) = T^{-1}(t).$$

COROLLARY 1.4. *Given the conditions of Lemma 1.3, $G(t) = T^{-1}(t)$ for all $t \le \pi - \arcsin(\mu)$.*[1]

We see that restrictions on G translate directly into restrictions on T^{-1}. Before proceeding, let us label the infimum and the supremum of the support of F as $\underline{t}$ and $\overline{t}$ (Lemma A.2 in Appendix A.1 states that

---

[1]We restrict arcsin to its principal branch. $\pi - \arcsin(\mu)$ is the second time at which $\kappa(1, t)$ is equal to μ. Since κ is smaller than μ for all x and all $t > \pi - \arcsin(\mu)$, the bubble never bursts after this point in equilibrium.

the support of F is an interval). Because no arbitrageur sells before $\underline{t}$ in equilibrium, we have that for all states $x < \underline{t}/\pi$ the bubble bursts when the absorption capacity returns to zero (see figure 2). This means that the function T^{-1} is equal to t/π for all $t < \underline{t}$. Given the previous corollary, we can rewrite the pay-off function v as

$$v(t) = e^{(g-r)t} \left[1 - T^{-1}(t) \right] + T^{-1}(t) - c.$$

Let us label the equilibrium pay-off as v^*. Since the pay-off has to be maximum for all strategies inside the equilibrium support, we must have that

$$T^{-1}(t) = \frac{e^{(g-r)t} - v^* - c}{e^{(g-r)t} - 1}$$

for all $\underline{t} \le t \le \bar{t}$ (Lemma A.1 in Appendix A.1 states that this condition is fulfilled at the boundaries). Because T is strictly increasing and continuous by Lemma 1.3, we know that some price path must decrease at each $\bar{t} < t \le \pi - \arcsin(\mu)$. This implies that the sample path of κ which touches s at $\bar{t}$ cannot be strictly increasing at $\bar{t}$. Otherwise, because s is flat for $t \ge \bar{t}$, there would be some interval $(\bar{t}, \bar{t} + \epsilon]$ (with $\epsilon > 0$) in which no price path decreases, contradicting Lemma 1.3. Given the functional form of the absorption capacity, all subsequent paths must intersect s for the first time while they are decreasing. Hence, $T^{-1}(t)$ is the solution to $\kappa(x, t) = \mu$ for all $t > \bar{t}$. In short: given the extreme points of the equilibrium support, the function T^{-1} is determined. The following lemma shows that such extreme points are, in fact, unique.

LEMMA 1.5. *All F fulfilling the conditions of Lemma 1.3 have the same support.*

Lemma 1.5 shows that the function T is unique for the class of equilibria that we consider. Hence, to find an equilibrium within this class amounts to find a mixed trigger-strategy which, trough (1.2), induces such T. Our main result is that there exists a unique mixed trigger-strategy which does the job.

PROPOSITION 1.6. *There is a unique equilibrium fulfilling the conditions of Lemma 1.3. In this equilibrium, each arbitrageur plays the mixed trigger-strategy*

$$F(t) = \frac{1}{\mu} \kappa \left(T^{-1}(t), t \right)$$

for all $t \le \pi - \arcsin(\mu)$.

We have included a plot of an equilibrium s in figure 2 where we can see, graphically, how the different pieces fit together. The bubble

bursts at the point at which the realized sample path of κ crosses s. We see that some sample paths of κ cross s while they are increasing, whereas others cross it while they are decreasing. For those who cross it after they have reached their maximum we may say, informally, that the bubble is burst by behavioral traders. The key to understand how the equilibrium works is to note that all paths which correspond to states fulfilling $x > \mu$ belong to this set, that is, that all potential manias take place. Arbitrageurs start selling at $\underline{t}$, but their aggregate selling pressure accumulates slowly enough to feed the bubble, to let it grow.

Some arbitrageurs sell earlier than others. The ones who rush out of the market have more chances to sell at the pre-crash price and the ones who wait have the opportunity to earn higher profits, but they all expect the same pay-off ex-ante. The coordination achieved in equilibrium is remarkable, but far from perfect. Most sample paths of κ cross s before they reach their maximum, which means that arbitrageurs, as a group, could have done it better; it is not optimal to burst a bubble when there are still some behavioral traders willing to inject more money into the market. This is the curse of competition, the same that may induce an immediate collapse as the one described in Proposition 1.2.

1.4. Concluding Remarks

Our model is a caricature which reflects a particular view on how the strategic side of a bubble is perceived by real-world sophisticated investors. As any caricature, it gives a very simplified and biased account of the situation which, nonetheless, we believe, retains its very essence. Having said this, we want to add a remark about the way we interpret the equilibrium in Proposition 1.6. We have focused on symmetric equilibria though it is easy to see that there are uncountably many asymmetric equilibria which share the same s. The reason is that we find that an asymmetric equilibrium is an unnatural solution concept within our essentially symmetric context. We would find it hard to justify why otherwise identical arbitrageurs played different strategies in equilibrium, how did they know which strategy should they play, and so on. But symmetry cannot be more than a rough approximation to reality, however appealing and convenient. Asymmetries surely play their role in the workings of financial markets, though we do not bring them to the center of the discussion. We interpret our mixed strategy equilibrium from the Bayesian perspective, that is, from the view that it serves as an approximation to a more complex world in

which each arbitrageur harbors doubts about privately known characteristics of other arbitrageurs. Rather than to a classical randomizing interpretation of mixed strategies, we subscribe to the modern view in which our arbitrageurs would in fact be playing pure strategies, with mixed-strategies representing their uncertainties about others. Since every arbitrageur is negligible, we do not see why anyone would want to hide his choice.

Propositions 1.2 and 1.6 provide two legitimate solutions of the game. Proposition 1.6 hints on the logic of the persistence of bubbles and suggests how this phenomenon can be reconciled with a good deal of rationality in the market. The standard equilibrium in Proposition 1.2, we believe, loses strength when compared to it. The reason is that now it looks quite paranoid to quit the market right away on the fear that other arbitrageurs will do the same. Why don't you wait a while to confirm that your beliefs are right? You have nothing to lose. The standard equilibrium is quite appealing when it is unique, but not that much when there is another option available that everybody would prefer. It is quite natural to think that, after some particularly good news about the stocks, arbitrageurs will choose to wait for a while to see whether a bubble rises up.

CHAPTER 2

Currency Speculation in a Game-Theoretic Model of International Reserves

ABSTRACT. This paper is concerned with the ability of speculation to generate a currency crisis. We consider a game-theoretic setting between a unit mass of speculators and a government that holds foreign currency reserves. We analyze conditions under which the speculators may be able to force the government to devaluate the currency. Among these conditions, we analyze the role of heterogeneous beliefs, transaction costs, the level of international reserves, and the widening of currency bands. The explicit consideration of international reserves in our model makes speculators' actions to be strategic substitutes—rather than strategic complements. This is a main analytical departure with respect to related global games of currency speculation not including reserve holdings [e.g., Morris and Shin (1998)]. Our simple framework with international reserves becomes suitable to review some long-standing policy issues.

2.1. Introduction

2.1.1. Currency crises. Foreign currency reserves allow governments to follow exchange rate policies by intervention in the foreign exchange market. In a currency peg, these international reserves are used to absorb balance of payments deficits and to provide a cushion against other market forces. But currency speculation may also occur: If a mass of trades considers that the stock of international reserves is too low then they may rush to short the currency. The stock of reserves may be depleted—and the government is forced to leave the peg and float the currency. A currency attack can result in a sudden devaluation with severe negative effects on the financial and real sectors; these effects may stem from collateral requirements and other financial frictions, and price rigidities. A currency crisis may then emerge.

There are numerous examples of currency attacks, and there are long-standing issues regarding the optimal amount of transparency, transaction costs, and other regulations, to protect the value of a currency. Some of these issues became apparent in the last three most important currency crises.

Since its inception in 1979, the European exchange rate mechanism (ERM) experienced constant tensions that translated into a substantial number of currency realignments. After a swing of devaluations affecting some major currencies (e.g., the French franc, British pound, and Italian lira) the ERM essentially collapsed in 1993 as it moved to a much broader currency band. Then, currency values stabilized. Most models of exchange rate determination are not suited to assess the influence of currency bands on currency speculation. We shall study a simple extension of our model in which it becomes harder to attack the currency under a broader currency band.

The 1994 currency crisis of the Mexican peso brought up some transparency issues. For instance, in several papers Calvo [e.g., see Calvo (1998)] argued that with uncertainty on the fundamentals, economic crises may spread by contagion and herding behavior. The International Monetary Fund (IMF) has set up the Special Data Dissemination Standards (SDDS) for all member countries. Disclosure practices of foreign currency reserves and other macro variables have varied over time and across member countries, but it is often argued that it is desirable to adhere to the highest possible standards of transparency (see *op. cit.*). Global games with heterogeneous beliefs provide a natural setting to deal with uncertainty on the fundamentals (Carlsson and Damme, 1993). Equilibrium effects from changes in heterogeneity of beliefs have already been addressed in Morris and Shin (1998) and several other papers. Our interest is to see how these conclusions from the existing literature may survive in our setting with explicit consideration of international reserves.

In the Asian crisis that started in 1997, the Thai government spent billions of dollars of its foreign currency reserves to defend its baht against speculative attacks. The lack of timely response by the IMF and other institutions such as the US Fed was blamed to be a hitch at the onset of the crisis. In this paper we address how the borrowing of international reserves could be effective to deter currency speculation. Indeed, the effectiveness of these interventions is going to depend on the degree of coordination of the speculators: If there is common knowledge about the fundamentals then we know from some simple models of exchange rate determination that over a certain region external aid from other central banks may be quite inoperative (e.g., see Figure 1 below). These models generate multiple equilibria over a wide range of parameter values and hence they lack predictive power. Transaction costs, taxes on capital gains, and the size of currency depreciation have actually no apparent effect. Hence, an obvious policy prescription of some of these models with multiple equilibria would be to shut

down international capital markets. The current Greek debt crisis is another case in point. Massive coordination by European countries has proved to lower risk premiums—albeit the reprieve may only last for a few years. The effectiveness of these coordination efforts is not clear since in some cases the whole private sector could short more currency than global entities can ever supply. Hence, the question is whether or not external borrowing of international reserves becomes more effective under asymmetric information to prevent a currency crisis.

2.1.2. Self-fulfilling currency attacks. In a fixed exchange rate, the government bears the risk of a speculative attack as it is willing to exchange the currency at a predetermined price. Although there are associated benefits of fixing the value of a currency, the costs could be prohibitive. For speculators it is of paramount importance that the government wants to resist the attack; it is precisely this foreseen resistance what motivates their actions in the first place. That is, speculators would like to short the domestic currency at the pegged price—and later undo the trading at a lower equilibrium price.

Currency attacks may be self-fulfilling. The mere belief on an imminent attack may induce speculators to flee from the currency. A frenzied rush of capital outflows is then vindicated by a devaluation that confirms the initial beliefs. The point has been neatly discussed in Obstfeld (1996). Obstfeld proposed a game in which two private holders of domestic currency must decide whether to sell or to hold the currency. The government owns reserves to sustain the peg, yet a 50% devaluation sets off if reserves are depleted. Let us assume the following conditions: (a) The government owns 10 units of reserves, (b) The pegged rate is 1-1, (c) Each holder has 6 units of currency, and (d) Each holder bears a cost of 1 upon selling. Then, we get the following pay-off matrix:

	Hold	*Sell*
Hold	$0,0$	$0,-1$
Sell	$-1,0$	$3/2,3/2$

FIGURE 1. The intermediate reserve game in Obstfeld (1996).

Note that none of the two holders can break the peg unilaterally. Hence, an individual holder alone cannot recover the transaction cost. But if both traders sell, then there is a capital gain from the 50% devaluation which outweighs the transaction cost. Therefore, this game has two pure-strategy equilibria; one in which both holders sell and another one in which no holder sells. The players' actions are *strategic*

complements because selling is profitable only if the other holder sells. But this game also reflects the idea that the total gains from speculation depend on the amount of reserves released for sale by the government; or put it somewhat differently, in the (*sell, sell*)-equilibrium a trader would be better off if the other holder had only 4 units of currency.

As pointed out by Morris and Shin (1998), a main problem with various models with multiple equilibria is that the immediate reasons behind the actual onset of an attack are left unexplained.

2.1.3. Our results. At this venture it may be helpful to provide a cursory review of our results with those of Figure 1. In this figure there are two equilibria: (*hold, hold*) and (*sell, sell*). Many models of currency speculation have emphasized the existence of multiple equilibria, and the need of coordination devices over those equilibria. Coordination may actually come in the form of sound economic policies [cf. Kaminsky et al. (1998)] that direct traders to non-speculative equilibria. We must note, however, that in a corresponding extensive form representation of the game the (*sell, sell*) equilibrium is sub-game perfect. Hence, this is a focal point of the game: For low transaction costs any sensible perturbation of the game will point towards this equilibrium. This simple observation seems to be absent in the so called second-generation models of currency crises that simply stress multiplicity of equilibria without regard to further properties of these equilibria. But the problem with the simple game of Figure 1 is that (*sell, sell*) remains the preferred equilibrium outcome regardless of the size of transactions costs and the benefits from speculation. This is a really odd result that comes from perfect information. The only prescriptions of this game is for the government to meet the amount of reserves, or to shut down the economy from international capital flows. It should nevertheless be pointed out that the (*sell, sell*) equilibrium pinpoints the inherent instability of fixed exchange rate regimes, since traders are motivated by the gains of speculation.

The above game allows each agent to be placed in the position of the other player: There is common knowledge about the fundamentals. Asymmetric information will certainly change the picture. Players with diffused information about the stock of reserves may not be so sure about shorting the currency and bear the transaction cost if there is a certain probability that other players may not move to short the currency. That is, each player has to guess the beliefs of other players, and everyone will be guessing about others' guesses, and so on. This is a complex topic that leads us to the literature of global games.

2.1.4. The model of Morris and Shin. These authors propose a two-stage game played by the government and a continuum of speculators. In the first stage, each speculator decides whether or not to sell short one unit of the domestic currency at a certain cost $t > 0$, whereas in the second stage the government decides whether or not to defend the peg e^*. If the government defends, the price stays at the original level e^* and the speculators who attack earn nothing and pay the cost of short-selling. If the government does not defend and floats the currency, the price falls to $f(\theta)$, where f is increasing in the state θ of the fundamentals, and the speculators exchanging the currency earn the price difference minus the cost: $e^* - f(\theta) - t$. The government's pay-off upon defending is written as,

$$v - c(\alpha, \theta).$$

This value increases with the state θ of the fundamentals and goes down with the mass α of speculators who attack.

As in the example above, Morris and Shin show that under common knowledge there is a wide range of θ in which the game has two equilibria; one in which no speculator attacks and the government maintains the peg, and another one in which all speculators attack and the government accommodates. Morris and Shin show that this multiplicity of equilibria is not robust: Asymmetric information about the state of the fundamentals induces a unique equilibrium.

Conditioning upon the government not defending, the pay-off that a speculator receives from attacking is independent of the mass of speculators who attack. The actions of the speculators are thus strategic complements because the chances of a devaluation increase with the mass of speculators who attack. Therefore, Morris and Shin assume that *all* speculators can sell the domestic currency at the pegged price if the government does not defend. In this model it is not really clear who buys the domestic currency from speculators since the pay-off of each speculator is independent of the mass of speculators who attack. Further, the domestic currency always depreciates if the government does not defend, which implies that there must be an excess supply of the domestic currency at the pegged price—even if no speculator attacks.

In our model below, speculators' actions are *strategic substitutes*: When the peg is abandoned the gains from trading that accrue to each speculator are inversely related to the mass of speculators shorting the currency. As is well known from the global games literature, strategic substitutability may generate additional technical problems for existence and uniqueness of equilibria.

2.1.5. Other related work.

2.1.5.1. *Bank runs.* Banks play an important role as providers of liquidity insurance. Demand-deposit contracts pool idiosyncratic risks to finance more attractive long-term investments. But the early interruption of long-term investments typically entails a loss. If idiosyncratic liquidity shocks are sufficiently uncommon and independent across the population, banks can improve upon the autarkic allocation.

Hence, banks are also vulnerable to runs that may cause them to vanish. As in the case of currency attacks, the fear of an imminent run may propel massive withdrawals—vindicating the initial beliefs. As is well known, Diamond and Dybvig (1983) provide a model of demand-deposit contracts in which there are two equilibria; an efficient equilibrium in which only investors facing liquidity shocks withdraw early, and a bank run equilibrium in which all investors withdraw and the bank fails.

The model of Diamond and Dybvig is subject to the same criticisms as models of currency attacks with multiple equilibria. Goldstein and Pauzner proposed a model with asymmetric information à la Morris and Shin and show that the multiplicity of equilibria in Diamond and Dybvig (1983) washes out. In the model of Goldstein and Pauzner (2005) the actions of the depositors are not strategic complements everywhere. This is because conditioning on the bank failing, as more depositors withdraw their funds, the lower is their share on the bank's liquidation value. There are, however, $one-sided$ $strategic$ $complementarities$, because if the bank survives then early withdrawals reduce the pay-offs to the depositors that stay with the bank. Goldstein and Pauzner build their proof of uniqueness upon this property of the pay-offs.

2.1.5.2. *Bubbles.* Our work is also related to the theory of bubbles in behavioral finance. In Abreu and Brunnermeier (2003) a continuum of speculators must decide at each instant whether or not to sell an overpriced stock. They face a mass of behavioral traders who are responsible for the abnormal price growth. It is assumed that the price of the stock will continue to grow at the bubbly rate as long as the mass of speculators who sell remains below the mass of behavioral traders (who buy); once trading surpasses this threshold, the bubble bursts out immediately.

Although the problem is framed in a richer, dynamic setting, the similarities between their problem and ours are evident. A speculator should sell immediately in the belief of imminent collapse and wait otherwise; moreover, the belief of a sudden crash is self-confirming. However, the technical approach of Abreu and Brunnermeier to this

problem was very different. They did not follow the line of the global games literature because, in their own words: "In the richer strategy set of our model strategic complementarity is not satisfied and the global games approach does not apply" (Abreu and Brunnermeier, 2003, page 177).

In their model, speculators have an incentive to preempt others because the pay-off from selling at the date of bursting of the bubble is decreasing in the mass of speculators who sell. The pay-offs exhibit strategic substitutability at the date of bursting because then the amount of speculators who sell outweighs the amount of behavioral traders who buy and the market must clear—behavioral traders play here the role played by the government in a model of currency crisis. Given the similarities between the two problems, we expect that our work will serve as a first step towards the incorporation of global games to the theory of bubbles.

2.2. The Model

The state of the world is given by the amount R of international reserves that the government has ready to defend the peg. The government operates here as a passive player who buys the domestic currency until it runs out of reserves. The amount of reserves may be interpreted as the government's degree of commitment to the exchange rate defense rather than as an exogenous limit (as in Obstfeld, 1996). That is, R must be thought of as the outcome of a previous, yet not modeled, deliberation by the government; e.g., its ability to draw in funds from international capital markets. This information is usually hard to guess by both the government and the traders as it may depend on unexpected external forces.

Intervention is necessary because the government's desired exchange rate, the pegged rate, differs from the equilibrium rate: There is an excess supply of s_e units of the domestic currency which would cause a devaluation if the government did not intervene.

We consider a simultaneous-move game played by a continuum of speculators of unit mass. Each speculator can short one unit of the domestic currency at a cost $c > 0$. If the mass of speculators who short the currency, s, plus the excess supply, s_e, exceed the government's reserves, $s + s_e > R$, the domestic currency depreciates by a fraction $\delta \in (0, 1)$. Otherwise, the peg survives.

In a devaluation, the total amount of reserves is shared among those who sell (or, equivalently, they all have equal chances to sell before the devaluation). Therefore, the pay-off to a speculator who attacks is $-c$

if the peg survives, and

$$\frac{R}{s + s_e} \delta - c$$

if it does not survive. A speculator who does not attack gets zero in any case.

In summary, the gains from speculation stem from selling short the domestic currency at the pegged rate and then purchasing back the same currency at the ensuing equilibrium rate after the depreciation. The overall gains from speculation are the total reserves times the rate of depreciation. The government will try to sustain the peg, but the amount of reserves is limited. We assume that these reserves are equally shared by all traders executing the transaction. Then, the actions of the speculators are strategic substitutes if $s + s_e > R$: Conditioning on the peg being abandoned, the pay-off to a speculator from attacking decreases with the mass of speculators who attack.

2.2.1. Perfect information. Let us begin with the simple case in which the amount of reserves held by the government is common knowledge among speculators. Depending on the size of R we can identify three different types of games (as in Obstfeld, 1996):

- If $R < s_e$ we are in a low reserve game. The government does not have enough reserves to defend the peg even if no speculator attacks. Therefore, a devaluation will come for sure. Assuming that the cost associated to short selling is sufficiently small;

(2.1) $$c < \frac{R}{1 + s_e} \delta,$$

  we can ensure that attacking is a dominant strategy whenever $R < s_e$. In this case, there is a unique equilibrium in which all speculators attack.

- If $s_e \leq R < 1 + s_e$ we are in an intermediate reserve game. The peg will be abandoned depending on the mass of speculators who attack; $s \in [0, 1]$. Attacking is the optimal choice for all speculators who believe that $s + s_e > R$, and not attacking is the optimal choice for those who believe that $s + s_e \leq R$. Moreover, both beliefs are self-confirming because they end up being right if they are held equally across the population of speculators. There are, thus, two equilibria in pure strategies within this range of reserves: One in which all speculators attack and the peg is abandoned, and another one in which no speculator attacks and the peg survives.

- If $R \geq 1 + s_e$ we are in a high reserve game. Here the government has enough reserves to defend the peg even if all speculators attack. The peg will thus survive, and so attacking becomes a strictly dominated strategy. There is a unique equilibrium in which no speculator attacks.

2.2.2. Imperfectly observed reserves. Let us now assume that speculators do not observe R directly, but hold certain beliefs. Suppose that each speculator holds a uniform prior over the interval $[\underline{R}, \overline{R}]$, where $R = \underline{R}$ fulfills (2.1) and $\overline{R} > 1 + s_e$. Each speculator receives a conditionally independent signal x which is also distributed uniformly over the interval $[R - \varepsilon, R + \varepsilon]$ (with $\varepsilon > 0$).[1] Under these assumptions, the posterior belief about R of a speculator who receives the signal x is uniform over the interval $[x - \varepsilon, x + \varepsilon]$.

Note that under this specification parameter ε is both a measure of the precision of each signal and the degree of informational asymmetry among speculators since signals are conditionally independent; varying the degree of dependence between the signals would allow us to disentangle both features. More importantly, it is crucial to realize that only the event $[\underline{R}, \overline{R}]$ is common knowledge among speculators, no matter how small ε might be. Note that an event $E \subset [\underline{R}, \overline{R}]$ is nth-order mutual knowledge at $R \in E$ only if $E \supseteq [R - 2n\varepsilon, R + 2n\varepsilon] \cap [\underline{R}, \overline{R}]$, which means that there is always some n for which the last inclusion fails to hold.

As shown below, small departures from common knowledge lead to very different results. Indeed, the two pure strategy equilibria of the intermediate reserve game in the previous subsection require a high degree of coordination among speculators. A speculator must predict the behavior of speculators who receive signals which are an ε away from this speculator, which in turn depends on their beliefs about the behavior of speculators who are an ε away from them, and so on. This is how a small seed of noise infects the whole range of signals.

A strategy for a speculator is now a function from the set of signals to the set of actions. Let $\pi(x)$ denote the proportion of speculators who attack from those who have received the signal x. Adding up across signals, the aggregate short sales under the stock of reserves R we get:

$$s(R, \pi) = \frac{1}{2\varepsilon} \int_{R-\varepsilon}^{R+\varepsilon} \pi(x) \, dx.$$

---

[1]The limits of this interval should obviously be adjusted if $x < \underline{R} + \varepsilon$ or $x > \overline{R} - \varepsilon$.

Given π, the peg is abandoned in the event:

$$A(\pi) := \{R : s(R, \pi) + s_e > R\}.$$

And the expected pay-off from attacking to a speculator who receives the signal x must be:

$$(2.2) \qquad u(x, \pi) = \frac{1}{2\varepsilon} \int_{A(\pi) \cap [x-\varepsilon, x+\varepsilon]} \frac{R}{s(R, \pi) + s_e} \delta \, dR - c.$$

An equilibrium of the game occurs if $\pi(x) = 1$ whenever $u(x, \pi) > 0$, and $\pi(x) = 0$ whenever $u(x, \pi) < 0$.

2.3. Results

2.3.1. Threshold equilibrium.
A threshold equilibrium is an equilibrium in which there is a R^* such that: (a) The peg is abandoned for all $R < R^*$ and (b) The peg survives for all $R \geq R^*$. We will see below that functions π and s have both a particularly simple form in a threshold equilibrium. This will be of great help in order to show that there is a unique threshold equilibrium.

Suppose that R^* characterizes a threshold equilibrium. For every signal $x \leq R^* - \varepsilon$ we have that $u(x, \pi) > 0$ because speculator x believes that the peg will be abandoned with probability one. For every signal $x \geq R^* + \varepsilon$ we have that $u(x, \pi) = -c$ because speculator x believes that the peg will survive with probability one. Moreover, $u(x, \pi)$ is strictly decreasing in x in the interval $(R^* - \varepsilon, R^* + \varepsilon)$ since as we move to the right within this interval, the integral in (2.2) adds up states in which the pay-off is $-c$ and leaves off states in which it is positive. By the continuity of the integral, there is a unique x^* fulfilling $u(x^*, \pi) = 0$. Therefore, we have shown that, in any threshold equilibrium, π must have the form:[2]

$$(2.3) \qquad I_x(z) = \begin{cases} 1 & \text{if } z \leq x \\ 0 & \text{if } z > x. \end{cases}$$

If $\pi = I_x$, we know that $s(R, \pi)$ is equal to one if $R \leq x - \varepsilon$ and equal to zero if $R > x + \varepsilon$; moreover, it decreases at the rate of $1/2\varepsilon$ between these two points. In short:

$$s(R, I_x) = \begin{cases} 1 & \text{if } R \leq x - \varepsilon \\ \frac{1}{2} - \frac{1}{2\varepsilon}(R - x) & \text{if } x - \varepsilon < R \leq x + \varepsilon \\ 0 & \text{if } R > x + \varepsilon. \end{cases}$$

---

[2]The value of $I_x(x)$ can be chosen arbitrarily from [0,1].

Consequently, event $A(\pi)$ becomes $A(I_x) = [\underline{R}, \rho(x))$, where

$$\rho(x) = \frac{1}{1 + 2\varepsilon}[x + (1 + 2s_e)\varepsilon].$$

Every x fulfilling $u(x, I_x) = 0$ characterizes an equilibrium. We show now that there is exactly one such x. Considering $u(x, I_x)$ as a function of x alone, we see that if ε is not too big[3] this function is positive at the lower end of the set of signals and negative at the upper end. As one moves to the right, two opposite effects are in action: (i) More speculators are required to cause a devaluation; and (ii) The individual benefit from shorting the currency goes down with the mass of speculators attacking the currency. The first effect is not present in models of bank runs (Goldstein and Pauzner, 2005), the second is not present in models with global strategic complementarities (Morris and Shin, 1998). The expression for $u(x, I_x)$ is

$$u(x, I_x) = \delta \left\{ (1 + 2\varepsilon)\rho(x) \log\left(\frac{1 + s_e}{\rho(x)}\right) - [\rho(x) - (x - \varepsilon)] \right\} - c.$$

Taking its second derivative with respect to x,

$$\frac{\partial^2 u(x, I_x)}{\partial x^2} = -\frac{\delta}{(1 + 2\varepsilon)\rho(x)},$$

we see that $u(x, I_x)$ is strictly concave, which, in turn, implies that there is a unique x^* fulfilling $u(x^*, I_{x^*}) = 0$. We have just proved the following proposition.

PROPOSITION 2.1. *There is a unique threshold equilibrium. In this equilibrium there is a signal x^* such that: (a) All speculators who receive a signal $x < x^*$ attack, (b) All speculators who receive a signal $x > x^*$ do not attack, (c) The peg is abandoned for all $R < \rho(x^*)$ and (d) The peg survives for all $R \geq \rho(x^*)$.*

REMARK 2.2. Strictly speaking, there is a continuum of threshold equilibria which only differ in a set of measure zero (at x^*).

2.3.2. Iterated deletion of dominated strategies. A remarkable property of the model of Morris and Shin is that the equilibrium strategies are the only ones surviving iterated deletion of strictly dominated strategies. This property is a direct consequence of global strategic complementarities; we now show that it does not hold for more general pay-off structures.

Attacking is a dominant action for all speculators who receive a signal below $s_e - \varepsilon$ because they believe that the peg will be abandoned

---

[3]A sufficient condition is $2\varepsilon < \min\{s_e - \underline{R}, \overline{R} - (1 + s_e)\}$.

for sure. This fact has an effect on the behavior of speculators who receive signals above $s_e - \varepsilon$, since they now know that $\pi(x) = 1$ for all $x < s_e - \varepsilon$. That is, some speculators to the right of $s_e - \varepsilon$ may find that because all speculators below $s_e - \varepsilon$ attack then this is a sufficient condition for them to attack as well. More generally, we are interested in the lowest pay-off that speculator x can expect from attacking, provided that $\pi(z) = 1$ for all $z < x$. If such expected pay-off is positive, we know that speculator x will attack.

PROPOSITION 2.3. *Not attacking does not survive the iterated deletion of strictly dominated strategies for all signals below x^*.*

PROOF. The proof proceeds in three steps:

Step 1: We first show that the expected pay-off for speculator x, provided that $\pi(z) = 1$ for all $z < x$, is bounded below by the one derived from some threshold function I_{x_0}. If $\pi(z) = 1$ for all $z < x$, we know that $s(R, \pi)$ is weakly decreasing in the interval $(x - \varepsilon, x + \varepsilon)$. Since we are looking for the minimum expected pay-off, it has to be the case that $s(R_0, \pi) + s_e = R_0$ for some R_0 in $(x - \varepsilon, x + \varepsilon)$. Now, choose the x_0 that makes $s(R_0, I_{x_0}) + s_e = R_0$. We must have that $u(x, I_{x_0}) \leq u(x, \pi)$ since I_{x_0} lies above π on $(x - \varepsilon, R_0)$.

Step 2: The next step is to show that $u(x, I_x) \leq u(x, I_{x_0})$. The first integrates over the interval $[x - \varepsilon, \rho(x)]$, whereas the second integrates over $[x - \varepsilon, \rho(x_0)]$, which is larger. What we will do is to compare, moving to the left from the right limit of each interval, the pay-offs at each state. The pay-offs for I_x can be written as:

$$(2.4) \qquad \frac{R}{\rho(x) + \frac{1}{2\varepsilon}[\rho(x) - R]}\delta - c.$$

Pairing each state in $[x - \varepsilon, \rho(x)]$ with the corresponding state in $[x - \varepsilon, \rho(x_0)]$ (recall that we are moving to the left from the right end of each interval), the pay-offs for I_{x_0} are:

$$(2.5) \qquad \frac{R + \rho(x_0) - \rho(x)}{\rho(x_0) + \frac{1}{2\varepsilon}[\rho(x) - R]}\delta - c$$

if $R > x_0 - \varepsilon + \rho(x) - \rho(x_0)$, and

$$(2.6) \qquad \frac{R + \rho(x_0) - \rho(x)}{1 + s_e}\delta - c$$

otherwise. Subtracting (2.4) from (2.5) gives us

$$\frac{2\varepsilon(1 + 2\varepsilon)[\rho(x) - R][\rho(x_0) - \rho(x)]}{[(1 + 2\varepsilon)\rho(x) - R][\rho(x) + 2\varepsilon\rho(x_0) - R]}\delta \geq 0.$$

Since (2.6) is greater than (2.5), the proof is complete.

Step 3: We have just shown that the least that speculator x can get from attacking, provided that $\pi(z) = 1$ for all $z < x$, is $u(x, I_x)$. Since $u(x, I_x) > 0$ for all $x < x^*$, this implies our result. $\qquad\square$

In the previous proof we have shown that, if $\pi(z) = 1$ for all $z < x$, the case in which speculator x gets the least from attacking is when $\pi(z) = 0$ for all $z > x$. If we had global strategic complementarities, the converse would be true when $\pi(z) = 0$ for all $z > x$. Then, it would be immediate to see that the threshold equilibrium is the unique equilibrium. This is not our case, however. If we look now for the most that speculator x can get from attacking if $\pi(z) = 0$ for all $x > z$, we see that the answer is not given by the threshold function I_x. The reason is that, conditioning on the peg being abandoned, the more speculators who attack, the lower is the pay-off that they get from attacking.

PROPOSITION 2.4. *There is a minimal $x^\diamond$, with $x^\diamond > x^*$, such that attacking does not survive the iterated deletion of strictly dominated strategies for all signals above $x^\diamond$.*

PROOF. The proof proceeds in three steps:
Step 1: We first construct an upper bound for the expected pay-off of speculator x whenever $\pi(z) = 0$ for all $z > x$. If $\pi(z) = 0$ for all $z > x$, then $s(R, \pi)$ is weakly decreasing in $(x - \varepsilon, x + \varepsilon)$. Since we are looking for the maximum expected pay-off, it has to be the case that $s(R_0, \pi) + s_e = R_0$ for some R_0 in $(x - \varepsilon, x + \varepsilon)$, and also that $s(R, \pi)$ is constant in $(x - \varepsilon, R_0)$. Conditioning on $s(R, \pi) + s_e = R_0$ in $(x - \varepsilon, R_0)$, the expected pay-off at x is

$$(2.7) \qquad \frac{1}{2\varepsilon} \int_{x-\varepsilon}^{R_0} \frac{R}{R_0} \, \delta \, dR - c,$$

which is strictly increasing in R_0. Then, the maximum expected pay-off is attained at the maximum R_0, namely, $\rho(x)$ (recall that $\pi(z) = 0$ for all $z > x$). Hence, any π fulfilling $\pi(z) = 0$ for $z \in (x - 2\varepsilon, \rho(x) - \varepsilon)$, and $\pi(z) = 1$ for $z \in (\rho(x) - \varepsilon, x)$, attains the maximum expected pay-off.
Step 2: Substituting $\rho(x)$ for R_0 in (2.7), we see that the maximum expected pay-off is positive if x is sufficiently small and negative if it is sufficiently large. We can compute its second derivative with respect to x,

$$\frac{-2(1 + s_e)^2 \varepsilon (1 + 2\varepsilon)}{[x + (1 + 2s_e)\varepsilon]^3} \, \delta,$$

and see that it is strictly concave. This means that there is exactly one $x^\diamond$ at which the maximum expected payoff is zero, that it is positive for

$x < x^\diamond$, and that it becomes negative for $x > x^\diamond$. Therefore, attacking does not survive the iterated deletion of strictly dominated strategies if, and only if, $x > x^\diamond$.

Step 3: That $x^\diamond > x^*$ comes directly from the fact that the maximum expected pay-off constructed in Step 1 for each x is always larger than $u(x, I_x)$.   □

2.3.3. Uniqueness of the threshold equilibrium. Our main result concerns the uniqueness of the threshold equilibrium.

PROPOSITION 2.5. *The equilibrium in Proposition 2.1 is the only equilibrium.*

PROOF. The proof is by contradiction. We suppose that there exists an equilibrium which is not a threshold equilibrium and then show that this is impossible. Suppose that π characterizes an equilibrium. Define $\overline{x}$ as

$$\overline{x} := \sup\{x | \pi(x) > 0\}$$

and $\underline{x}$ as

$$\underline{x} := \begin{cases} \overline{x} & \text{if } \pi(x) = 1 \text{ for all } x < \overline{x} \\ \sup\{x < \overline{x} | \pi(x) < 1\} & \text{otherwise.} \end{cases}$$

Note that, if the equilibrium is not a threshold equilibrium, we must have that $\underline{x} < \overline{x}$. Also, by continuity, we must have that $u(\underline{x}, \pi) = u(\overline{x}, \pi) = 0$. We presently show that this is impossible.

Step 1: The peg survives for all $R > \rho(\overline{x})$. We know that $s(R, \pi)$ must be weakly decreasing in $(\overline{x} - \varepsilon, \overline{x} + \varepsilon)$. Since $u(\overline{x}, \pi) = 0$, there must be a R_0 in $(\overline{x} - \varepsilon, \overline{x} + \varepsilon)$ at which $s(R_0, \pi) + s_e = R_0$. The expected pay-off $u(x, \pi)$ is strictly positive within the interval $[R_0 - \varepsilon, \overline{x})$ since, as we move to the left from its right end, we are excluding states in which the peg survives (and adding some in which it is abandoned). Therefore, $\pi(x) = 1$ for all x in $[R_0 - \varepsilon, \overline{x})$ which, in turn, implies that $R_0 = \rho(\overline{x})$. Furthermore, for the same reason, $s(R, \pi)$ must start decreasing at the fastest rate before $\rho(\overline{x})$.

Step 2: If $\underline{x} < \overline{x} - 2\varepsilon$, then $\underline{x} = \overline{x}$. If $\underline{x} < \overline{x} - 2\varepsilon$, we have that $u(\overline{x}, \pi) = u(\overline{x}, I_{\overline{x}})$, which is zero only if $\overline{x} = x^*$ (Proposition 1). But, then, what we have is a threshold equilibrium.

Step 3: If $\underline{x} \geq \overline{x} - 2\varepsilon$, then $u(\underline{x}, \pi) > u(\overline{x}, \pi)$. In order to compute $u(\underline{x}, \pi)$ and $u(\overline{x}, \pi)$ we must integrate over the intervals $[\underline{x} - \varepsilon, \underline{x} + \varepsilon]$ and $[\overline{x} - \varepsilon, \overline{x} + \varepsilon]$. Both intervals overlap and, therefore, we only need to compare the pay-off accumulated at both sides of the common subinterval.

Step 3.1: To accomplish this task, we first find a lower bound to the pay-off accumulated on the left-hand side subinterval. From (2.2) we know that such a lower bound can be obtained by (a) reducing the size of the set $A(\pi)$, and (b) substituting the denominator $s(R, \pi) + s_e$ by a larger quantity at each state R. Let $s_0 = s(\overline{x} - \varepsilon, \pi)$ and let $R_0 = \max\{R < \overline{x} - \varepsilon | s(R, \pi) + s_e = R\}$. First, we reduce the set $A(\pi)$ by assuming that the peg survives for all states in $[\underline{x} - \varepsilon, R_0)$. Second, we find an upper bound for the denominator $s(R, \pi) + s_e$ within $[R_0, \overline{x} - \varepsilon]$. We know that $s(R, \pi)$ is weakly increasing in $(\underline{x} - \varepsilon, \overline{x} - \varepsilon)$, which means that $s(R, \pi) \leq s_0$ within this interval. The pay-off accumulated in $[\underline{x} - \varepsilon, \overline{x} - \varepsilon]$ is bounded below by

$$(2.8) \qquad \frac{1}{2\varepsilon} \left[\int_{R_0}^{R_1} \frac{R}{R_0 + \frac{1}{2\varepsilon}(R - R_0)} \delta \, dR + \int_{R_1}^{\overline{x} - \varepsilon} \frac{R}{s_0} \delta \, dR - c \right],$$

where $R_1 = \min\{R_0 + 2\varepsilon(s_0 - R_0), \overline{x} - \varepsilon\}$. The denominator inside the first integral corresponds to an upward sloping line, starting at the point R_0 on the 45-degree line and growing at the maximum feasible rate until the upper limit s_0 is reached; in the second integral the curve becomes flat. Thus, we have constructed the largest admissible denominators given the constraints: $s(R_0, \pi) + s_e = R_0$ and $s(R, \pi) \leq s_0$ in $[\underline{x} - \varepsilon, \overline{x} - \varepsilon]$.

We now show that (2.8) is decreasing in R_0. A sufficient condition for (2.8) to be decreasing in R_0 is that the first integral is so when $R_1 = R_0 + 2\varepsilon(s_0 - R_0)$. The derivative of the first summand in (2.8) in this case is negative if

$$\log\left(\frac{s_0}{R_0}\right) \leq \frac{1}{1 - 2\varepsilon}.$$

Since $s_0 < 1 + s_e$ and $R_0 > \rho(x^*)$, a sufficient condition for this to be true is that

$$(2.9) \qquad \rho(x^*) \geq \frac{1 + s_e}{e},$$

where $e = 2.7182\ldots$ On the other hand, we know that the derivative of $u(x, I_x)$ with respect to x,

$$\frac{\partial u(x, I_x)}{\partial x} = \delta \left[\log\left(\frac{1 + s_e}{\rho(x)}\right) - \frac{1}{1 + 2\varepsilon} \right],$$

has to be negative at $x = x^*$. Since this implies (2.9), we have shown that (2.8) decreases with R_0. Therefore, substituting R_0 by a larger number in (2.8) gives us a lower bound for the pay-off accumulated on the left-hand side subinterval.

Step 3.2: Let $\varrho(x)$ be the point at which $s(R, 1 - I_x) + s_e = R$, i.e.:

$$\varrho(x) = \frac{1}{1 - 2\varepsilon}[x - (1 + 2s_e)\varepsilon].$$

We know that $\varrho(\underline{x}) \geq R_0$ because $s_0 \geq \rho(\overline{x})$. Then, the pay-off accumulated on the left-hand side is bounded below by

(2.10) $$\frac{1}{2\varepsilon}\left[\int_{\varrho(\underline{x})}^{\overline{x}-\varepsilon} \frac{\varrho(\underline{x})}{\frac{1}{2} + \frac{1}{2\varepsilon}(R - \underline{x}) + s_e} \delta \, dR - c\right].$$

On the other hand, the pay-off accumulated on the right-hand side is bounded above by

(2.11) $$\frac{1}{2\varepsilon}\left[\int_{\underline{x}+\varepsilon}^{\rho(\overline{x})} \frac{\rho(\overline{x})}{\frac{1}{2} - \frac{1}{2\varepsilon}(R - \overline{x}) + s_e} \delta \, dR - c\right].$$

We have that (2.10) is larger than (2.11) if

(2.12) $$\varrho(\underline{x}) \log\left(\frac{\frac{\overline{x}-\underline{x}}{2\varepsilon} + s_e}{\varrho(\underline{x})}\right) \geq \rho(\overline{x}) \log\left(\frac{\frac{\overline{x}-\underline{x}}{2\varepsilon} + s_e}{\rho(\overline{x})}\right).$$

But the function

$$x \log\left(\frac{a}{x}\right)$$

is decreasing if

$$x \geq \frac{a}{e}.$$

In our case

$$\varrho(\underline{x}) \geq \frac{\frac{\overline{x}-\underline{x}}{2\varepsilon} + s_e}{e},$$

and so

$$\varrho(\underline{x}) \geq \frac{1 + s_e}{e}$$

suffices for (2.12) to be true. Combining $\varrho(\underline{x}) > \rho(x^*)$ and (2.9) we have that $u(\underline{x}, \pi) > u(\overline{x}, \pi)$. $\qquad\square$

From these results we can now show:

PROPOSITION 2.6. *In the limit, as ε goes to zero, x^* is obtained as the solution of the following equation:*

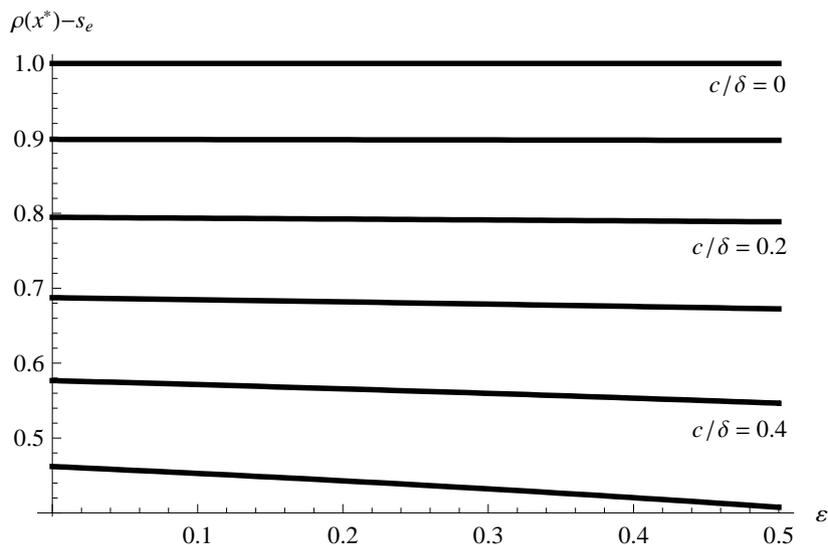$$x \log\left(\frac{1 + s_e}{x}\right) = \frac{c}{\delta}.$$

FIGURE 2. Various plots of $\rho(x^*) - s_e$ as a function of ε for $s_e = 3$ and $c/\delta < 1/2$.

2.3.4. Comparative statics. The previous proposition becomes fundamental to perform comparative static exercises. As suggested in the introduction, most important results refer to variations in the size ε of the noise term, which can be interpreted as a measure of the lack of transparency of the monetary policy. In particular, we shall be interested in the behavior of the quantity $\rho(x^*) - s_e$, which gives the proportion of states in which the peg is abandoned, as the noise ε becomes small. Our results show that, in general, this quantity gets bigger as the size of the noise decreases. That is, an increase in the transparency of the monetary policy tends to enlarge the set of states in which currency attacks succeed. In any case, this effect is of quantitative little importance for the majority of parameter values (in fact, it is almost zero for the most interesting cases).

Figure 2 presents several plots of $\rho(x^*) - s_e$ as a function of ε for different values of the ratio c/δ and $s_e = 3$. We have chosen the maximum value of ε to be 0.5 because this is the value that makes speculator $s_e + 1/2$ believe that all states in the intermediate reserve region are possible. The values for the ratio c/δ are all smaller that $1/2$ in order to fulfill our parameter's restrictions [(2.1) and $\underline{R} < s_e - 2\varepsilon$].

Our results are in sharp contrast with the findings of Morris and Shin, who write on the same issue:[4]

---

[4]Morris and Shin, 1998, page 595.

Above all, our analysis suggests an important role for public announcements by the monetary authorities, and more generally, the transparency of the conduct of monetary policy and its dissemination to the public. If it is the case that the onset of currency crises may be precipitated by higher-order beliefs, even though participants believe that the fundamentals are sound, then the policy instruments which will stabilize the market are those which aim to restore transparency to the situation, in an attempt to restore common knowledge of the fundamentals.

Finally, we see that for fixed ε, both a lower cost and a higher depreciation rate imply a larger range of states in which the peg is abandoned. But he general picture does not change for variations in s_e; in this case the game becomes roughly an invariant translation: An increase in s_e must come forth with the same increase in R.

2.4. Extensions

2.4.1. Determinants of exchange rates. So far we have considered that at the prevailing exchange rate e there is an excess supply s_e. Besides the own exchange rate, e, this excess supply could actually be a function of some fundamental value, θ. Variable θ can be conformed by various internal and external market forces. A change in θ may trigger a move away from the exchange rate fundamental value e^* as reflected by a variation in the excess supply s_e. In light of the preceding analysis, there is no loss of generality to assume that both θ and s_e are common knowledge, since the case of heterogeneous beliefs on θ can be easily accommodated under our framework.

We may even suppose that the fundamental value of the currency e^* depends on the amount of speculation, s, and the amount of reserves, R. The idea is that extensive speculation may lead to an undershooting of the currency, and an increased amount of reserves may instill confidence in the economy. Again, these considerations can easily be integrated into the above framework.

2.4.2. Currency bands. In a currency band, the exchange rate is allowed to fluctuate within certain margins. In our simple model, the risk of speculation becomes smaller the longer the currency is away from the floor, say $\underline{e}$. Let us assume that the value of the currency is at some point $e_0 > \underline{e}$. Then, in a speculative attack the exchange rate would have to move first from e_0 to $\underline{e}$. This move puts downward pressure on both the excess supply, s_e, and the gains from speculation,

δ, as the government would trade reserves at the lower exchange rate, $\underline{e}$. Likewise, other managed float schemes may seem appropriate to minimize the benefits of speculation and boost currency stability.

2.4.3. Large trader, and sequential games. These important extensions are considered in Corsetti et al. (2004). If the actions of the large trader are not observed, it does not seem to be so obvious what would be their effects on the behavior of the small traders. Nevertheless, Proposition 5 above suggests that for very small transaction costs, speculators will attack the currency whenever $s_e + 1 > R$. Hence, in those cases it seems that the existence of a large trader will not change the results. But as Figure 2 shows the degree of asymmetric information matters when transaction costs are larger.

2.5. Concluding Remarks

In this paper we study a game-theoretic model of currency speculation with asymmetric information. There is a continuum of speculators that can short the currency and a government that holds a stock of international reserves to sustain a currency peg. Unlike most global games, speculators' actions aren't global complements. We nevertheless establish existence of a unique threshold equilibrium, and this is the only equilibrium of the game.

In various numerical exercises we observe that variations in the degree of asymmetric information have mild effects in equilibrium outcomes. The strategic substitutability condition embedded in the game seems to lead to more active speculation behavior as the degree of asymmetric information vanishes. These asymmetric information effects are more pronounced under large transaction costs (or small gains from currency depreciation) where traders with diffused priors become less forthcoming about the benefits of the currency attack.

Under asymmetric information, both transactions costs and capital gains influence the required amount of reserves to deter a currency attack. But as seen by Proposition 2.6 above, for small transaction costs, as the noise goes to zero, such required amount of reserves has to exceed the borrowing capacity of speculators. Furthermore, our numerical experiments suggests that this required quantity remains invariant to changes in the degree of asymmetric information. Therefore, our results point at the inherent instability of fixed exchange rate regimes. A shock to the economy that generates an excess supply of currency may need to be accommodated by a corresponding increase in international reserves: If international reserves cannot be spared then a speculative attack would be the likely outcome. Therefore, policy

coordination among central banks and other global institutions would be the most effective tool to avoid currency attacks by enlarging the borrowing capacities of the economy. Of course, these attacks will be more intensive the further away is the peg from the fundamental value; likewise, the so called "sand-in-the-wheels" in the way of taxes or other transaction costs will have mitigating effects on currency speculation.

CHAPTER 3

Technology Adoption with Learning by Doing and Switching Costs

ABSTRACT. We present a dynamic, finite-time model in which two long-lived sellers compete at each period for a short-lived buyer. One of the sellers has the option to adopt a new technology for production which exhibits both switching costs and learning by doing. We show that some efficient technologies are not adopted in equilibrium. Switching costs and learning by doing give incentives to the second seller to undercut prices and render the adoption unprofitable. We characterize the set of technologies which are adopted in equilibrium and show that those technologies which are learned faster—and not necessarily those which are more efficient—are more likely to be adopted.

3.1. Introduction

In this paper we study a dynamic model of technology adoption and price competition. Our main goal is to examine whether efficient technologies are introduced in equilibrium or not. The adoption of a new technology is characterized by two main features: First, adopting a new technology entails a switching cost with respect to the old technology; Second, the productivity of the new technology is advanced through learning by doing.

In our model, two long-lived sellers compete at each of a finite number of dates for a short-lived buyer with unit demand. At each date, one of the sellers has the option to adopt a given new technology—if he has not done that yet. Then, after observing the adoption decision, the sellers compete by simultaneously declaring prices to the living buyer, who buys from the seller who offers the highest consumer surplus.

The main idea that the model intends to capture is simple. As adopting a new technology involves switching costs, the adopting seller will be, for some period, in a market disadvantage with respect to its rival. This temporary market disadvantage may become perpetual if his rival trades with the buyers that the adopting seller needs to learn the new technology. Hence, the prospect of future "monopoly" profits that the rival seller might obtain if a new technology remains unlearned discourages its adoption in the first place.

Our basic model has a number of interesting results. First, and in line with much empirical evidence [see Holmes and Schmitz (2010); Parente and Prescott (2002)], some efficient technologies are never adopted in equilibrium. Second, equilibrium adoption exhibits what we call the impatience property: Among technologies that have the same social surplus, the adoption is biased towards technologies whose benefits are received earlier rather than later. Put somewhat differently: the adoption of a new technology is determined not only by the social surplus it creates but also by the inter-temporal distribution of this surplus. For the same social surplus, some technologies might be adopted and others might not; those who are adopted tend to exhibit smaller switching costs and also smaller productivity improvements over time.

Our basic model can also be extended to explore the value of competition. That is, the effect that the presence of an additional seller has on the equilibrium set of adopted technologies.

3.1.1. Related Literature. Our results show that Pareto technologies may remain unused even in the absence of sunk adoption costs. In its interest to understand adoption by a seller, our paper is connected with Arrow (1962) who was the first to compare the incentives to adopt a technology under competition and monopoly. However, Arrow and the literature that follows [see, for instance, Gilbert and Newbery (1982)] assume away both the presence of switching costs and of learning by doing.

Our paper is also related to a series of papers in which learning by doing and switching cost do play an important role. Jovanovic and Nyarko (1996) discuss a model of a single decision maker who must decide how fast to switch to a new—and potentially more productive—technology. Switching entails a productivity loss that might prevent an agent who is already skilled in the old technology to adopt a new one, even when the latter is superior. Chari and Hopenhayn (1991) also exploit the idea that capital is specific to a given technological vintage in a perfectly competitive model to explain adoption and diffusion of new technologies. Clearly, our paper shares some of these features. We inspect, however, the adoption of Pareto technologies in an strategic setting. This allows us to exploit the idea that switching costs represent an opportunity for the non-adopting seller to undercut prices and thus increase the cost of trading with the adopting seller in the long-run. This distinctive feature of our model is absent in those related papers.

In its emphasis on learning by doing and how it affects the dynamics of leadership in an industry our model is related to Cabral and Riordan (1994). Although their main goal is to understand whether

market dominance may emerge as an equilibrium outcome, there is a link between our results and theirs. In our model a technology is never adopted when the adopting seller anticipates that his rival will become (or continue being) the market leader. Put somewhat differently: Adopted technologies are those which guarantee market leadership to the adopting seller.

Two more papers are related to ours: Bergemann and Välimäki (2006) and Holmes et al. (2008). Our framework shares some formal similarities with that of Bergemann and Välimäki (2006). Notably, we also assume that the sellers compete by declaring prices in a transferable utility economy. Moreover, in our set-up it also results crucial for the efficiency of the equilibrium how the surplus created by a new technology is divided among buyers and sellers. A difference that should be noted is that buyers are short-run agents in our framework while there is a unique long-run buyer interacting with the sellers in their dynamic game. Our assumption of short-lived buyers intends to capture the notion of negligible and perfectly competitive buyers.

Holmes et al. (2008) is closest to our paper. These authors explain why a more competitive environment might lead to higher incentives to adopt new technologies. Key for their results is the presence of switching cost. At the heart of our model lies the interplay between switching cost and learning by doing in a dynamic model of price competition. This trade-off allows us to complement their results. Hence we are able to offer a sharp characterization of what types of technologies could be adopted in equilibrium. The next target in our agenda is to understand under which conditions competition has a positive value.

3.2. The Model

We consider a transferable utility (TU) economy which lasts finitely many dates $t = 0, \ldots, T$. There is one short-lived buyer alive at each date who can buy at most one good from one of two long-lived sellers $s \in \{i, j\}$. If the buyer alive at date t buys from seller s, they create the (static) gain from trade $\theta_s \geq 0$. The gains from trade (hereafter, gains) are just the difference between the utility obtained by the buyer and the production cost of the seller. We use this terminology because, as will become clear later, it is irrelevant whether a difference in gains arises from a higher utility or a lower cost—we allow for heterogeneous goods.

Seller i has the option of adopting a new technology. The gains associated to the new technology are given by a function

$$\theta : \{0, \ldots, T\} \longrightarrow \mathbb{R}_+$$

of the past cumulative sales x made by seller i. Since this function contains all the relevant information about the new technology, we refer to the latter as the technology θ. We assume that θ fulfills

$$(3.1) \qquad\qquad\qquad \theta(0) \leq \theta_j,$$

$$(3.2) \qquad\qquad\qquad \theta(x+1) \geq \theta(x)$$

and

$$(3.3) \qquad\qquad\qquad \sum_{x=0}^{T} \theta(x) \geq (T+1) \max_s\{\theta_s\}.$$

The inequalities (3.1) and (3.2) say that, whereas the first sale made by seller i with the new technology creates a smaller gain than the one created with seller j, the gain from each additional sale increases with the amount of past cumulative sales made by seller i. In other words, the inequality (3.1) captures the idea of switching costs and the inequality (3.2) learning by doing. The last inequality states that producing the $T+1$ goods with the new technology is a weak Pareto improvement (recall that in a TU economy an allocation is Pareto optimal if, and only if, maximizes the social surplus). The set of all technologies fulfilling (3.1), (3.2) and (3.3) for a given pair of old technologies $\{\theta_i, \theta_j\}$ is called a Pareto set. We denote it by Θ. The following example illustrates:[1]

EXAMPLE 3.1. Let $T = 1$. Suppose that the utility to each buyer from a unit of consumption is u. Let c_s be the cost of production of seller s and (c^0, c^1) be costs of production with the new technology if zero and one cumulative sales have been made. That is: $\theta_s = u - c_s$, $\theta(0) = u - c^0$ and $\theta(1) = u - c^1$. The set of cost pairs $C = \{(c^0, c^1) \in \mathbb{R}_+^2 : c^0 \geq c_j, c^1 \leq c^0, c^0 + c^1 \leq 2\min_s\{c_s\}\}$ characterizes the largest subset of the Pareto set where $\theta(1) \leq u$. This set is depicted in Figure 1 for the case $c_i = c_j = 1$.

We assume that seller i has, at each date, the option to switch from the old technology θ_i to the new technology θ if he has not switched at any previous date. This choice is irreversible; once seller i has adopted the new technology he cannot switch back to the old one.

At each date, and after observing the adoption decision, the two sellers compete by simultaneously declaring prices to the living buyer. Then, the buyer buys from the seller offering him the highest surplus. We assume that the sale is always made by the seller who is willing to offer the highest surplus. In case of a tie the sale is made by seller
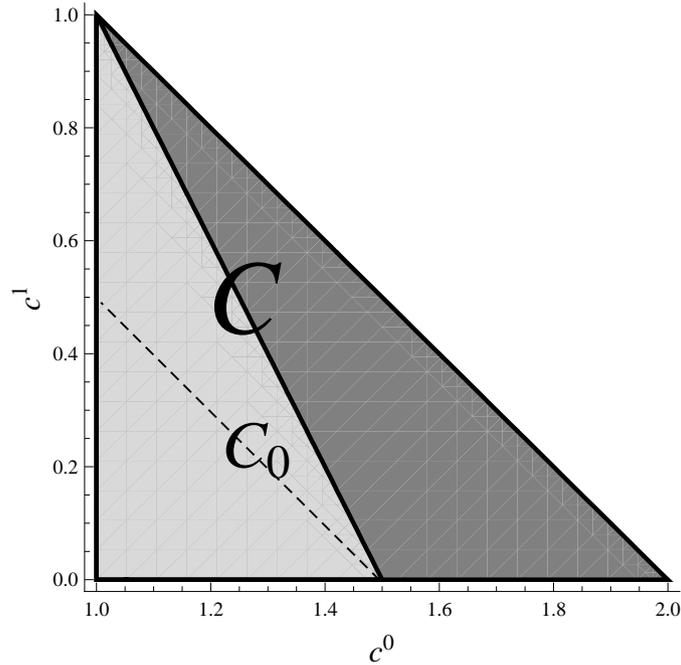
---

[1]All the examples below continue on this one.

FIGURE 1. Example with $T = 1$, constant utility and costs $c_i = c_j = 1$.

i. Both sellers wish to maximize the sum of their short-run profits—without discounting.

3.3. The Equilibrium

We restrict our attention to Markov perfect equilibria. A MPE is a perfect equilibrium in which the strategies of the players only depend on the pay-off relevant information. In our economy this information is: (a) the current date t, (b) whether seller i adopted the new technology at some previous date and (c) the number x of sales made by seller i with the new technology.

3.3.1. Dynamic price competition. We solve for equilibria by backwards induction. Because the adoption decision is irreversible, we solve first for all states in which the new technology has been adopted already. That is, we abstract momentarily from the adoption decision and concentrate on the dynamic price competition that starts once the new technology has been adopted. This is a necessary preliminary step before addressing optimal adoption decisions; one needs to compute the value from adopting a new technology to decide whether one should

adopt it or not. To simplify the notation, we will refer to a pair (x, t) as a state within this subsection.

Once the new technology has been adopted, the sales made by each seller have a prominent strategic role. The key point is that each sale accrued by seller i strengthens his strategic position against seller j. Note that because of (3.3), there must be a minimum amount of sales $\underline{x}$ such that $\theta(x) \geq \theta_j$ for all $x \geq \underline{x}$, formally:

$$\underline{x} := \min\{x : \theta(x) \geq \theta_j\}.$$

If seller i manages to accrue $\underline{x}$ sales, then he sells at a profit at every subsequent period. But every sale made before $\underline{x}$ should entail a short-run loss because the technology of seller j is better whenever $x < \underline{x}$. On the other hand, seller j can prevent seller i from learning the new technology by incurring in short-run losses. If seller j manages to reach a state (x, t) for which the periods remaining are not more than the sales before $\underline{x}$; if $T + 1 - t \leq \underline{x} - x$, then he also sells at a profit at every subsequent period. In summary, the value from making a sale transcends short-run considerations and both sellers are, in principle, willing to incur in short-run losses to improve their strategic position.

For the exposition, we find it convenient to view the sellers as offering surplus to the buyers rather than declaring prices. From this perspective, sellers can be seen as bidding for the buyers in a second-price auction at each state. Let $V_s(x, t)$ denote the profit (value) of seller s at state (x, t). The maximum amount of utility that each seller is willing to transfer to the buyer at state (x, t) is given by the following *bidding functions*:

$$b_i(x, t) = \theta(x) + V_i(x + 1, t + 1) - V_i(x, t + 1),$$
$$b_j(x, t) = \theta_j + V_j(x, t + 1) - V_j(x + 1, t + 1).$$

They read as follows: suppose that seller i had to transfer $b_i(x, t)$ to the seller to make a sale. If he sells, he earns the short-run gain $\theta(x)$, minus the transfer $b_i(x, t)$, plus the continuation value $V_i(x + 1, t + 1)$; if he does not sell, he earns his continuation value $V_i(x, t + 1)$. The bidding functions give us the amount that leaves each seller indifferent between selling or not at (x, t).

We have assumed that seller i sells at (x, t) if, and only if, $b_i(x, t) \geq b_j(x, t)$. Since the winner must pay the bid of the loser, the value of each seller at (x, t) is written as:

$$V_i(x, t) = \max\{\theta(x) - b_j(x, t) + V_i(x + 1, t + 1), V_i(x, t + 1)\}$$
$$V_j(x, t) = \max\{\theta_j - b_i(x, t) + V_j(x, t + 1), V_j(x + 1, t + 1)\}.$$

EXAMPLE 3.2. Suppose that seller i introduced the new technology at $t = 0$. The bidding functions at $t = 1$ are: $b_i(0, 1) = u - c^0$, $b_i(1, 1) = u - c^1$ and $b_j(0, 1) = b_j(1, 1) = u - c_j$. This implies that the values at $t = 1$ are: $V_i(1, 1) = c_j - c^1$, $V_j(0, 1) = c^0 - c_j$ and $V_i(0, 1) = V_j(1, 1) = 0$, because the price at $t = 1$ must equal the highest cost. Going one period backwards, we have that $b_i(0, 0) = u - c^0 + c_j - c^1$ and $b_j(0, 0) = u - c_j + c^0 - c_j$. This implies that $V_i(0, 0) = \max\{3c_j - 2c^0 - c^1, 0\}$ and $V_j(0, 0) = \max\{2c^0 + c^1 - 3c_j, 0\}$, because the price at $t = 0$ must equal the loser's reservation price: $\min\{2c_j - c^0, c^0 + c^1 - c_j\}$. The set of cost pairs $C_0 = \{(c^0, c^1) \in \mathbb{R}_+^2 : c^0 \geq c_j, c^1 \leq c^0, 2c^0 + c^1 \leq 3c_j\}$ for $c_j = 1$ is depicted in Figure 1.

The previous example shows that—if $T = 1$ and seller i introduces the new technology at $t = 0$—either seller i sells at both periods or at none of them. The following proposition states that this is true in general for the sub-game that starts when the new technology is adopted.

PROPOSITION 3.3. *Consider the sub-game that starts when the new technology is adopted. In a MPE of this sub-game the same seller sells at each date.*

On the other hand, in Example 3.2 we have that seller i is the winning seller if, and only if, $2c^0 + c^1 \leq 3c_j$. That is, there is a unique equilibrium for each specification of the costs. This is also true in general, as stated in the next proposition.

PROPOSITION 3.4. *The sub-game that starts when the new technology is adopted has a unique MPE.*

3.3.2. The adoption decision. We turn now to the adoption decision. We assume that if seller i is indifferent between adopting or not the new technology he will not adopt it. This rules out some uninteresting equilibria. Under this assumption, our first result states that the new technology is never adopted with delay.

PROPOSITION 3.5. *In a MPE the new technology is either adopted at $t = 0$ or never.*

Therefore, seller i adopts the new technology only if he finds it profitable at $t = 0$. The next proposition gives a necessary and sufficient condition for this to happen.

PROPOSITION 3.6. *Let Z be*

$$Z(\theta) := \sum_{k=0}^{T} (T + 1 - k)[\theta(k) - \theta_j].$$

In a MPE the new technology is adopted if, and only if,

$$(3.4) \qquad Z(\theta) > (T+1) \max\{\theta_i - \theta_j, 0\}.$$

Note that the function Z puts more weight on the lower values of k for which $\theta(k) - \theta_j$ is smaller (negative). Hence, Z is typically smaller than $\sum_k [\theta(k) - \theta_j]$. We will see below that, for this reason, some efficient technologies are not adopted in equilibrium. The following result is an immediate consequence of the Propositions 3.3–3.6.

COROLLARY 3.7. *There is a unique MPE for each $\theta \in \Theta$. If (3.4) holds the new technology is adopted at $t = 0$ and seller i sells at every date. If (3.4) does not hold the new technology is never adopted. In this latter case, seller i sells at every date if $\theta_i \geq \theta_j$; otherwise, seller j sells at every date.*

3.4. Efficiency

In this section we are interested in the comparison of different technologies that have the same total gain. In doing so, we will isolate the role of the inter-temporal distribution of gains (i.e., the speed of the learning process) in determining equilibrium outcomes. To begin with, we partition the set of Pareto technologies according to their total gain. Let $\Gamma(g)$ denote the class of Pareto technologies with total gain g:

$$\Gamma(g) := \left\{ \theta \in \Theta : \sum_{k=0}^{T} \theta(k) = g \right\}.$$

Our first result states that some classes are so productive that all technologies within these classes are adopted in equilibrium—no matter how the total gain is distributed across dates.

PROPOSITION 3.8. *If $g > (T+1)(\max_s\{\theta_s\} + \frac{T}{2}\theta_j)$, then all technologies in $\Gamma(g)$ are adopted in equilibrium.*

For the rest of classes, however, we can always find a Pareto technology which is not adopted in equilibrium. These classes exhibit what we shall call the *impatience property*. The word "impatience" refers to the fact that those technologies which are learned faster (in the precise sense of having a larger Z) are the ones that are adopted in equilibrium.

PROPOSITION 3.9. *All classes $\Gamma(g)$ with $g \leq (T+1)(\max_s\{\theta_s\} + \frac{T}{2}\theta_j)$ exhibit the impatience property (excluding $g = (T+1)\max_s\{\theta_s\}$ if $T = 1$).*

The impatience property implies, for example, that if seller i is offered to choose among several technologies, he may choose to adopt a

less efficient one just because its gains are more biased towards earlier periods. Note that this preference appears in a world without time discounting and without borrowing constraints. The technologies which are learned faster leave less profits to seller j when he prevents seller i from reaching the minimum scale $\underline{x}$ (of course, this occurs out of the equilibrium path). For this reason, they are easier to implement in equilibrium.

EXAMPLE 3.10. Consider the set of cost pairs $\{(c^0, c^1) \in \mathbb{R}_+^2 : c^0 \geq 1, c^1 \leq c^0, c^0 + c^1 = \kappa\}$ for $\kappa \in [1, 2]$. These classes are lines of slope -1 in Figure 1 (parallel displacements of the dashed line). We see that, if $\kappa < 1.5$, all technologies in the same class are adopted in equilibrium (the whole line lies in C_0). If $\kappa \geq 1.5$, however, some technologies lie outside C_0 and, therefore, are not adopted in equilibrium. This is an elementary illustration of the impatience property: Let $\kappa_1 > \kappa_2 > 1.5$. Seller i would prefer a less efficient technology—with $c^0 + c^1 = \kappa_1$—that is adopted in equilibrium to a more efficient one—with $c^0 + c^1 = \kappa_2$—which is not adopted.

3.5. Future Work

The next step in our agenda will be the study of *the value of competition*. We will get into this issue by introducing a third seller m in our model. The next example gives us an idea of the results that we may expect.

EXAMPLE 3.11. Suppose that we introduce a new seller m with a cost of production $c_m \in [c_j, c^0]$ (and that $c_i = c_j$). If seller i introduces the new technology and sells at $t = 0$, he will also sell at $t = 1$ at the price c_j. This means that he is willing to lose, at most, $c_j - c^1$ at $t = 0$, which implies that his reservation price at $t = 0$ is $c^0 + c^1 - c_j$. If seller j has sold at $t = 0$ and seller i has introduced the new technology, seller j will also sell at $t = 1$ at the price c_m. This means that he is willing to lose, at most, $c_m - c_j$ at $t = 0$, which implies that his reservation price at $t = 0$ is $2c_j - c_m$. Therefore, the new technology will be adopted if, and only if, $c_m + c^0 + c^1 \leq 3c_j$. Since $c_m \leq c^0$, the set of technologies which are adopted with the new seller is never smaller than with two sellers.

Introducing a third seller in our example reduces the profits that seller j will get in the second period (provided that seller i introduced the new technology and that seller j sold at $t = 0$). Seller j has to lower the price to compete with seller m, which reduces his incentives to sell at $t = 0$. On the other hand, the value from selling at $t = 0$ for seller i remains unchanged because $c_m \leq c_j$ (he does not have to

lower his second-period price). This explains why more technologies are adopted in equilibrium.

If a new technology is adopted, the price at $t = 0$ is equal to the reservation price of the seller j, that is, equal to $2c_j - c_m$. Since $c_m \leq c^0$, this leads us to the surprising result that increasing competition may lead to higher prices at $t = 0$ (recall that the same price with two sellers was $2c_j - c^0$). The reason for this is that the new seller reduces the seller j's willingness to transfer utility to the buyer alive at $t = 0$, which, in turn, increases the market power of seller i. Since the second-period price does not change, consumers are generally worse-off with the new seller for those technologies which were also adopted with two sellers. If a technology was not adopted with two sellers, the price at both periods was equal to $\max_s\{c_s\}$. Consider a technology which was not adopted with two seller but it is now adopted with three. Because $c_m \geq c_j$, the consumer surplus with two sellers $2(u - \max_s\{c_s\})$ is never larger than the consumer surplus with three sellers $2u + c_m - 3c_j$. Lastly, for the technologies which are not adopted even with three sellers the equilibrium prices do not change.

Details of Chapter 1

A.1. Proofs

Two properties of the functions G and v will be used repeatedly. First: because G is a distribution function, it is right-continuous, which, in turn, implies that v is also right-continuous. Second: v is increasing whenever G is constant and vice versa.

A.1.1. Preliminary results.

LEMMA A.1. *Given the conditions of Lemma 1.3, $v(\underline{t}) = v(\bar{t}) = v^*$.*

PROOF. If either the infimum or the supremum is an isolated point of the support, then it is a mass point of the distribution F and v at that point must be v^*.

Because v is right-continuous, $\lim_{t \downarrow \underline{t}} v(t) = v(\underline{t})$, which, if $\underline{t}$ is not isolated, implies that $v(\underline{t}) = v^*$.

Suppose that $v(\bar{t}) < v^*$. If $\bar{t}$ is not isolated, this can happen only if v has a downward jump at $\bar{t}$, which, by (1.3), only occurs if G jumps at $\bar{t}$. G jumps only if s (and F) also have a jump at the same point; a point in which s first surpasses a strictly positive mass of sample paths of κ. But this means that $\bar{t}$ is a mass point of F, and so $v(\bar{t}) = v^*$. □

LEMMA A.2. *Given the conditions of Lemma 1.3, $\underline{t} > 0$, and the support of F is an interval.*

PROOF. We know from Proposition 1.2 that F is non-degenerate. Because G is right-continuous and $G(0) < 1$, we know that for every $\epsilon > 0$ there must exist some δ such that $t < \delta$ implies that $G(t) < G(0) + \epsilon$. Take any $\epsilon < 1 - G(0)$ and see that $v(\delta) > 1 - c$. This shows that $v(\underline{t}) = v^* > 1 - c$, and so $\underline{t} > 0$.

We will use the following properties of the function s. We must have both $s(t) < \kappa(1, t)$ for all $0 < t \le \arcsin(\mu)$ and $s(t) = \mu$ for all $t \ge \pi - \arcsin(\mu)$. Otherwise, the bubble bursts no later than $\bar{t}$ with probability one, which would imply that $v(\bar{t}) = v^* = 1 - c$.

Now we can show that F is strictly increasing for all $\underline{t} \le t \le \bar{t}$. Suppose that $F(t_0) = F(t_1) = k_0/\mu$ for some $\underline{t} \le t_0 < t_1 \le \bar{t}$, that is, $s(t) = k_0$ for all $t \in [t_0, t_1]$. We will see that this leads to

a contradiction. Let $t_- = \inf\{t | F(t) = k_0/\mu\}$, that is, t_- belongs to the support of F. Let x_- be the solution to the equation $\kappa(x, t_-) = k_0$ (that this equation has a solution is implied by the previous paragraph). The derivative of $\kappa(x_-, t)$ with respect to t, evaluated at t_-, cannot be strictly positive. If this were the case, then, there would exist some $\epsilon > 0$ such that $s(t) = k_0 < \kappa(x_-, t)$ for all $t_- < t \leq t_- + \epsilon$, which, in turn, implies that no market price path decreases therein. But this contradicts that fact that t_- belongs to the equilibrium support; any point fulfilling $t_- < t \leq t_- + \epsilon$ pays more than $v(t_-) = v^*$ because $G(t) = G(t_-)$.

Let $t_+ = \sup\{t | F(t) = k_0/\mu\}$, that is, t_+ belongs to the support of F. Since some price path has to decrease in $(t_-, t_- + \epsilon]$ for every $\epsilon > 0$, and since the aforementioned derivative has to be either negative or zero, we have that a price path decreases at each instant $t_- < t < t_+$, that is, a sample path of κ intersects s for the first time at each of those instants. Let x_+ be solution to the equation $\kappa(x, t_+) = k_0$. What we have just shown is that the function T is

$$T(x) = \left[\pi - \arcsin\left(\frac{k_0}{x} \right) \right] x$$

for all $x_- < x < x_+$. Therefore, we can make the change of variable $t = T(x)$ and rewrite the pay-off function v for $t_- < t < t_+$ as

(A.1)
$$e^{(g-r)T(x)}(1 - x) + x - c,$$

where $x_- < x < x_+$. Since the function (A.1) is uni-modal for $k_0 \leq x \leq 1$ (provided that $k_0 > 0$), we must have that v is decreasing for $t_- < t < t_+$; otherwise, the pay-off would not be maximized at t_-. The pay-off function v can only have downward jumps (which correspond to upward jumps of G) and, therefore, we must have that $v(t_-) > v(t_+)$, which contradicts the fact that t_+ belongs to the equilibrium support. This shows that F is strictly increasing in $[\underline{t}, \overline{t}]$. □

LEMMA A.3. *Given the conditions of Lemma 1.3, G is continuous and strictly increasing for all $t \leq \pi - \arcsin(\mu)$.*

PROOF. Since $s(t) = 0$ for all $t < \underline{t}$, we have that $T(x) = \pi x$ for all $x < \underline{t}/\pi$, which implies that $G(t) = t/\pi$ for all $t < \underline{t}$. Thus, G is strictly increasing and continuous for all $t < \underline{t}$.

Suppose that there were $\underline{t} \leq t_0 < t_1 \leq \overline{t}$ such that $G(t_0) = G(t_1) = h_0$. Then, $v(t_0) < v(t_1)$, contradicting the fact that $v(t_0) = v^*$, as implied by the previous lemma. Also, if G had a jump at some $\underline{t} \leq t_0 \leq \overline{t}$, then v would have a downward jump at t_0, which cannot happen

if $v(t_0) = v^*$. Thus, G is strictly increasing and continuous for all $\underline{t} \leq t \leq \overline{t}$.

There cannot be any $\epsilon > 0$ such that no price path decreases in $(\overline{t}, \overline{t}+\epsilon]$ because that would mean that $G(\overline{t}+\epsilon) = G(\overline{t})$, that is, $v(\overline{t}+\epsilon) > v(\overline{t})$, which contradicts the fact that $v(\overline{t}) = v^*$. Since s is flat for all $t \geq \overline{t}$, any sample path of κ which intersects s for the first time in that range has to be decreasing. This implies that one price path decreases at each $\overline{t} < t \leq \pi - \arcsin(\mu)$. Thus, G is strictly increasing and continuous for all $\overline{t} < t \leq \pi - \arcsin(\mu)$. $\qquad\square$

A.1.2. Results in the main text.

LEMMA 1.3. *Suppose that there is a symmetric equilibrium in mixed trigger-strategies which fulfills $G(0) < 1$. Then, the function T is strictly increasing and continuous.*

PROOF. T is monotone increasing because $\kappa(x_1, t) \geq \kappa(x_0, t)$ for all t and all $x_1 > x_0$. Let us define

$$\xi(t) := \sup \{x | T(x) \leq t\}.$$

It is clear from the definition that $T(x) \leq t$ implies $x \leq \xi(t)$ and, hence, $\mathbf{P}(T(X) \leq t) \leq \mathbf{P}(X \leq \xi(t))$. If $x < \xi(t)$, then $T(x) \leq t$; since T is monotone increasing, $T(x) > t$ would contradict the definition of $\xi(t)$. Therefore, there is at most one point which could satisfy both $x \leq \xi(t)$ and $T(x) > t$, namely $x = \xi(t)$. Because the standard uniform is a continuous distribution, any singleton has probability zero. Thus, we can write:

$$(A.2) \qquad G(t) = \mathbf{P}(T(X) \leq t) = \mathbf{P}(X \leq \xi(t)) = \xi(t).$$

Since T is monotone increasing, it can only have jump discontinuities. Suppose that T had a discontinuity at x_0 and let $t_0 < t_1$ be the one-sided limits of T as x approaches to x_0. Then, we should have that $\xi(t) = \xi(t_0)$ for all $t_0 \leq t < t_1$, which, by Lemma A.3 and (A.2), cannot happen (ξ is strictly increasing for all $t \leq \pi - \arcsin(\mu)$). This shows that T is continuous.

To show that T is strictly increasing, suppose that $T(x_0) = T(x_1) = t_0$ for some $x_0 < x_1$. Let $x_- = \inf\{x | T(x) = t_0\}$ and $x_+ = \sup\{x | T(x) = t_0\}$. Because T is increasing, $\xi(t_0) = x_+$ and $\xi(t_0-\epsilon) \leq x_-$ for all $\epsilon > 0$. This means that ξ must have a discontinuity at t_0, which cannot happen since Lemma A.3 together with (A.2) imply that ξ is continuous.

Since T is strictly increasing and continuous, then $\xi = T^{-1}$. $\qquad\square$

LEMMA 1.5. *All F fulfilling the conditions of Lemma 1.3 have the same support.*

PROOF. Using the change of variable $t = T(x)$, we can write

$$(A.3) \qquad v(T(x)) = e^{(g-r)[\pi - \arcsin(\mu/x)]x}(1 - x) + x - c$$

for all $x \geq T^{-1}(\bar{t})$. This function is uni-modal for $\mu \leq x \leq 1$. Since $v(t)$ must be smaller or equal than v^* for all $t > \bar{t}$, we know that $T^{-1}(t)$ cannot be smaller that the point $\bar{x}$ which maximizes (A.3). Otherwise, (A.3) would be strictly increasing at $T^{-1}(\bar{t})$, which contradicts the fact that $v(\bar{t}) = v^*$.

On the other hand, we know that $s(t) \leq \mu$ for all t. This implies that, for all $x > \mu$, the inequalities $\arcsin(\mu/x)x < T(x) < [\pi - \arcsin(\mu/x)]x$ cannot both be true: within the interval of time $(\arcsin(\mu/x)x, [\pi - \arcsin(\mu/x)]x)$—the mania—the absorption capacity is strictly greater than μ for all $x > \mu$, which means that no price path can decrease therein. The inequality $T(x) \leq \arcsin(\mu/x)x$ can neither be true. Suppose that there exists some $x_0 > \mu$ such that $T(x_0) \leq \arcsin(\mu/x_0)x_0$. We know that for all $x > x_0$ and all $\arcsin(\mu/x_0)x_0 \leq t \leq [\pi - \arcsin(\mu/x_0)]x_0$, $\kappa(x, t) > \mu$. Therefore, no price path decreases within the interval of time $(\arcsin(\mu/x_0)x_0, [\pi - \arcsin(\mu/x_0)]x_0)$, which is incompatible with T being continuous. In short, we must have that $T(x) \geq [\pi - \arcsin(\mu/x)]x$ for all $\mu < x \leq T^{-1}(\bar{t})$. Going back to v, we must have that

$$(A.4) \qquad v(T(x)) \geq e^{(g-r)[\pi - \arcsin(\mu/x)]x}(1 - x) + x - c$$

for all $\mu < x \leq T^{-1}(\bar{t})$. This implies that $T^{-1}(\bar{t})$ cannot be greater than $\bar{x}$, since $\bar{x}$ would then contradict the inequality. Hence, there is only one admissible $T^{-1}(\bar{t})$, namely, $\bar{x}$, and $\bar{t} = [\pi - \arcsin(\mu/\bar{x})]\bar{x}$.

Clearly, $\bar{t}$ gives us v^*. On the other hand, $\underline{t}$ is the solution to

$$(A.5) \qquad e^{(g-t)t}\left(1 - \frac{t}{\pi}\right) + \frac{t}{\pi} - c = v^*$$

for $t < \bar{t}$. Such solution always exists and is unique because the left-hand side of (A.5), with the change of variable $t = \pi x$, is uni-modal in $[0, 1]$, equal to zero at $t = 0$, and lies above the right-hand side of (A.3) for all $x > \mu$. $\qquad \square$

PROPOSITION 1.6. *There is a unique equilibrium fulfilling the conditions of Lemma 1.3. In this equilibrium, each arbitrageur plays the mixed trigger-strategy*

$$F(t) = \frac{1}{\mu}\kappa\left(T^{-1}(t), t\right)$$

for all $t \leq \pi - \arcsin(\mu)$.

PROOF. We first show that F characterizes an equilibrium. We start showing that F is a distribution function. Because $T^{-1}(t) = t/\pi$ for all $t \leq \underline{t}$, $F(t) = 0$ for all $t \leq \underline{t}$, and because $T^{-1}(t)$ is the solution to $\kappa(x, t) = \mu$ for all $t > \bar{t}$, $F(t) = 1$ for all $t \geq \bar{t}$. To prove that F is non-decreasing one should take the derivative of

$$\frac{1}{\mu} \frac{e^{(g-r)t} - v^* - c}{e^{(g-r)t} - 1} \sin\left(\frac{e^{(g-r)t} - 1}{e^{(g-r)t} - v^* - c} t\right)$$

with respect to t and show that it is non-negative for all $\underline{t} \leq t \leq \bar{t}$. The problem is that $\bar{t}$ is obtained from a maximization problem for which there is no closed-form solution, whereas $\underline{t}$ is obtained from $\bar{t}$ via an equation which neither has a closed-form solution. Since the derivative changes sign, the precise location of both points is crucial. We have computed F for a fine grid of parameter values $0 < \mu < 1$ and $0 < g - r < 1$ to find that it is indeed strictly increasing inside the equilibrium support for all of them. A program which produces an animated plot of F for such range of parameter values is printed in Appendix A.2.

We now show that F induces T, that is,

(A.6) $$T(x) = \inf\{t | \mu F(t) \geq \kappa(x, t), t \geq 0\}.$$

Take any $0 \leq x_0 \leq 1$. It is obvious that $\mu F(T(x_0)) = \kappa(x_0, T(x_0))$. Also, for all $t < T(x_0)$, we have that $T^{-1}(t) < x_0$ and, hence, $\mu F(t) < \kappa(x_0, t)$.

We have seen in the proof of Lemma 1.5 that $v(t)$ is strictly increasing for all $t < \underline{t}$ and strictly decreasing for all $t > \bar{t}$. That is, all trigger-strategies outside the equilibrium support pay less than the equilibrium pay-off v^*. It only remains to show that no other pure strategy pays more that v^*. Consider an arbitrary pure strategy involving $N > 1$ transactions (since all arbitrageurs liquidate a some time, every pure strategy which is not a trigger-strategy involves more than one transaction). Let $(\mathbf{t}, \mathbf{z})$ be the vector specifying the transaction dates $(t_1, \ldots, t_N)$ and the positions held between transactions $(z_1, \ldots, z_N)$, where $z_N = 1$. Such strategy is a plan of action; it tells us what the arbitrageur will do as long as the bubble persists. If the bubble bursts between to transaction dates, the arbitrageur liquidates. The pay-off of a general pure strategy involving finitely many transactions

is, therefore,

$$V((\mathbf{t}, \mathbf{z})) := \sum_{n=1}^{N} \left[(z_n - z_{n-1})e^{(g-r)t_n} - c \right] [1 - G(t_n)]$$
$$+ (1 - z_{n-1} - c)[G(t_n) - G(t_{n-1})]\mathbf{1}_{[0,1)}(z_{n-1}),$$

where $z_0 = t_0 = 0$. It is a simple matter of algebra to show that

$$V((\mathbf{t}, \mathbf{z})) < \sum_{n=1}^{N} [z_n - z_{n-1}]v(t_n)$$

whenever $(\mathbf{t}, \mathbf{z})$ does not correspond to a trigger-strategy. We can rewrite the right-hand side as

$$v(t_N) + \sum_{n=1}^{N-1} z_n[v(t_n) - v(t_{n+1})].$$

Now let $N_1 = \max\{n | t_n \leq \bar{t}\}$. Since v is non-decreasing for all $t < \bar{t}$, we have that the last expression is bounded above by

$$v(t_N) + \sum_{n=N_1}^{N-1} z_n[v(t_n) - v(t_{n+1})],$$

which, because $0 \leq z_n \leq 1$ and $v(t_n) - v(t_{n+1}) \geq 0$ for all $n \geq N_1$, cannot be greater than v^*.

We now show that there is no other equilibrium F fulfilling (A.6). Because any F has to be continuous from the right, we know that

$$t_0 = \inf\{t | \mu F(t) \geq \kappa(x_0, t), t \geq 0\}$$

if, and only if, (a) $\mu F(t) < \kappa(x_0, t)$ for all $t < t_0$ and (b) $\mu F(t_0) \geq \kappa(x_0, t_0)$. Thus, any other equilibrium F must fulfil (b) with strict inequality for some x_0. This means that F must have a jump at t_0. Now, let x_1 be the solution to $\kappa(x, t_0) = \mu F(t_0)$. We must have both $x_1 > x_0$ and $T(x_1) = T(x_0)$ (because (a) also implies that $\mu F(t) < \kappa(x_1, t)$ for all $t < t_0$), which contradicts Lemma 1.3. $\square$

A.2. Program

The following *Mathematica 6*[1] program produces an animated plot of F for parameter values $10^{-4} \leq \mu \leq 1$ and $10^{-4} \leq g - r \leq 1$. The functions w, a, and b in the program correspond to the quantities $v^* + c$, $T^{-1}(\underline{t})$, and $T^{-1}(\bar{t})$; g, μ, and F correspond to $g - r$, μ, and F in the model; ϕ and Φ are auxiliary functions.

---

[1] www.wolfram.com

$\text{Clear}[\phi, w, b, a, \Phi, F]$

$\phi[\text{g\_}, \mu\text{\_}]:=\phi[g, \mu] = \text{Maximize}\left[\left\{e^{g\left(\pi-\text{ArcSin}\left[\frac{\mu}{x}\right]\right)x}(1-x)+x, \mu \leq x \leq 1\right\}, x\right]$

$w[\text{g\_}, \mu\text{\_}]:=\phi[g, \mu][[1]]$

$b[\text{g\_}, \mu\text{\_}]:=\phi[g, \mu][[2, 1, 2]]$

$a[\text{g\_}, \mu\text{\_}]:=a[g, \mu] = \text{FindRoot}\left[e^{g\,\pi\,x}(1-x)+x == w[g, \mu], \{x, 0\}\right][[1, 2]]$

$\Phi[\text{t\_}, \text{g\_}, \mu\text{\_}]:=\frac{1}{\mu}\frac{e^{g\,t}-w[g,\mu]}{e^{g\,t}-1}\text{Sin}\left[\frac{e^{g\,t}-1}{e^{g\,t}-w[g,\mu]}t\right]$

$F[\text{t\_}, \text{g\_}, \mu\text{\_}]:=\begin{cases} 0 & 0 \leq t < \pi\,a[g, \mu] \\ \Phi[t, g, \mu] & \pi\,a[g, \mu] \leq t < \left(\pi - \text{ArcSin}\left[\frac{\mu}{b[g,\mu]}\right]\right)b[g, \mu] \\ 1 & t \geq \left(\pi - \text{ArcSin}\left[\frac{\mu}{b[g,\mu]}\right]\right)b[g, \mu] \end{cases}$

$\text{Manipulate}[\text{Plot}[F[t, g, \mu], \{t, 0, \pi\}], \{\mu, 0.0001, 1\}, \{g, 0.0001, 1\}]$

APPENDIX B

Details of Chapter 3

PROPOSITION 3.3. *Consider the sub-game that starts when the new technology is adopted. In a MPE of this sub-game the same seller sells at each date.*

PROOF. We shall prove the result for the case in which adoption takes place at $t = 0$. For later adoption dates we only need to change the number of periods remaining—put $T + 1 - t$ instead of $T + 1$. The proof proceeds in four steps:

Step 1. Let us define the function τ as

$$\tau(x) := \max\{t : b_i(x, t) \geq b_j(x, t)\}$$

if the set $\{t : b_i(x, t) \geq b_j(x, t)\}$ is not empty, and $\tau(x) := x - 1$ if it is empty. We see that $\tau(x) = T$ for all $x \geq \underline{x}$, because, for all states (x, t) with $x \geq \underline{x}$, we have that $\theta(x) \geq \theta_j$. That is, we know that seller i will sell at every successor state of (x, t).

Step 2. We show now that $\tau(x + 1) > \tau(x)$ for all $x < \underline{x}$. The proof is by contradiction. Suppose that $\tau(x + 1) \leq \tau(x)$ for some $x < \underline{x}$. Then, we have that $b_j(x, t) > b_i(x, t)$ and $b_j(x + 1, t) > b_i(x + 1, t)$ for all $t > \tau(x)$, which implies that $b_j(x, \tau(x)) > b_i(x, \tau(x))$ because seller i knows at $(x, \tau(x))$ that he will not sell at any later period. But this contradicts the definition of τ.

Step 3. For every state (x, t) such that $t < \tau(x)$ (and $x < \underline{x}$), we have that $b_i(x, t) \geq b_j(x, t)$. Furthermore, we have that $V_j(x, t) = 0$ and $V_i(x, t) = w_i(x, t)$. Let us start at the state $(\underline{x} - 1, \tau(\underline{x} - 1) - 1)$. At this state, seller j knows that he will not sell at any future date—note that, even if he sells, since $b_i(\underline{x} - 1, \tau(\underline{x} - 1)) \geq b_j(\underline{x} - 1, \tau(\underline{x} - 1))$, he knows that the state will move from $(\underline{x} - 1, \tau(\underline{x} - 1))$ to $(\underline{x}, \tau(\underline{x} - 1) + 1)$ in the next period. Therefore, he is willing to bid, at most, the short-run gain: $b_j(\underline{x} - 1, \tau(\underline{x} - 1) - 1) = \theta_j$. This implies that the value at this state for him must be zero: $V_j(\underline{x} - 1, \tau(\underline{x} - 1) - 1) = 0$, and that the value for seller i must be maximum: $V_i(\underline{x} - 1, \tau(\underline{x} - 1) - 1) = w_i(\underline{x} - 1, \tau(\underline{x} - 1) - 1)$. Going one period backwards, we see that seller j also knows at state $(\underline{x} - 1, \tau(\underline{x} - 1) - 2)$ that he will not sell at any future date. Repeating the same argument over and over again, we

find that $b_i(\underline{x} - 1, t) \geq b_j(\underline{x} - 1, t)$ for all $t < \tau(\underline{x} - 1)$. Furthermore, $V_j(\underline{x} - 1, t) = 0$ and $V_i(\underline{x} - 1, t) = w_i(\underline{x} - 1, t)$ for all $t < \tau(\underline{x} - 1)$.

Because τ is strictly increasing, seller j knows at $(\underline{x} - 2, \tau(\underline{x} - 2) - 1)$ that he will not sell at any future date. Using the same reasoning of the previous paragraph, we can establish the corresponding result for all the states $(\underline{x} - 2, t)$ such that $t < \tau(\underline{x} - 2)$.

Repeating the same argument over and over again for each $x < \underline{x} - 2$ we see that, for all states (x, t) such that $t < \tau(x)$, we have: $b_i(x, t) \geq b_j(x, t)$, $V_j(x, t) = 0$ and $V_i(x, t) = w_i(x, t)$.

Step 4. We have just seen that $b_i(x, t) \geq b_j(x, t)$ for all states (x, t) such that $t \leq \tau(x)$. This, combined with the fact that τ is strictly increasing, implies that either (a) seller i sells at every period (if $\tau(0) \geq 0$), or (b) seller j sells at every period (if $\tau(0) = -1$). $\qquad\square$

PROPOSITION 3.4. *The sub-game that starts when the new technology is adopted has a unique MPE.*

PROOF. We shall prove the result for the case in which adoption takes place at $t = 0$. For later adoption dates we only need to change the number of periods remaining—put $T + 1 - t$ instead of $T + 1$. Our proof constructs the equilibrium. Let us start defining the auxiliary functions w_i and w_j as:

$$w_i(x, t) := \sum_{k=0}^{T-t} [\theta(x + k) - \theta_j],$$
$$w_j(x, t) := (T + 1 - t)[\theta_j - \theta(x)].$$

We proceed in five steps:

Step 1. For every state (x, t) such that $x \geq \underline{x}$, we have that

$$b_i(x, t) = \theta(x) + w_i(x + 1, t + 1) - w_i(x, t + 1)$$
$$b_j(x, t) = \theta_j$$

and

$$V_i(x, t) = w_i(x, t)$$
$$V_j(x, t) = 0.$$

This is true because both sellers know that the seller i will be the winning seller at every future date.

Step 2. Let $x < \underline{x}$. For every state (x, t) such that $t \geq \tau(x + 1)$, seller i knows that he will not sell at any later period. This is true because τ is strictly increasing if $x < \underline{x}$. Therefore, we have that $b_i(x, t) = \theta(x)$ and $V_i(x, t) = 0$ for all states (x, t) such that $t \geq \tau(x + 1)$. This gives us

that $V_j(x,t) = \theta_j - \theta(x) + V_j(x, t+1)$ for each (x,t) with $t \geq \tau(x+1)$, which, in turn, implies that $V_j(x,t) = w_j(x,t)$ for all $t \geq \tau(x+1)$.

Step 3. So far, we have taken the function τ as given. What we will do now is to construct such function (given the technology parameters θ and θ_j). As a by-product, we will also obtain the values of V_i and V_j for all states (x,t) such that $\tau(x) \leq t < \tau(x+1)$ (and $x < \underline{x}$).

Define the auxiliary function z as

$$(\text{B.1}) \qquad z(x,t) := \sum_{k=0}^{T-t} [T + 1 - (t+k)][\theta(x+k) - \theta_j].$$

The following two properties of z will be used below

$$(\text{B.2}) \qquad z(x,t) = z(x+1, t+1) - w_j(x,t)$$

$$(\text{B.3}) \qquad = z(x, t+1) + w_i(x,t).$$

By definition, we have that $\tau(\underline{x} - 1) = t_0$ if, and only if, $b_j(\underline{x} - 1, t_0 + k) > b_i(\underline{x} - 1, t_0 + k)$ for $k = 1, \ldots, T - 1 - t_0$ and $b_i(\underline{x} - 1, t_0) \geq b_j(\underline{x} - 1, t_0)$. Let us start with $t_0 = T - 1$. At state $(\underline{x} - 1, T - 1)$ each seller knows that if he does not sell now, he will neither sell in the next (and last) period. Clearly, then, $V_i(\underline{x}, T) = w_i(\underline{x}, T)$ and $V_j(\underline{x} - 1, T) = -z(\underline{x} - 1, T)$, whereas $V_i(\underline{x} - 1, T) = V_j(\underline{x}, T) = 0$. Therefore, $b_i(\underline{x} - 1, T - 1) = \theta(\underline{x} - 1) + w_i(\underline{x}, T)$ and $b_j(\underline{x} - 1, T - 1) = \theta_j - z(\underline{x} - 1, T)$. We have thus that $b_i(\underline{x} - 1, T - 1) - b_j(\underline{x} - 1, T - 1) = \theta(\underline{x} - 1) - \theta_j + w_i(\underline{x}, T) + z(\underline{x} - 1, T)$. Because of (B.3) and the definition of w_i we know that

$$V_i(\underline{x} - 1, T - 1) = \max\{z(\underline{x} - 1, T - 1), 0\}$$
$$V_j(\underline{x} - 1, T - 1) = \max\{-z(\underline{x} - 1, T - 1), 0\}.$$

Note that $z(\underline{x} - 1, T - 1) \geq 0$ is a necessary and sufficient condition for seller i to be the winning seller at $(\underline{x} - 1, T - 1)$ (that is, for $\tau(\underline{x} - 1) = T - 1$). Otherwise, seller j would be the winning seller at $(\underline{x} - 1, T - 1)$. Repeating the same argument over and over again for earlier periods, we see that $\tau(\underline{x} - 1) = t_0$ if, and only if, (a) $z(\underline{x} - 1, t) < 0$ for all $t > t_0$ and (b) $z(\underline{x} - 1, t_0) \geq 0$.

Condition (a) can be simplified to $z(\underline{x} - 1, t_0 + 1) < 0$ since this inequality implies the rest. Let us show why. In the first place, we have that $z(x,t) < 0$ for all states $(\underline{x} - k, T - k - l)$ with $k = 1, \ldots, \underline{x}$ and $l = 1, \ldots, k$. This is because at these states all the summands in (B.1) are negative. Second, for all other states the last summand in (B.1) is nonnegative. This implies that z is decreasing in t for those states. In summary, $z(x,t) < 0$ implies that $z(x, t+k) < 0$ for $k = 1, \ldots, T - t$.

Given $\tau(x+1)$, we can compute $\tau(x)$ in the same fashion by backwards induction. By definition, we know that $\tau(x) = t_0$ if, and only if, $b_j(x, t_0 + k) > b_i(x, t_0 + k)$ for $k = 1, \ldots, \tau(x+1) - 1 - t_0$ and $b_i(x, t_0) \geq b_j(x, t_0)$. Let us start with $t_0 = \tau(x+1) - 1$. We know that $V_i(x+1, \tau(x+1)) = z(x+1, \tau(x+1))$ by our induction hypothesis. Clearly, $V_j(x, \tau(x+1)) = w_j(x, \tau(x+1))$ (Step 2) and $V_i(x, \tau(x+1)) = V_j(x+1, \tau(x+1)) = 0$. Therefore, $b_i(x, \tau(x+1) - 1) = \theta(x) + z(x+1, \tau(x+1))$ and $b_j(x, \tau(x+1) - 1) = \theta_j + w_j(x, \tau(x+1))$. We have thus that $b_i(x, \tau(x+1) - 1) - b_j(x, \tau(x+1) - 1) = \theta(x) - \theta_j - w_j(x, \tau(x+1)) + z(x+1, \tau(x+1))$. Because of (B.2) and the definition of w_j we know that

$$V_i(x, \tau(x+1) - 1) = \max\{z(x, \tau(x+1) - 1), 0\}$$
$$V_j(x, \tau(x+1) - 1) = \max\{-z(x, \tau(x+1) - 1), 0\}.$$

Note that $z(x, \tau(x+1) - 1) \geq 0$ is a necessary and sufficient condition for seller i to be the winning seller at $(x, \tau(x+1) - 1)$ (that is, for $\tau(x) = \tau(x+1) - 1$). Otherwise, seller j would be the winning seller at that state. Repeating the same argument over and over again for earlier periods, we see that $\tau(x) = t_0$ if, and only if, (a) $z(x, t_0 + 1) < 0$ and (b) $z(x, t_0) \geq 0$. (Note that we have used the result of the previous paragraph).

In summary, we have shown that

$$\tau(x) = \max\{t : z(x, t) \geq 0\}$$

if $\{t : z(x, t) \geq 0\}$ is not empty, and $\tau(x) = x - 1$ otherwise.

Step 4. Besides constructing τ, we have also shown that:

(1) If $\tau(x) < t < \tau(x+1)$, then $V_i(x, t) = 0$ and $V_j(x, t) = -z(x, t)$.
(2) $V_i(x, \tau(x)) = z(x, \tau(x))$ and $V_j(x, \tau(x)) = 0$.

Because z is increasing in t for $t \leq \tau(x+1)$ (and $x < \underline{x}$)—see Step 3—we have that $z(x+1, \tau(x+1) - k) \geq 0$ for $k = 1, \ldots, \tau(x+1) - x - 1$. This implies that for all states (x, t) with $\tau(x) < t < \tau(x+1)$ (and $x < \underline{x}$), $z(x+1, t+1) \geq 0$, which, through (B.2), implies that $-z(x, t) \leq w_j(x, t)$. On the other hand, since $z(x, \tau(x) + 1) < 0$ we have, through (B.3), that $z(x, \tau(x)) < w_i(x, \tau(x))$. Therefore, the value functions V_i and V_j are bounded above by the functions w_i and w_j.

Step 5. We saw in Step 1 that the technology parameters uniquely determine the pay-offs for all states (x, t) with $x \geq \underline{x}$. We have seen (Step 3) that the technology parameters uniquely determine the function τ through the auxiliary function z, and that z determines also the values V_i and V_j for all states (x, t) such that $\tau(x) \leq t < \tau(x+1)$. In Step 2 we proved that, in combination with τ, the technology parameters

determine the pay-offs for states (x, t) such that $t \geq \tau(x + 1)$. This is also true for the states fulfilling $t < \tau(x)$ (and $x < \underline{x}$), as we saw in Step 3 of the proof of Proposition 3.3. Therefore, the equilibrium is unique. $\qquad\square$

REMARK B.1. We can summarize what we have learned about V_i and V_j as follows:
- (1) If $z(x, t) \geq 0$ $(t \leq \tau(x))$, then $V_j(x, t) = 0$ and:
  - (1a) If $z(x, t + 1) \geq 0$ $(t < \tau(x))$, then $V_i(x, t) = w_i(x, t) \leq z(x, t)$.
  - (1b) If $z(x, t + 1) < 0$ $(t = \tau(x))$, then $V_i(x, t) = z(x, t) < w_i(x, t)$.
- (2) If $z(x, t) < 0$ $(t > \tau(x))$, then $V_i(x, t) = 0$ and:
  - (2a) If $z(x + 1, t + 1) < 0$ $(t > \tau(x + 1))$, then $V_j(x, t) = w_j(x, t) < -z(x, t)$.
  - (2b) If $z(x+1, t+1) \geq 0$ $(\tau(x) < t \leq \tau(x+1))$, then $V_j(x, t) = -z(x, t) \leq w_j(x, t)$.

Or, in a more compact but less informative fashion:
- (i) If $z(x, t) \geq 0$, then $V_i(x, t) = \min\{z(x, t), w_i(x, t)\}$ and $V_j(x, t) = 0$.
- (ii) If $z(x, t) < 0$, then $V_i(x, t) = 0$ and $V_j(x, t) = \min\{-z(x, t), w_j(x, t)\}$.

PROPOSITION 3.5. *In a MPE the new technology is either adopted at $t = 0$ or never.*

PROOF. The value for seller i from adopting the new technology at date t is

(B.4) $$t \max\{\theta_i - \theta_j, 0\} + V_i(0, t),$$

whereas the value from never adopting is $(T + 1) \max\{\theta_i - \theta_j, 0\}$. We have assumed that the new technology is adopted at t only if $V_i(0, t) > 0$. First, note that if $z(0, 0) < 0$ the new technology will never be adopted because this implies that $V_i(0, t) = 0$ for all t. Second, if $V_i(0, 0) = z(0, 0)$, we also know that $V_i(0, t) = 0$ for all $t \geq 1$. Therefore, in this case we have that the new technology will be adopted at $t = 0$ if

$$z(0, 0) > (T + 1) \max\{\theta_i - \theta_j, 0\},$$

or it will never be adopted.

If $V_i(0, 0) = w_i(0, 0)$ and $\theta_i \leq \theta_j$, (B.4) is bounded above by $w_i(0, t)$, which is itself strictly smaller than $w_i(0, 0)$ if $t > 0$, meaning that the new technology will be adopted at $t = 0$. Suppose now that $V_i(0, 0) = w_i(0, 0)$ and $\theta_i > \theta_j$. We know that (B.4) is bounded above by

$$t(\theta_i - \theta_j) + w_i(0, t),$$

which is equal to

$$t\theta_i + \sum_{k=0}^{T-t} \theta(k) - (T+1)\theta_j.$$

This is strictly smaller than $w_i(0,0)$ because of (3.1) and (3.3). Therefore, the new technology is adopted $t = 0$. Note that we have already proved Proposition 3.6. $\square$

PROPOSITION 3.8. *If* $g > (T+1)(\max_s\{\theta_s\} + \frac{T}{2}\theta_j)$, *then all technologies in* $\Gamma(g)$ *are adopted in equilibrium.*

PROOF. Consider the optimization problem

$$\min_{\theta \in \Gamma(g)} Z(\theta).$$

The minimum is attained by the technology for which all learning takes places at the last period:

$$\theta(k) = \begin{cases} 0 & \text{if } k = 0, \ldots, T-1 \\ g & \text{if } k = T. \end{cases}$$

For this technology, Z is equal to

$$Z(\theta) = g - \frac{T+2}{2}(T+1)\theta_j,$$

which is larger than $(T+1)\max\{\theta_i - \theta_j, 0\}$ (see Proposition 3.6) if, and only if,

$$g > (T+1)\left(\max_s\{\theta_s\} + \frac{T}{2}\theta_j\right).$$

$\square$

PROPOSITION 3.9. *All classes* $\Gamma(g)$ *with* $g \leq (T+1)(\max_s\{\theta_s\} + \frac{T}{2}\theta_j)$ *exhibit the impatience property (excluding* $g = (T+1)\max_s\{\theta_s\}$ *if* $T = 1$).

PROOF. From Proposition 3.8 we know that we can always find a technology which is not adopted. Now consider the optimization problem

$$\max_{\theta \in \Gamma(g)} Z(\theta).$$

The maximum is attained by the technology for which all learning occurs between $t = 0$ and $t = 1$:

$$\theta(k) = \begin{cases} \theta_j & \text{if } k = 0 \\ \frac{g-\theta_j}{T} & \text{if } k = 1, \ldots, T. \end{cases}$$

For this technology, Z is equal to

$$Z(\theta) = \frac{1}{2}(T+1)[g - (T+1)\theta_j],$$

which, because of (3.3), is bounded below by

$$\frac{1}{2}(T+1)^2(\max_s\{\theta_s\} - \theta_j).$$

This is strictly larger than $(T+1)\max\{\theta_j - \theta_i, 0\}$ if $T > 1$. If $T = 1$, the class $\Gamma((T+1)\max_s\{\theta_s\})$ is excluded. $\qquad\square$

Bibliography

Abreu, D. and M. K. Brunnermeier (2003). Bubbles and Crashes. *Econometrica 71*(1), 173–204.

Arrow, K. J. (1962). Economic Welfare and the Allocation of Resources for Inventions. In R. Nelson (Ed.), *The Rate and Direction of Inventive Activity*. Princeton, New Jersey: Princeton Universuty Press.

Bergemann, D. and J. Välimäki (2006). Dynamic Price Competition. *Journal of Economic Theory 127*(1), 232 – 263.

Cabral, L. M. B. and M. H. Riordan (1994). The Learning Curve, Market Dominance, and Predatory Pricing. *Econometrica 62*(5), pp. 1115–1140.

Calvo, G. A. (1998). Contagion in Emerging Markets: When Wall Street is a Carrier. University of Maryland.

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

Chari, V. V. and H. Hopenhayn (1991). Vintage Human Capital, Growth, and the Diffusion of New Technology. *The Journal of Political Economy 99*(6), pp. 1142–1165.

Corsetti, G., A. Dasgupta, S. Morris, and H. S. Shin (2004). Does One Soros Make a Difference? A Theory of Currency Crises with Large and Small Traders. *The Review of Economic Studies 71*(1), pp. 87–113.

Diamond, D. W. and P. H. Dybvig (1983). Bank Runs, Deposit Insurance, and Liquidity. *The Journal of Political Economy 91*(3), pp. 401–419.

Friedman, M. (1953). The Case for Flexible Exchange Rates. In *Essays in Positive Economics*. University of Chicago Press.

Galbraith, J. K. (1994). *A Short History of Financial Euphoria*. Penguin Books.

Gilbert, R. J. and D. M. G. Newbery (1982). Preemptive Patenting and the Persistence of Monopoly. *The American Economic Review 72*(3), pp. 514–526.

Goldstein, I. and A. Pauzner (2005). Demand-Deposit Contracts and the Probability of Bank Runs. *The Journal of Finance 60*(3), 1293–1327.

Greenwood, R. and S. Nagel (2009). Inexperienced Investors and Bubbles. *Journal of Financial Economics 93*(2), 239 – 258.

Holmes, T. J., D. K. Levine, and J. A. Schmitz (2008). Monopoly and the Incentive to Innovate when Adoption Involves Switchover Disruptions. *NBER Working Paper*.

Holmes, T. J. and J. A. Schmitz (2010). Competition and Productivity: A Review of Evidence. Research Department Staff Report 439, Federal Reserve Bank of Minneapolis.

Jovanovic, B. and Y. Nyarko (1996). Learning by Doing and the Choice of Technology. *Econometrica 64*(6), pp. 1299–1310.

Kaminsky, G., S. Lizondo, and C. Reinhart (1998). Leading Indicators of Currency Crises. *Staff Papers-International Monetary Fund 45*(1), 1–48.

Kindleberger, C. P. and R. Z. Aliber (2005). *Manias, Panics and Crashes: A History of Financial Crises* (fifth ed.). Palgrave Macmillan.

Morris, S. and H. S. Shin (1998). Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks. *The American Economic Review 88*(3), 587–597.

Obstfeld, M. (1996). Models of Currency Crises with Self-Fulfilling Features. *European Economic Review 40*(3-5), 1037 – 1047. Papers and Proceedings of the Tenth Annual Congress of the European Economic Association.

Parente, S. L. and E. C. Prescott (2002). *Barriers to Riches.* Walras-Pareto Lectures. The MIT Press.

Shiller, R. J. (2005). *Irrational Exuberance* (second ed.). Princeton University Press.

Shleifer, A. (2000). *Inefficient Markets: An Introduction to Behavioral Finance.* Clarendon Lectures in Economics. Oxford University Press.