

Bubbles and Crashes*

Dilip Abreu[†]
Department of Economics
Princeton University

Markus K. Brunnermeier[‡]
Department of Economics
Princeton University

first version: October 2000
this version: November 2001

Abstract

We present a model in which an asset bubble can persist despite the presence of rational arbitrageurs. The resilience of the bubble stems from the inability of arbitrageurs to temporarily coordinate their selling strategies. This *synchronization problem* together with the individual incentive to *time the market* results in the persistence of bubbles over a substantial period of time. Since the derived trading equilibrium is unique, our model rationalizes the existence of bubbles in a strong sense. The model also provides a natural setting in which public events, by enabling synchronization, can have a disproportionate impact relative to their intrinsic informational content.

Keywords: Bubbles, Crashes, Temporal Coordination, Synchronization, Market Timing, ‘Overreaction’, Limits to Arbitrage, Behavioral Finance

*We thank especially Stephen Morris, Ming Huang, José Scheinkman, Andrei Shleifer, Haluk Ergin and seminar participants at the London School of Economics, Stanford University, UC Berkeley, the Stockholm School of Economics, Harvard University, the Wharton School, NBER Behavioral Finance Conference at the University of Chicago, CEPR financial markets meetings in Gerzensee, Cornell University, Northwestern University, Columbia University and New York University for helpful comments. The authors acknowledge financial support from National Science Foundation and Princeton’s CEPS program, respectively.

[†]Department of Economics, Princeton University, Princeton, NJ 08544-1021, e-mail: dabreu@princeton.edu, <http://www.princeton.edu/~dabreu>

[‡]Department of Economics, Bendheim Center for Finance, Princeton University, Princeton, NJ 08544-1021, e-mail: markus@princeton.edu, <http://www.princeton.edu/~markus>

1 Introduction

Share prices of internet stocks and telecommunication stocks reached extremely high values in March 2000. In the subsequent months, the CBOE Internet Index lost more than 75 % of its value. Dramatic price increases followed by sharp price drops are not a new phenomena. Famous bubbles include the Dutch Tulip Mania of the 1630's, the Mississippi Bubble and the South Sea Bubble of the 1720's.

These episodes are hard to reconcile with the predictions of standard neoclassical economic theory wherein *all* market participants are assumed to be fully rational. This theory predicts that rational traders will typically eliminate a bubble before it has a chance to develop. Bubbles are ruled out by a classic backwards induction argument in finite horizon models and by the transversality condition in infinite horizon models. This counterfactual prediction of rational models has led some scholars like Kindleberger (1978) to emphasize the role of irrational behavior, and indeed even adopt the opposite viewpoint that all market participants behave irrationally during such episodes.

Proponents of the efficient market hypothesis, such as Fama (1965) are quite willing to admit that behavioral/boundedly rational traders are active in the market place. However, they argue that the existence of sufficiently many well-informed and well-financed arbitrageurs guarantees that any potential mispricing induced by behavioral traders will be corrected. Hence, the efficient market theory also rules out the persistence of bubbles.

Our objective is to examine the validity of the efficient market perspective. In particular, we investigate whether asset bubbles can survive in the presence of rational arbitrageurs. Our conclusions suggest that arbitrage 'ultimately' works, though it might be ineffectual over substantial periods. In our setting, a bubble persists even though rational agents know that the bubble must burst with probability one in finite time.

Imagine a world in which there are *some* 'behavioral' agents variously subject to animal spirits, fads and fashions, overconfidence and related psychological biases which might lead to momentum trading, trend chasing and the like. There is by now a large literature which documents and models such behavior.¹ We do not investigate why behavioral biases needing rational corrections arise in the first place; for this we rely on the body of work very partially documented in footnote 1. We study the impact of rational arbitrage in this setting.

Suppose that the bubble asset price grows at a faster rate than the risk free interest rate, until the bubble bursts. Rational arbitrageurs know that eventually the market will collapse but meanwhile would like to ride the bubble as it continues to grow. Ideally, they would like to exit the market just prior to the crash. However, market timing

¹Cognitive biases are illustrated in books and articles, such as Daniel, Hirshleifer, and Subrahmanyam (1998), Hirshleifer (2001), Odean (1998), Thaler (1991), Shiller (2000), and Shleifer (2000).

is a difficult task. Our investors understand that they will, for a variety of reasons, come up with different solutions to this optimal timing problem. This dispersion of exit strategies and the consequent lack of synchronization are precisely what permit the bubble to thrive and grow. This is despite the fact that the bubble bursts as soon as a sufficient mass of traders sells out.

We present a model which formalizes the above synchronization problem and yields a new perspective on the existence, persistence, and collapse of bubbles. Our approach emphasizes two elements; *dispersion of opinion* among rational arbitrageurs and the need for *coordination*.

We assume that the price surpasses the fundamental value at a random point in time t_0 . Thereafter, arbitrageurs become sequentially aware of the fact that the price surpassed the fundamental value. Arbitrageurs do not know whether they have learnt this information early or late relative to other rational arbitrageurs. Since they become sequentially aware of the mispricing, their trading strategies are initialized at different random times. In addition to its literal interpretation the assumption of sequential awareness can be viewed as a metaphor for a variety of factors such as dispersion of opinion and information which, in particular, find expression in the factor we seek to explore and emphasize, that is, temporal miscoordination. Arbitrageurs have common priors about the underlying structure of the model.

Second, in our model selling pressure only bursts the bubble when a sufficient mass of arbitrageurs have sold out. Arbitrageurs face financial constraints which limit their stock holding as well as their maximum short-position. This limits the price impact of each arbitrageur. Large price movements can only occur if the accumulated selling pressure exceeds some threshold κ . In other words, a permanent shift in price levels requires a coordinated attack. In this respect, our model shares some features with the static second generation models of currency attacks in the international finance literature (Obstfeld 1996). However, these currency attack models focus exclusively on the question of *whether* to attack or not but ignore the important *temporal* aspect of coordination. Coordinating on a given action is complicated by the need to coordinate both the action and the time at which it is taken. Thus, speculators need to decide both whether or not to attack - the problem that is traditionally emphasized in the currency attack literature - and also 'when' to attack. It would be futile to simply coordinate on the 'whether' if it were not possible to also coordinate on the 'when'. This temporal element exacerbates the underlying coordination problem.

Our model has both elements of cooperation and of competition. On the one hand, at least a fraction κ of arbitrageurs need to sell out in order for the bubble to burst; this is the coordination aspect. On the other hand, arbitrageurs are competitive since at most a fraction $\kappa < 1$ of them can leave the market prior to the crash.²

In the equilibrium of our model, arbitrageurs stay in the market until the subjective

²Potential further applications of the dynamic coordination game developed in this paper include currency attacks, bank-runs and investments games.

probability that the bubble will burst in the next trading round is sufficiently high. Arbitrageurs who get out of the market just prior to the crash make the highest profit. Arbitrageurs who leave the market very early make some profit, but forgo much of the higher rate of appreciation of the bubble. Arbitrageurs who stay in the market too long lose all capital gains which result from the bubble's appreciation.

The result on market timing seems to fit well with popular accounts of the behavior of hedge fund managers during the recent internet bubble.³ For example, when Stanley Druckenmiller, who managed George Soros' \$8.2 billion Quantum Fund, was asked why he didn't get out of internet stocks earlier even though he knew that technology stocks were overvalued, he replied that he thought the party wasn't going to end so quickly. In his words "We thought it was the eighth inning, and it was the ninth."⁴ Faced with mounting losses, Druckenmiller resigned as Quantum's fund manager in April 2000.

However, fund managers can also not afford to simply stay away from a rapidly growing bubble. Julian Roberts, manager of the legendary Tiger Hedge Fund, refused to invest in technology stocks since he thought they were overvalued. The Tiger Fund was dissolved in 1999 because its returns could not keep up with the returns generated by dotcom stocks. A Wall Street analyst who has dealt with both hedge fund managers vividly summarized the situation: "Julian said, 'This is irrational and I won't play,' and they carried him out feet first. Druckenmiller said, 'This is irrational and I will play,' and they carried him out feet first."

Rational arbitrageurs ride the bubble even though they know that the bubble will burst for exogenous reasons by some time $t_0 + \bar{\tau}$ if it has not succumbed to endogenous selling pressure prior to that time. Here t_0 is the unknown time at which the price path surpasses the fundamental value and arbitrageurs start getting aware of the mispricing. Our analysis shows that if arbitrageurs' opinion is sufficiently dispersed, there exists an equilibrium in which the bubble never bursts prior to $t_0 + \bar{\tau}$. Even long after the bubble begins and after all agents are aware of the bubble, it is nevertheless the case that endogenous selling pressure is never high enough to burst the bubble. Moreover, this equilibrium is *unique*.⁵ Thus, there is a striking failure of the backwards induction argument which would yield immediate collapse in a standard model. The persistence of the bubble in our model relies on dispersion in traders' viewpoints about when the bubble emerged. Presumably, this dispersion is specially high at times of significant

³Kindleberger (1978) notes that even Isaac Newton tried to ride the South Sea Bubble in 1720. He got out of the market at £7,000 after making a £3,500 profit, but he decided to re-enter it thereby losing £20,000 at the end. Frustrated with his experience, he concluded: "I can calculate the motions of the heavenly bodies, but not the madness of people."

⁴New York Times, April 29, 2000, "Another Technology Victim; Soros Fund Manager Says He 'Overplayed' Hand".

⁵We restrict attention to perfect Bayesian equilibria in which when an agent sells out, all agents who become aware of the bubble prior to the agent in question, are also out of the market. We view this "monotonicity" restriction as both natural and innocuous in the context of the issues we seek to investigate.

technological changes, such as the invention of the steam boat, telegraph, internet, etc. Structural breaks such as large scale financial liberalization programs can be another breeding ground for bubbles, as illustrated by the economic developments in the Scandinavian countries during the late eighties.

Second, we show that while arbitrageurs never burst a bubble if their opinions are sufficiently dispersed or if the absorption capacity by the behavioral momentum traders κ is very large, for smaller dispersion of opinion or for smaller κ , endogenous selling pressure advances the date at which the bubble eventually collapses. Nevertheless, the bubble grows for a substantial period of time. Again, we also show that the symmetric equilibrium where each arbitrageur sells out after waiting for a certain number of periods after he becomes aware of the mispricing is unique.

The model also provides a natural setting in which news events can have a disproportionate impact relative to their intrinsic informational content. This is because news events make it possible for agents to synchronize their exit strategies. Of course, large price drops are themselves significant public events, and we investigate how an initial price drop may lead to a full-fledged collapse. Thus the model yields a rudimentary theory of ‘overreaction’ and price ‘cascades’ and suggests a rationale for psychological benchmarks such as ‘resistance lines’. In addition, our model provides a framework for understanding fads in information such as the (over-)emphasis on trade figures in the eighties and on interest rates in the nineties. Finally, our model supports arguments in favor of centralized news dissemination since news which is received sequentially over a long interval is much less likely to be reflected in the price.

Overall, the idea that the bursting of a bubble requires synchronized action by rational arbitrageurs, who might lack both the incentive and the ability to act in a coordinated way, has important implications. It provides theoretical support for empirical observations on the existence and the pervasiveness of bubbles. It undermines the central presumption of the efficient market perspective that not all agents need to be rational for prices to be efficient. Finally, our model provides further support for behavioral finance models which do not explicitly model rational arbitrageurs.

The remainder of the paper is organized as follows. Section 2 illustrates how the analysis relates to the literature. In Section 3 we introduce the primitives of the model and define the equilibrium. Section 4 shows that all arbitrageurs employ trigger strategies in any trading equilibrium. Section 5 demonstrates that if the dispersion of opinion among arbitrageurs is sufficiently large, they never burst the finite horizon bubble and it only crashes for exogenous reasons at its maximum size. For smaller dispersion of opinion the bubble also persists but arbitrageurs burst it before it reaches its maximum size. This section clarifies why a lack of common knowledge leads to a failure of backwards induction. Section 6 highlights the special role of public events and discusses the fragility of bubbles with respect to different forms of public events. Section 7 focuses on price cascades and market rebounds. Section 8 concludes.

2 Related Literature

Since finite horizon bubbles can be easily ruled out by backwards induction in a symmetric information setting with rational traders, most bubble models consider an infinite horizon setup.⁶ Blanchard and Watson (1982) present a particularly tractable formulation, in which the bubble bursts with a constant probability in each trading period for exogenous reasons. Like the value of any non-dividend paying asset, the bubble ought to grow in expectations in perpetuity. However, Santos and Woodford (1997) demonstrated that growing bubbles can be ruled out because they violate the transversality condition that results from the agents' optimization problem. They conclude that bubbles cannot arise in a setting with symmetric information with the exception of a few stylized cases.

Tirole (1982) shows that bubbles can also be ruled out in a rational expectations model where risk-neutral traders hold different pieces of information. Risk-neutrality guarantees that any allocation is interim Pareto efficient. This makes trading a zero-sum game and hence buying an overvalued 'bubble' asset is a negative-sum game. Allen, Morris, and Postlewaite (1993) highlight three necessary conditions for bubbles in addition to the interim Pareto inefficiency requirement. They call a mispricing a bubble if all traders know that the price is too high. However, mutual knowledge of the bubble does not imply that all traders know that all traders know that the price exceeds its fundamental value. They provide illustrative examples which satisfy their conditions and which support bubbles. In Allen and Gorton (1993) fund managers "churn bubbles." They take on an overvalued asset even though they know that they might be last in line and hence unable to unload the asset. While they share in positive profits they do not in losses, leading to a positive sum game for managers at the expense of client investors.

All the papers described above assume that all agents are fully rational. In contrast, our model falls into the class of models wherein rational arbitrageurs interact with boundedly rational behavioral traders in the market place.⁷ In DeLong, Shleifer, Summers, and Waldmann (1990b) rational arbitrageurs push up the price after some initial good news in order to induce behavioral feedback traders to aggressively buy stocks in the next period. This delayed reaction by feedback traders allows the arbitrageurs to unload their position at a profit. In DeLong, Shleifer, Summers, and Waldmann (1990a) arbitrageurs' risk aversion and short horizons limit their ability to correct the mispricing. In contrast, in our model arbitrageurs initially do not even attempt to lean

⁶A more extensive review of the literature on bubbles can be found in Brunnermeier (2001).

⁷There is also a growing literature that tries to explicitly model the trading patterns of behavioral traders. This goes beyond the scope of the current paper since our focus is on the inability of rational arbitrageurs to correct any mispricing caused by behavioral traders. The interested reader is referred to Barberis and Thaler (forthcoming) and Hirshleifer (2001) for an overview of this literature as well as to books by Thaler (1991), Shiller (2000), and Shleifer (2000) among others.

against the mispricing even though they are risk-neutral and infinitely lived. In Shleifer and Vishny (1997), professional fund managers forgo profitable long-run arbitrage opportunities because the price might depart even further from the fundamental value in the intermediate term. In that case, the fund manager would have to report intermediate losses causing client investors to withdraw part of their money which forces him to liquidate at a loss. Knowing this might happen in advance, the fund manager only partially exploits the arbitrage opportunity.

All these papers build on the insight that rational arbitrageurs do not have the collective ability to correct a mispricing either because of their risk aversion or because of other exogenously assumed capital constraints. In contrast in our paper, the aggregate resources of all arbitrageurs is sufficient to bring the price back to its fundamental value. Thus, the weakness of arbitrage in our model is particularly striking because arbitrageurs can jointly correct the mispricing, but it nevertheless persists.

Our assumption that a critical mass of speculators is needed to burst a bubble has its roots in the currency attack literature. Obstfeld (1996) highlights the necessity of coordination among speculators to break a currency peg and points out the resulting multiplicity of equilibria in a setting with symmetric information. Morris and Shin (1998) introduce asymmetric information and derive a unique equilibrium by applying the *global games* approach of Carlsson and van Damme (1993). Both currency attack models are static in the sense that speculators only decide whether to attack now or never.⁸ Our model is, of course dynamic. The dispersion of opinion among arbitrageurs, which we model as sequential awareness, can be related to the notion of “asynchronous clocks.” Asynchronous clocks were first introduced in the computer science literature by Halpern and Moses (1990). Morris (1995) applies the concept to a dynamic coordination problem in labor economics. There are several differences between our papers. His model satisfies strategies complementarity and the global games approach applies. Our model has elements of both coordination and competition. The latter leads to a pre-emption motive which plays a central role in our analysis. It is important for his result that players can only condition on the individual clocks and not on calendar time or on their payoffs, whereas in our model they can. In particular,

⁸There are also other dynamic coordination failure games without the preemption motive. They focus on the elimination of multiple equilibria. Chamley (1999) is able to derive a unique equilibrium in a dynamic setting with asymmetric information by allowing agents to condition on past outcomes. In contrast to our model, agents live only for one period in Chamley (1999) and hence do not face a timing decision problem. A key feature of Burdzy, Frankel, and Pauzner (2001) and Frankel and Pauzner (2000) - the latter within the specific framework of Matsuyama’s (1991) dynamic two sector model - is that agents are subject to frictions; in particular they can only change their actions at random times. This inertia, together with other assumptions regarding permanent payoff shocks generate unique equilibrium behavior often by iterative dominance arguments. Asymmetric information typically does not play a role in the latter papers. Our own work relies on asymmetric information; however, the asymmetric information primarily concerns the temporal structure of the model and leads to a synchronization problem.

traders learn from the existence of the bubble in our setting. In the richer strategy set of our model strategic complementarity is not satisfied. See footnote 18 for a discussion of this point.

Finally we note that some of the key elements of our model echo themes from Keynes (1936). The connections are elaborated upon in the discussion following the presentation of the model setup.

3 Model

3.1 Model Setup

Historically bubbles have often emerged in periods of productivity enhancing structural change. Examples include the railway boom, the electricity boom, and the recent internet and telecommunication boom. In the latter case, and in many of the historical examples, sophisticated market participants gradually understood that the immediate economic impact of these structural changes was limited and that their full implementation would take a long time. They also realized that only a few of the companies engaged in the new technology would survive in the long run. On the other hand, less sophisticated traders over-optimistically believed that a ‘paradigm shift’ or a ‘new economy’ would lead to permanently higher growth rates.

We assume the price process depicted in Figure 1.

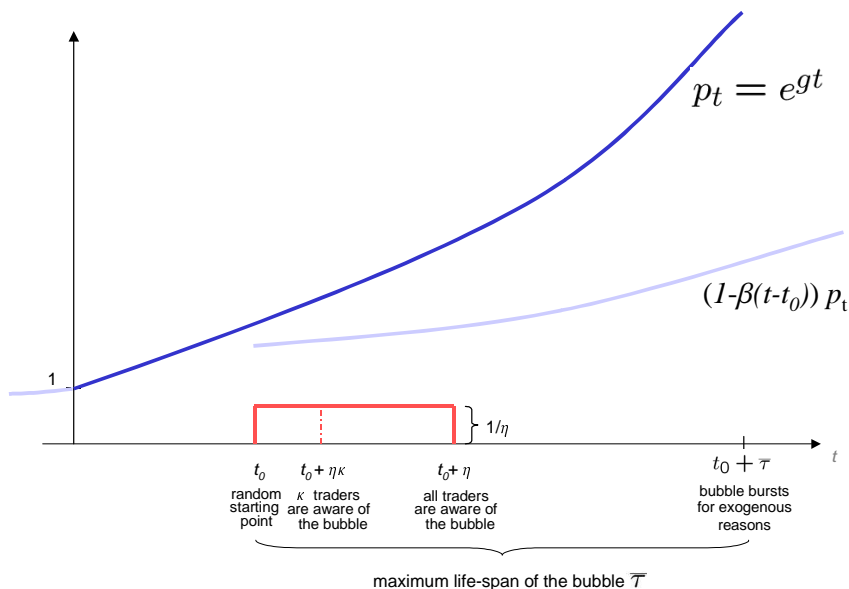


Figure 1: Illustration of price paths

This price process reflects the scenario outlined above, and may be interpreted as follows. Prior to $t = 0$ the stock price index coincides with its fundamental value

which grows at the risk-free interest rate r and rational arbitrageurs are fully invested in the stock market. Without loss of generality, we normalize the starting point to $t = 0$ and the stock market price at $t = 0$ to $p_0 = 1$. From $t = 0$ onwards the stock price p_t grows at a rate of $g > r$, that is $p_t = e^{gt}$.⁹ This higher growth rate may be viewed as emerging from a series of unusual positive shocks which gradually make investors more and more optimistic about future prospects. Until some random time t_0 , the higher price increase is justified by the fundamental development. We assume that t_0 is exponentially distributed on $[0, \infty)$ with the cumulative distribution function $\Phi(t_0) = 1 - e^{-\lambda t_0}$.¹⁰ Nevertheless, the price continues to increase at the faster rate g even after t_0 . Hence, from t_0 onwards, only the fraction $(1 - \beta(\cdot))$ of the price is justified by the fundamentals, while the fraction $\beta(\cdot)$ reflects the “bubble component.” We assume that $\beta(\cdot) : [0, \bar{\tau}] \mapsto [0, \bar{\beta}]$ is a strictly increasing and continuous function of $t - t_0$, the time elapsed since the price departed from fundamentals. For the special case where the fundamental value $e^{gt_0+r(t-t_0)}$ always grows at a rate r from t_0 onwards, $\beta(t - t_0) = 1 - e^{-(g-r)(t-t_0)}$.

The price $p_t = e^{gt}$ is kept above its fundamental value by irrationally exuberant behavioral traders. They believe in a “new economy paradigm” and think that the price will grow at a rate of g in perpetuity.¹¹ As soon as the rational arbitrageur’s selling pressure exceeds κ , the absorption capacity of momentum traders, the price drops by a fraction $\beta(\cdot)$ to its post-crash price. From this point onwards, the price grows at a rate of r . In other words, the bubble bursts as soon as a fraction κ of arbitrageurs sell out their stock holdings.¹² Even if the selling pressure never exceeds κ we assume that the bubble bursts for exogenous reasons as soon as it reaches its maximum size $\bar{\beta}$. This translates to a final date, since $\beta(\cdot)$ is strictly increasing. Let us denote this date by $t_0 + \bar{\tau}$. Note that this assumption of a final date is arguably the least conducive to the persistence of bubbles. In a classical model it would lead to an immediate collapse for the usual backwards programming reasons.

Another important element of our analysis is that rational arbitrageurs become sequentially informed that the fundamental value has not kept up with the growth of the stock price index. More specifically, a new cohort of rational arbitrageurs of mass

⁹We will illustrate later that the analysis can easily be extended to a setting with a stochastic price process.

¹⁰Kindleberger (1978) refers to the phase of fundamentally good news as ‘displacement.’ We focus on the consecutive subgame.

¹¹The price process we assume is a modeling simplification which facilitates a clean and simple analysis. An explicit derivation would take us too far afield. One, admittedly simplistic, setting which rationalizes this process is a world with behavioral traders who are risk-neutral, but wealth constrained and who require a rate of return of g in order to invest in the stock market. Their wealth constraint limits their aggregate absorption capacity to κ .

¹²We refer to fundamental value as the price which will emerge after the bubble bursts. Notice that for $\beta(0) = 0$, there is no drop in fundamentals at t_0 . Nevertheless, the trader who becomes aware of the bubble at t_0 , thinks that the current price is too high, since he does not know that she is the first one, who hears of the mispricing.

$\frac{1}{\eta}$ becomes ‘aware’ of the mispricing in each instant t from t_0 until $t_0 + \eta$. That is, $[t_0, t_0 + \eta]$ forms the ‘awareness window.’ Since t_0 is random, an individual arbitrageur does not know how many other arbitrageurs have received the signal before or after her. An agent who becomes aware of the bubble at t_i has a posterior distribution for t_0 with support $[t_i - \eta, t_i]$. Each agent views the market from the relative perspective of her own t_i . Viewed more abstractly, arbitrageurs’ types are determined by $t_i \in [0, \infty)$, the date when they become aware of the bubble. We denote by \hat{t}_i the arbitrageur who learns of the mispricing at time t_i . Figure 2 depicts the distributions of t_0 for arbitrageurs \hat{t}_i , \hat{t}_j and \hat{t}_k .




Figure 2: Sequential Awareness

Nature’s choice of t_0 determines the ‘active’ types $\{\hat{t}_i | t_i \in [t_0, t_0 + \eta]\}$ in the economy. As noted earlier, we view this specification as a modeling device which captures temporal miscoordination arising from differences of opinion and information. The date t_i at which agent \hat{t}_i becomes ‘aware’ of the mispricing may be more generally thought of as the date at which a player’s strategy is ‘initialized’. We assume that λ is sufficiently small such that, $\frac{\lambda}{1 - e^{-\lambda\eta\kappa}} < \frac{q-r}{\beta(\eta\kappa)}$. This guarantees that arbitrageurs do not wish to sell out before they become aware of the mispricing. When λ is very high the model is uninteresting, since arbitrageurs will sell out right away.

The price $p_t = e^{gt}$ exceeds its post-crash (fundamental) value from t_0 onwards. However, only a few arbitrageurs are aware of the mispricing at this point. From $t_0 + \eta\kappa$ onwards, the mispricing is mutual known to a large enough mass of arbitrageurs who collectively are able to correct it. We label any persistent mispricing beyond $t_0 + \eta\kappa$ as a *bubble*.

Given such an environment, we want to identify the best strategy of a rational player \hat{t}_i . Each arbitrageur can sell all or part of her stock holding or even go short till she reaches a certain limit where her financial constraint is binding. Each trader can also buy back shares. A trader may exit from and return to the market multiple times. However, each arbitrageur is limited in the number of shares each arbitrageur can go short or long. Without loss of generality, we can normalize the action space to be the continuum between $[0, 1]$, where 0 indicates the maximum long position and 1 the maximum short position each arbitrageur can take on. Given the focus of our analysis on selling pressure, this sign convention is convenient.

Let $\sigma(t, \hat{t}_i)$ denote the selling pressure of arbitrageur \hat{t}_i at t and the function $\sigma : [0, \infty) \times [0, \infty) \mapsto [0, 1]$ a strategy profile. The strategy of a trader who became aware

of the bubble at time t_i is given by the mapping $\sigma(\cdot, \hat{t}_i) : [0, t_i + \bar{\tau}] \mapsto [0, 1]$. Note that $[1 - \sigma(t, \hat{t}_i)]$ is trader \hat{t}_i 's stock holding at time t . For arbitrary σ , the function $\sigma(t, \cdot)$ need not be measurable in t_i . We confine attention to strategy profiles which guarantee a measurable function $\sigma(t, \cdot)$. Notice that no individual deviation from such a profile has any impact on the measurability property. In equilibrium, it will turn out that the $\sigma(t, \cdot)$ functions have an extremely simple structure. By the definition of a trading equilibrium, the set of agents with strictly positive sales is an interval. Furthermore, it will turn out that when agents sell shares they sell out completely. The aggregate selling pressure of all arbitrageurs at time $t \geq t_0$ is given by $s(t, t_0) = \int_{t_0}^{\min\{t, t_0 + \eta\}} \sigma(t, \hat{t}_i) dt_i$. Let $T^*(t_0) = \inf\{t | s(t, t_0) \geq \kappa \text{ or } t = t_0 + \bar{\tau}\}$ denote the bursting time of the bubble for a given realization of t_0 . Recall $\Phi(\cdot | t_i)$ denotes arbitrageur \hat{t}_i 's beliefs of t_0 given that $t_0 \in [t_i - \eta, t_i]$. Hence, his beliefs about the bursting date are given by

$$\Pi(t | \hat{t}_i) = \int_{T^*(t_0) < t} d\Phi(t_0 | \hat{t}_i).$$

Given the structure of the game, the actions of the other traders affect trader \hat{t}_i 's payoff only if these actions cause the bubble to burst. Each arbitrageur's payoff depends on the prices at which he sells and buys his shares minus the transactions costs for executing the order. The execution prices of arbitrageurs' orders are either the pre-crash price $p(t)$ or the post-crash price $(1 - \beta(t - t_0))p(t)$. In the special case where an arbitrageur submits his sell order exactly at the instant when the bubble bursts, his order is executed at the pre-crash price $p(t) = e^{gt}$ as long as the accumulated selling pressure is smaller than or equal to κ . If the accumulated selling pressure exceeds κ , then only the first randomly chosen orders will be executed at the pre-crash price while the remaining orders are only executed at the post-crash price. In other words, the *expected* execution price

$$(1 - \alpha)p(t) + \alpha(1 - \beta(t - t_0))p(t)$$

is a convex combination of both prices, where $\alpha > 0$ if the selling pressure is strictly larger than κ , and $\alpha = 0$ if the selling pressure is exactly κ at the time of the bursting of the bubble.

We assume that arbitrageurs incur transactions costs whenever they alter their positions. This implies that each arbitrageur will only change her stock holdings at most a finite number of times; strategies entailing an infinite number of changes will be strictly dominated. Consequently, we confine attention to strategies which involve at most an (arbitrary) finite number of changes of position. Indeed, this is the primary role that transactions costs play in the model. We rule out prohibitively high transactions costs which preclude arbitrageurs from ever selling out despite the fear of a bursting bubble. It is algebraically convenient to assume that the present value of transactions cost is a constant c . That is, transactions cost at t equals ce^{rt} . This formulation together with

the former assumption guarantees that equilibrium behavior is independent of c , the size of the transactions costs.

Before fully specifying the payoffs, let us define $B(t)$ the set of t_0 which lead to a bursting of the bubble strictly prior to t . That is, $B(t) = \{t_0 | T^*(t_0) < t\}$, $B^c(t)$ is its complement and $b(t) = \{t_0 | T^*(t_0) = t\}$ contains all t_0 which lead to a bursting exactly at t .

The value at time t^1 of a position consisting of $x^1(t^1)$ stocks which is maintained unchanged until the bubble bursts is (from the viewpoint of arbitrageur \hat{t}_i):

$$V^1((x^1, t^1) | t_i) = x^1 \int_{B(t_i + \bar{\tau})} e^{-r(T^*(t_0) - t^1)} [1 - \beta(T^*(t_0) - t_0)] p(T^*(t_0)) d\Phi(t_0 | t_i, B^c(t^1)),$$

where $\Phi(\cdot | t_i, B^c(t^1)) = \frac{\Phi(\cdot | t_i)}{1 - \Phi(t^1 | t_i)}$ is arbitrageur \hat{t}_i 's distribution over t_0 , conditional upon the time t_i when arbitrageur \hat{t}_i became aware of the mispricing, and conditional upon the existence of the bubble at t^1 .¹³

Proceeding iteratively from the final change of the position, the value at $t^2 < t^1$ of x^2 stocks of which $(x^2 - x^1)$ are sold at t^1 is given by

$$\begin{aligned} & V^2((x^2, t^2), (x^1, t^1) | t_i) \\ = & x^2 \int_{B(t^1) \setminus B(t^2)} e^{-r(T^*(t_0) - t^2)} [1 - \beta(T^*(t_0) - t_0)] p(T^*(t_0)) d\Phi(t_0 | t_i, B^c(t^2)) \\ & + [1 - \Phi(B(t^1) | t_i, B^c(t^2))] \{-e^{rt^2} c + \\ & \quad + e^{-r(t^1 - t^2)} [V^1((x^1, t^1) | t_i) + (x^2 - x^1) p(t^1) - A(t^1 | t_i, B^c(t^2))]\} \end{aligned}$$

The adjustment term $A(t^1 | t_i, B^c(t^2))$ takes care for the special case where the bubble bursts exactly at t^1 and the order of $(x^2 - x^1)$ shares is not executed at the price $p(t^1)$ but at $[1 - \alpha_t(t_0) \beta(t - t_0)] p(t)$. More precisely, $A(t^1 | t_i, B^c(t^2))$ is given by the expectations $(x^2 - x^1) p(t^1) \int_{b(t^1)} \alpha_t(t_0) \beta(t - t_0) d\Phi(t_0 | t_i, B^c(t^2))$.

Replacing superscript 2 by $k + 1$ and superscript 1 with superscript k gives us a general recursive definition of the payoff function. Arbitrageur \hat{t}_i 's expected payoff from a strategy involving $K - 1$ changes in her portfolio is then given by the function

$$V^K((x^K, t^K), \dots, (x^1, t^1)),$$

where $(x^K = 1, t^K = t_i)$ denotes her initial position.

¹³Section 4 shows that each arbitrageur's trading strategy involves only one change in asset position in any trading equilibrium. This dramatically simplifies the payoff specification. However, before we derive this 'trigger-property', we need to specify the payoffs for strategies involving many changes in asset positions.

3.2 Discussion of the Model

We now discuss and put into perspective some key elements of our model. The question we address - do professional arbitrageurs correct mispricing? - is a very old one and goes back to at least Keynes' "General Theory of Employment, Interest and Money" (1936) in which he wrote:

It might have been supposed that *competition between expert professionals*, possessing judgment and knowledge beyond that of the average private investor, *would correct* the vagaries of the ignorant individual left to himself. (*italics added*)

Our model setup relies crucially on two elements, which can also be traced back to Keynes (1936). First, professional arbitrageurs want to ride the bubble as long as possible given that the bubble grows at a rate g which is larger than the riskfree rate r . Professional traders do not want to forego the capital gains as long as the bubble grows, but, of course, seek to exit before (ideally just before) the crash. Referring back to Keynes:

The actual, *private* object of the most skilled investment to-day is "*to beat the gun*", as the Americans so well express it, to outwit the crowd, and to pass the bad, or depreciating, half-crown to the other fellow. (*italics added*)

The second element is that large price changes require a certain degree of coordination. In our model the bubble only bursts if more than κ traders attack the bubble at the same time. Therefore, each professional arbitrageur only attacks at times when she believes that the other arbitrageurs attack as well. Hence, it is more important to focus on other arbitrageurs' trading than on the fundamentals of an asset. This is in the same spirit as "Keynes' Beauty Contest":¹⁴

... professional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole; so that each competitor has to pick, not those faces which he himself finds prettiest, but those which he thinks likeliest to catch the

¹⁴Other authors have also found inspiration in this famous quote in Keynes (1936). Froot, Scharfstein, and Stein (1992) focus on the endogenous information acquisition decision of traders. Their analysis relies on the fact that traders are short-lived and might be forced to unwind their position before the collected information is reflected in the price. Consequently, traders have an incentive to collect information which other traders will also collect. In contrast, we focus on the persistence of bubbles in the presence of rational arbitrageurs without restricting their horizons.

fancy of the other competitors, *all of whom are looking at the problem from the same point of view. (italics added)*

Indeed, in our relativistic framework, each trader looks at the problem from the same point of view, however, *relative to the date*, when she became aware of the bubble's existence. Combining these two elements with our earlier assumption that traders become aware of the bubble in a sequential random order leads to our results.

We have assumed that the bubble only bursts when the selling pressure exceeds κ . Hence, we implicitly also assume that rational agents do not become aware of selling pressure by other rational agents until it crosses a certain threshold. These are simplifying assumptions. We relax them somewhat in Section 6 and 7 by allowing for intermediate price drops prior to the final crash. Nevertheless, we do not fully endogenize the price process. In our model, the central uncertainty is about the random variable t_0 . A fuller model would entail noisy prices which would preclude the selling pressure by rational arbitrageurs from being inferred with certainty from the current price level. However, we believe that our principal results would be qualitatively preserved in such a setting, though their precise expression would be substantially more complicated.

4 Preliminary Analysis

This section shows that we can restrict the analysis without loss of generality to trading strategies where each arbitrageur sells her shares only once. We also derive the *sell-out condition* according to which each arbitrageur sells her shares exactly at the moment when the temptation to 'ride the bubble' balances the fear of its imminent collapse (Lemma 8).¹⁵

We use the following notion of equilibrium.

Definition 1 *A trading equilibrium is defined as a Perfect Bayesian Nash Equilibrium in which whenever a trader's stock holding is less than the maximum, then the trader (correctly) believes that the stock holding of all traders who became aware of the bubble prior to her are also less than their respective maximum long positions.*

This definition entails a restriction on beliefs which is a natural one in our setting. Indeed, an immediate conjecture is that this property of beliefs is an *implication* of equilibrium, but we have not been able to prove this. Note that we are **not** restricting attention to trigger-strategies in which an agent who sells out at t continues to attack the bubble at all times thereafter.

¹⁵Much of the analysis in this section is used to establish that without loss of generality attention may be restricted to symmetric trigger-strategies. Readers who wish to omit these details may after absorbing Lemma 8 move directly to Section 5.

Definition 2 *The function $T(t_i) = \inf \{t | \sigma(t, t_i) > 0\}$ denotes the **first instant** at which arbitrageur \hat{t}_i sells any of his shares.*

Lemma 1 states that, in equilibrium, an arbitrageur is either fully invested in the market, $\sigma(t, t_i) = 0$, or at his maximum short-position, $\sigma(t, t_i) = 1$.

Lemma 1 (No partial purchases or sell-outs) $\sigma(t, t_i) \in \{0, 1\} \forall t, t_i$.

This Lemma “essentially” reduces the per period action space to 0 or 1. It simplifies the analysis in our paper since the aggregate selling pressure is simply given by the mass of traders who are ‘out’ of the market. It is a consequence of risk-neutrality and the fixed component of transactions costs. However, it is not essential for the main results of our paper to hold. Risk-averse arbitrageurs would gradually leave the market. This would make it necessary to keep track of each arbitrageurs position to calculate the aggregate selling pressure. All this would add complexity without a corresponding increase in insight.

Lemma 1 together with the definition of a trading equilibrium, immediately implies Corollary 1. It states that when arbitrageur \hat{t}_i sells out her shares, all arbitrageurs \hat{t}_j where $t_j \leq t_i$ also have already sold, or will at that moment, sell all their shares. We refer to this feature as the ‘cut-off’ property.

Corollary 1 (Cut-off Property) $\sigma(t, t_i) = 1 \Rightarrow \sigma(t, t_j) = 1 \forall t_j \leq t_i$ and $\sigma(t, t_i) = 0 \Rightarrow \sigma(t, t_j) = 0 \forall t_j \geq t_i$.

Arbitrageur $\hat{t}_i = t_0 + \widehat{\eta\kappa}$ reduces her holdings for the first time at $T(t_0 + \eta\kappa)$. Since, by Corollary 1 all arbitrageurs who became aware of the mispricing prior to $t_0 + \eta\kappa$ are also completely out of the market at $T(t_0 + \eta\kappa)$ the bubble bursts when trader $t_0 + \widehat{\eta\kappa}$ sells out her shares, provided that it did not already burst earlier for exogenous reasons.

Corollary 2 *The bubble bursts at $T^*(t_0) = \min \{T(t_0 + \eta\kappa), t_0 + \bar{\tau}\}$.*

We will show that in any equilibrium, each arbitrageur \hat{t}_i can rule out certain realizations of t_0 at the time when she first sells out his shares. For this purpose, we define $\underline{t_0^{\text{supp}}}(t_i)$.

Definition 3 *The function $\underline{t_0^{\text{supp}}}(t_i)$ denotes the lower bound of the support of trader \hat{t}_i 's posteriors beliefs about t_0 , at $T(t_i)$.*

Lemma 2 derives a lower bound for $\underline{t_0^{\text{supp}}}(t_i)$ by making use of the fact that each arbitrageur has an incentive to pre-empt any possible crash which occurs with strictly positive probability.

Lemma 2 (Preemption) *In equilibrium, arbitrageur \hat{t}_i believes at time $T(t_i)$ that at most a mass κ of arbitrageurs became aware of the bubble prior to him. That is, $\underline{t_0^{supp}}(t_i) \geq t_i - \eta\kappa$.*

The Preemption-Lemma allows us to derive further properties of the bursting time $T^*(\cdot)$. Lemmas 3 and 4 in the Appendix show that $T^*(\cdot)$ is strictly increasing and continuous. It follows that $T^{*-1}(\cdot) : [T^*(0), \infty) \mapsto [0, \infty)$, the inverse of $T^*(\cdot)$ is well defined. Lemma 5 in the Appendix establishes that $T(\cdot)$ is also continuous.

Lemma 6 (Zero Probability) *For all $t_i > 0$, arbitrageur \hat{t}_i believes that the bubble bursts with probability zero at the instant $T(t_i)$. That is, $\Pr[T^{*-1}(T(t_i)) | t_i, B^c(T(t_i))] = 0$ for all $t_i > 0$.*

The following proposition establishes that in any equilibrium arbitrageurs necessarily use trigger-strategies. That is, each arbitrageur \hat{t}_i sells out at $T(t_i)$ and never re-enters the market.

Proposition 1 (Trigger-strategy) *In equilibrium, arbitrageur \hat{t}_i maintains the maximum short position for all $t > T(t_i)$, until the bubble bursts.*

Proof. Suppose there exists a non-trigger strategy equilibrium where arbitrageur \hat{t}_i sells out at $T(t_i)$ and re-enters the market later. Given transaction costs $c > 0$ and the preceding lemma, in equilibrium arbitrageur \hat{t}_i must stay out of the market for a strictly positive time interval. He will stay out of the market at least until type $\widehat{t_i + \varepsilon}$ exits the market, for some $\varepsilon > 0$ (independent of \hat{t}_i).¹⁶ By the cut-off property arbitrageur \hat{t}_i cannot re-enter the market until after arbitrageur $\widehat{t_i + \varepsilon}$ first re-enters or $t_0 + \bar{\tau}$ occurs. The same reasoning applies for $\widehat{t_i + \varepsilon}$ with respect to $\widehat{t_i + 2\varepsilon}$. Proceeding this way we conclude that arbitrageur \hat{t}_i stays in the market until the bubble bursts at $t_0 + \bar{\tau}$ or arbitrageur $\widehat{t_i + \eta\kappa}$ re-enters the market. Of course, the latest possible date at which the bubble bursts from arbitrageur \hat{t}_i 's viewpoint is when $\widehat{t_i + \eta\kappa}$ exits the market. ■

We have proved that $T^{*-1}(\cdot)$ exists and is strictly increasing and continuous. We confine attention to equilibria for which the latter function is *absolutely* continuous such that $\Pi(t) = \Phi(T^{*-1}(t))$ is also absolutely continuous. Let $\pi(t)$ denote its associated density. Recall that $\Phi(t_0) = 1 - e^{-\lambda t_0}$ and $\Pi(t|t_i)$ is arbitrageur \hat{t}_i 's conditional cumulative distribution function of the bursting date at time t_i . Similarly, $\pi(t|t_i)$ denotes the associated conditional density.

¹⁶Exiting the market and paying the transaction cost c can only be justified if the bubble bursts with a certain probability. Since the bubble bursts with zero probability exactly when trader \hat{t}_i enters, he has to stay out of the market till at least a certain mass of (younger) traders ε also exit the market. Note ε is independent of t_i .

The trigger-strategy characterization greatly simplifies the analysis of equilibrium and indeed even the specification of payoffs. The payoff to selling out at time t reduces to¹⁷

$$\int_{t_i}^t e^{-rs} (1 - \beta(s - T^{*-1}(s))) p(s) \pi(s|t_i) ds + e^{-rt} p(t) (1 - \Pi(t|t_i)) - c.$$

Differentiating the payoff function with respect to t yields the sell-out condition stated in Lemma 8. Note that $h(t|\hat{t}_i) = \frac{\pi(t|\hat{t}_i)}{1 - \Pi(t|\hat{t}_i)}$ is the hazard rate that the bubble will burst at t .

Lemma 8 (Sell out condition) *If arbitrageur \hat{t}_i 's subjective hazard rate is smaller than the 'cost-benefit ratio', i.e.*

$$h(t|\hat{t}_i) < \frac{\frac{p'(t)}{p(t)} - r}{\beta(t - T^{*-1}(t))}$$

trader \hat{t}_i will choose to hold the maximum long position at t . Conversely, if $h(t|\hat{t}_i) > \frac{\frac{p'(t)}{p(t)} - r}{\beta(t - T^{-1}(t))}$ she will trade to the maximum short position.*¹⁸

To understand the sell out condition intuitively, consider the first-order benefits and costs of attack at t versus $t + \Delta$, respectively. The benefits are given by $\Delta h(t|\hat{t}_i) [p(t) E[\beta(t - t_0) | t_i, t]]$, the size of the bubble times the probability that the bubble will burst over the small interval Δ . In the case that the bubble does not burst, the costs of being out of the market for a short interval Δ are $(1 - \Delta h(t|\hat{t}_i)) \left(\frac{p(t+\Delta) - p(t)}{\Delta} - r \right) \Delta$. Note that $\frac{p(t+\Delta) - p(t)}{\Delta} - r > 0$ since the bubble appreciates faster than the riskfree rate. Dividing by $\Delta p(t)$ and letting $\Delta \rightarrow 0$ yields the attack condition. Note that in the limit as $\Delta \rightarrow 0$, $E[\beta(t - t_0) | t_i, t] = \beta(t - T^{*-1}(t))$. Lemma 8 also implies that trader \hat{t}_i either wholeheartedly attacks or holds the maximum long position.

¹⁷One could easily extend the analysis to capture the *reputational penalty*, institutional investors face for staying out of the market while the bubble grows at the expected rate g . *Relative performance evaluation* of fund managers is one of the main causes of this cost. Tony Dye's case provides a vivid illustration. He was for many years the successful chief investment officer of Phillips and Drew, London. Nicknamed "Dr. Doom" he refused to bow to the fashion of investing in internet stocks fearing an imminent slump in the markets. He lost his job in March 2000 - just days before his warnings that the tech bubble would burst came true. In the words of a market commentator "The irony is he [Tony Dye] may well be right, but at the wrong time."

In a setting with a 'reputational penalty' equal to a fraction k of the price level, the term $\int_t^{t_i + \tau} [\int_t^s e^{-ru} k p(u) du] \pi(s|t_i) ds$ needs to be added to the payoff.

¹⁸With a reputational penalty the sell out condition generalizes to $h(t|\hat{t}_i) < (>) \frac{\frac{p'(t)}{p(t)} - r - k}{\beta(t - T^{*-1}(t))}$.

Random Price Process For simplicity we have assumed that the price process grows at a deterministic rate of g . The analysis can be extended to a setting which incorporates stochastic price processes with a finite number of possible price paths. In particular, we have a setting in mind where the price grows in expectations at a rate of g . That is, $E_s [p_t] = p_s e^{g(t-s)}$ for all t, s with $t > s$, and we maintain the assumption that the bubble is always a fraction $\beta(t - t_0)$ of the price. In this setup the sell out condition $h(t|t_i) = \frac{E[p'(t)|p(t)] - r}{\beta(t - T^{*-1}(t))}$.

5 Persistence of Bubbles

The impossibility of bubbles emerging in standard asset pricing models is most transparent and compelling for finite horizon bubbles because they are ruled out by a straightforward backward induction argument. If the latest possible date at which the bubble bursts is period \bar{t} , all traders will start selling the asset in the penultimate trading round. Consequently, the price of the asset will already drop in the previous period, causing arbitrageurs to sell even earlier. Iterating this argument precludes the emergence of bubbles in the first place.

To facilitate comparison with, and to sharply contrast our results to those obtained in the literature, we also assume that the bubble ultimately bursts when the mispricing equals a fraction $\bar{\beta}$ of the price. Given that the function $\beta(t - t_0)$ is strictly increasing, there exists a corresponding $\bar{\tau}$, such that the bubble ultimately bursts at $t_0 + \bar{\tau}$ for exogenous reasons. Though this assumption might not be specially realistic, it arguably makes it harder for bubbles to sustain in equilibrium.

Let us illustrate an alternative backwards induction argument to intuitively understand the main results of this section. This procedure entails the “iterative removal of non-best-response symmetric trigger-strategies.” Even if no arbitrageur sells her shares, the bubble ultimately bursts at $t_0 + \bar{\tau}$. Since each arbitrageur becomes aware of the mispricing only after t_0 , she knows for sure that the bubble will never last beyond $t_i + \bar{\tau}$. But it might burst even before this time since from arbitrageur t_i ’s point of view, the ultimate bursting date $t_0 + \bar{\tau}$ is distributed between $t_i + \bar{\tau} - \eta$ and $t_i + \bar{\tau}$. If the bubble bursts at $t_0 + \bar{\tau}$, arbitrageurs \hat{t}_i ’s best response is to ride the bubble before exiting the market. Let $\tau(t_i)$ denote the number of periods during which arbitrageur \hat{t}_i optimally rides the bubble after becoming aware of the mispricing at t_i .

Definition 4 *The function $\tau(t_i) = T(t_i) - t_i$ denotes the length of time arbitrageur \hat{t}_i waits after becoming aware of the mispricing, before selling her shares.*

More specifically, let τ^1 be the best response if arbitrageur \hat{t}_i conjectures that the bubble bursts only at $t_0 + \bar{\tau}$. In any symmetric equilibrium, all other arbitrageurs will also sell out their shares τ^1 periods after they became aware of the mispricing. As in

the case of backwards induction, we can derive the new bursting date. The new best response of each arbitrageur is to ride the bubble for τ^2 periods. Similarly, we can iteratively obtain τ^3 , τ^4 and so on. In the standard model this yields $\lim_{n \rightarrow \infty} \tau^n = 0$, which precludes the emergence of bubbles. We show in Section 5.1 that this backwards induction procedure has no bite if $\frac{\lambda}{1-e^{-\lambda\eta\kappa}} \leq \frac{g-r}{\beta}$. In particular, $\lim_{n \rightarrow \infty} \tau^n = \tau^1$ and furthermore, the bubble only bursts at $t_0 + \bar{\tau}$. Conversely, for $\frac{\lambda}{1-e^{-\lambda\eta\kappa}} > \frac{g-r}{\beta}$, we show in Section 5.2 that $\lim_{n \rightarrow \infty} \tau^n \leq \tau^1$ but $\lim_{n \rightarrow \infty} \tau^n > 0$. Hence, this backwards procedure does bite, but not as much as in the classical case. Note that this induction argument is restricted to symmetric strategies. The formal analysis below employs a different reasoning which also captures non-symmetric strategies.

5.1 Exogenous Crashes

Recall that arbitrageurs become aware of the bubble sequentially in a random order and furthermore have a non-degenerate posterior distribution over t_0 . All arbitrageurs become ‘aware’ of the bubble during the interval $[t_0, t_0 + \eta]$, where we have interpreted η to be a measure of differences in opinion and other heterogeneities across players. From $t_0 + \eta\kappa$ onwards, more than κ arbitrageurs are aware of the bubble and have collectively the ability to burst it.

We show that if η is not too small, that is, if $\frac{\lambda}{1-e^{-\lambda\eta\kappa}} \leq \frac{g-r}{\beta}$, then in the unique trading equilibrium the bubble only bursts for exogenous reasons when it reaches its maximum size $\bar{\beta}$ relative to the price. In this case the endogenous selling pressure of the rational arbitrageurs has absolutely no influence on the time at which the bubble bursts. It is worth noting that our result holds despite the fact that it is possible within our model for traders to coordinate selling out on particular dates, say Friday, 13th of April 2001 by adopting (asymmetric) strategies which entail non-constant $\tau(t_i)$.

Proposition 2 *Suppose $\frac{\lambda}{1-e^{-\lambda\eta\kappa}} \leq \frac{g-r}{\beta}$. Then there exists a unique trading equilibrium.*

In this equilibrium all traders sell out $\tau^1 = \bar{\tau} - \frac{1}{\lambda} \ln \left(\frac{g-r}{g-r-\lambda\bar{\beta}} \right) < \bar{\tau}$ periods after they became aware of the bubble and stay out of the market thereafter. Nevertheless, for all t_0 , the bubble bursts for exogenous reasons precisely when it reaches its maximum possible size $\bar{\beta}$.

Proof.

Step 1: Derive symmetric equilibrium τ^1 .

Suppose the bubble bursts at $t_0 + \bar{\tau}$ for all t_0 . Then arbitrageur \hat{t}_i 's posterior distribution function over bursting dates $t = t_i + \tau$ is $\Pi(t_i + \tau | t_i) = \frac{e^{\lambda\eta} - e^{\lambda(\bar{\tau}-\tau)}}{e^{\lambda\eta} - 1}$ at t_i . The associated hazard rate is $h(t_i + \tau | t_i) = \frac{\pi(t_i + \tau | t_i)}{1 - \Pi(t_i + \tau | t_i)} = \frac{\lambda}{1 - e^{-\lambda(\bar{\tau}-\tau)}}$. Let τ^1 solve $h(t_i + \tau^1 | t_i) = \frac{g-r}{\beta}$. Then $\tau^1 = \bar{\tau} - \frac{1}{\lambda} \ln \left(\frac{g-r}{g-r-\lambda\bar{\beta}} \right)$. By the Sell-out Lemma arbitrageur \hat{t}_i is optimally out of the market for $\tau > \tau^1$ and conversely for $\tau < \tau^1$ he holds on to his shares. Under

our assumed condition $t_0 + \tau^1 + \eta\kappa > t_0 + \bar{\tau}$. Hence, the symmetric trigger τ^1 indeed defines an equilibrium, which results in an exogenous crash at $t_0 + \bar{\tau}$.




Figure 3: Exogenous crash

Step 2: The uniqueness argument can be found in Appendix A.2. ■

In equilibrium, each arbitrageur optimally rides the bubble very long. At the time of the crash $t_0 + \bar{\tau}$, fewer than κ arbitrageurs will have sold out at the pre-crash price and thus, the bubble bursts purely for exogenous reasons. This is in contrast to the standard backwards induction reasoning where the anticipation of a crash prepones the price drop.

We note here that the standard iterative dominance proof of global games cannot be applied in our setting, since our game does not satisfy strategic complementarity. This is both because the assumption that $\kappa < 1$ introduces a competitive element and the fact that traders infer information from the fact that the bubble still exists.¹⁹

5.2 Endogenous Crashes

The previous section demonstrates that arbitrageurs never burst the bubble if the dispersion of opinion among them η and the absorption capacity of behavioral traders

¹⁹This is probably best illustrated by means of an example, wherein we restrict the strategy space to trigger strategies. Consider a trader \hat{t}_i who starts attacking the bubble at $t^{13} = t_i + \tau_i$, provided that all other traders attack immediately when they became aware of the bubble. Given this strategy profile, trader \hat{t}_i can infer a lower bound for t_0 from the fact that the bubble still exists. Compare this with a situation where other traders do not start attacking immediately when they become aware of the bubble but only at, say t^{13} . In this case trader \hat{t}_i cannot derive a lower bound for t_0 from the existence of the bubble. Consequently, she has a greater incentive to attack the bubble at t^{13} . This is exactly the opposite of what strategic complementarity would prescribe.

κ is sufficiently large. In this section we examine the opposite case when this condition is not satisfied. We show that our proposed backward iteration procedure does bite in this case. When no other arbitrageur ever sells out, the bubble bursts at $t_0 + \bar{\tau}$, which induces arbitrageurs to sell out at $t_i + \tau^1$. In this scenario, the bubble will burst at $t_0 + \tau^1 + \eta\kappa$, which is now strictly earlier than $t_0 + \bar{\tau}$ given the assumed smaller parameter values for η and κ . Given that the bubble bursts latest at $t_0 + \tau^1 + \eta\kappa$, arbitrageurs seek to sell out even earlier; $\tau^2 < \tau^1$ periods after they became aware of the bubble. Proceeding in this way leads to a decreasing sequence $\tau^1, \tau^2, \tau^3, \dots$ which converges to some τ^* which in fact defines the unique symmetric trigger-strategy Perfect Bayesian Nash equilibrium. The iteration of this argument does not eliminate bubbles. In contrast to the Efficient Market Hypothesis, the bubble in our model survives for a substantial period of time. The reason is that the iterative procedure loses bite gradually. As it prepones the bursting date, the size of the bubble also diminishes, which in turn increases the incentive to ride the bubble.

Proposition 3 derives a symmetric strategy equilibrium, where each arbitrageur sells his shares τ^* periods after he became aware of the bubble. It also demonstrates that this trading equilibrium is unique.

Proposition 3 *Suppose $\frac{\lambda}{1-e^{-\lambda\eta\kappa}} > \frac{g-r}{\beta}$. Then there exists a unique trading equilibrium, in which the arbitrageurs \hat{t}_i with $t_i \geq \eta\kappa$ leave the market $\tau^* = \beta^{-1} \left(\frac{g-r}{\frac{\lambda}{1-e^{-\lambda\eta\kappa}}} \right) - \eta\kappa$ periods after they become aware of the bubble. All arbitrageurs \hat{t}_i such that $t_i < \eta\kappa$ sell out at $\eta\kappa + \tau^*$. Hence, the bubble bursts when it is a fraction*

$$\beta^* = \frac{1 - e^{-\lambda\eta\kappa}}{\lambda} (g - r)$$

of the pre-crash price.

Proof.

Step 1: Derive symmetric equilibrium τ^* .

Suppose that all arbitrageurs with $t_i \geq \eta\kappa$ sell out their shares at $t_i + \tau^*$ and arbitrageurs with $t_i < \eta\kappa$ at $\eta\kappa + \tau^*$. Then the bubble bursts at $t_0 + \eta\kappa + \tau^*$ and when it is a fraction $\beta(\eta\kappa + \tau^*) = \beta^*$ of the price. Since in a symmetric equilibrium the hazard rate of an arbitrageurs who sells out his shares is $\frac{\lambda}{1-e^{-\lambda\eta\kappa}}$, the FOC $\frac{\lambda}{1-e^{-\lambda\eta\kappa}} = \frac{g-r}{\beta^*}$ implies that

$$\beta^* = \frac{1-e^{-\lambda\eta\kappa}}{\lambda} (g - r). \text{ Hence, } \tau^* = \beta^{-1} \left(\frac{g-r}{\frac{\lambda}{1-e^{-\lambda\eta\kappa}}} \right) - \eta\kappa.$$




Figure 4: Endogenous crash

To derive uniqueness we show in **Step 2** that the bubble always bursts for endogenous reasons when $\frac{\lambda}{1-e^{-\lambda\eta\kappa}} > \frac{g-r}{\beta}$ and in **Step 3** that the minimum and maximum of $\tau(t_i)$ coincide for $t_i \geq \eta\kappa$. Both steps are derived in Appendix A.3.

Step 4: For $t_i < \eta\kappa$, $T(t_i) = T(\eta\kappa)$. Clearly, no \hat{t}_i should sell out prior to $T(\eta\kappa)$ and by the cut-off property will sell out at $T(\eta\kappa)$. ■

Notice that the maximum ‘relative’ bubble size β^* increases as the dispersion of opinions among arbitrageurs η increases. Taking our model literally, η describes the time span (‘awareness window’) over which traders become sequentially aware of the bubble. It is also essential for our argument to work that individual traders do not know when they became aware of the bubble relative to others: individual traders become aware of the bubble in a sequential, random order. The larger the ‘awareness window’ η , the more uncertain is each arbitrageur about when other traders became aware of the bubble. Alternative model formulations show that the dispersion of timing is crucial for the emergence of the bubble and not the difference in the estimate of the fundamental value. The comparative statics for the absorption capacity of the momentum traders κ are the same as for η . A larger κ requires more coordination among arbitrageurs and thus extends the bubble size. As one expects, β^* is also increasing in the ‘excess growth rate’ of the bubble ($g-r$). The faster the bubble appreciates, the more tempting it is to ride the bubble.

5.3 Discussion

This subsection discusses why the standard arguments which usually enable one to rule out bubbles do not apply in our setting.

Lack of Common Knowledge To gain a better understanding for why a bubble persists even though the life span of the bubble is finite, it is useful to take a closer look at arbitrageurs' knowledge of the bubble as time evolves. The standard backwards induction argument - which usually rules out finite bubbles - requires a starting point which is common knowledge. Proposition 4 below shows that it is never common knowledge that at least κ traders are aware of the bubble; this basic observation provides an alternative perspective on the difference between our model and the classical literature on bubbles.

Proposition 4 *It is never common knowledge among κ traders that at least κ traders are aware of the bubble.*

Proof. It is sufficient to look at the first κ traders. At $t_0 + \eta\kappa$, at least κ traders know of the bubble. That is, it is mutual knowledge among κ traders at $t_0 + \eta\kappa$. However traders, in particular arbitrageurs who only became aware at $t_0 + \eta\kappa$, are not sure whether the other arbitrageurs are aware of the bubble too. At $t_0 + 2\eta\kappa$, the first κ traders know that a bubble exists and that at least a fraction κ of the arbitrageurs knows of the bubble. However, they do not know whether a fraction κ knows that a fraction κ knows that the bubble exists, etc. This is only the case at $t_0 + 3\eta\kappa$. More generally, let n be a positive integer, then at $t_0 + n\eta\kappa$, the κ th trader knows that at least κ traders know that at least κ know that ... and so on at most n -times. It will never be common knowledge among κ traders that there are at least a fraction κ traders who know of the bubble. ■

No Zero-Sum Argument Recall that Tirole (1982) ruled out bubbles by employing a zero-sum argument in a setting with asymmetric information. Trading is a zero-sum game if the initial allocation is interim Pareto efficient. In particular, if all traders are risk-neutral any allocation is ex-ante Pareto efficient and hence, any arbitrageur's trading profit is somebody else's loss. In other words, traders are not willing to buy a bubble asset since some traders have already realized their gains and have left a negative-sum game for the other traders.

In contrast, in our model both rational arbitrageurs as well as behavioral momentum traders operate in the market. The presence of these momentum traders makes trading a positive sum game for the arbitrageurs. The over-optimistic momentum traders lose in our setting. More interestingly, the 'synchronization problem' of the arbitrageurs works in their favor. The longer the bubble persists, the higher is the aggregate trading profit of the arbitrageurs. Viewed differently, the synchronization problem enables them to coordinate their riding of the bubble since it weakens the preemption motive for each individual arbitrageur. Hence, from the arbitrageurs' point of view, the lack of synchronization is not a "problem" but rather a blessing.

6 Synchronizing Public Events

During the life-span of the bubble, information flows unrelated to t_0 might not only lead to random price changes but could also trigger a bursting of the bubble. So far we have abstracted from these additional news events. They are the focus of this section. An important distinction emerges between unanticipated public news - analyzed in Subsection 6.1 - and pre-scheduled news announcements, which are covered in Subsection 6.2.

6.1 Unanticipated news events

Unanticipated public events might serve as a synchronization device for arbitrageurs to sell, thereby triggering a crash. To analyze this question, we extend the earlier model to allow for the arrival of public signals at a Poisson arrival rate θ . We assume that $\theta < \frac{g-r}{\beta}$ to ensure that the possible occurrence of a public event does not by itself justify selling out. As in Subsection 5.2 (endogenous crashes) we focus on the parameter values which satisfy $\frac{g-r}{\beta} < \frac{\lambda}{1-e^{-\lambda\eta\kappa}} < \frac{g-r}{\beta(\eta\kappa)}$. The latter inequality guarantees that no arbitrageur finds it optimal to sell out even before she becomes aware of the mispricing.

Since the primary emphasis of this section is to understand the role of information in causing price changes *beyond* its informational content, we restrict our formal analysis to sunspot public events which serve as pure coordination devices. In other words, we will not consider signals which reveal additional information about the stock value or t_0 . Informative public signals would not entail qualitative changes to the analysis presented below.

Arbitrageurs who are aware of the bubble become more and more wary as time goes by. Therefore, they increasingly look out for signals which might cause the bubble to burst even though these signals might be totally unrelated to the fundamentals. We try to capture this idea by assuming that sunspots are only observed by traders who became aware of the bubble more than τ_e periods ago. Traders who are either unaware of the bubble or only recently became aware of it do not observe sunspots.²⁰ Alternatively, one can also envision a more general setting where traders who became aware of the bubble less than τ_e periods ago observe the sunspot but do not attribute much importance to it. Note that $\tau_e = 0$ captures the special case where all aware traders observe this public event.

²⁰Note that, if the public event serves to make even ‘unaware’ arbitrageurs aware of the bubble and this fact is common knowledge then the model become degenerate: we are back in a world of symmetric information in which the bubble cannot possibly survive after the public event.

Public events both allow arbitrageurs to synchronize sell-outs and also convey valuable information in the event that attacks are unsuccessful. The bubble might survive a synchronized sell-out if not sufficiently many arbitrageurs have observed the public event, that is, if t_0 is sufficiently high. Let t_e be the date of such a public event. By the cut-off property there exists $\tilde{\tau}_e \geq \tau_e$ such that all arbitrageurs \hat{t}_i who observed the sunspot at least $\tilde{\tau}_e$ periods after becoming aware of the mispricing sell out after observing the public event at t_e . If the bubble does *not* burst then all arbitrageurs learn that $t_0 > t_e - \tilde{\tau}_e - \eta\kappa$. It follows that all arbitrageurs \hat{t}_i with $t_i \leq t_e - \tilde{\tau}_e$ will re-enter the market until type $t_e - \tilde{\tau}_e$ first exists in equilibrium, or until the next public event occurs. Even traders who left the market prior to the arrival of the public event buy back their shares. Thus, unlike in the previous section, trigger-strategies do not arise in equilibrium.

A related simplification is however available: in any trading equilibrium arbitrageurs re-enter and leave the market at most once between two consecutive public events. Define an *interim-trigger-strategy* to be one for which an agent follows a trigger strategy *between* two successive public events. We argue that, in equilibrium, agents use interim-trigger-strategies. Let $H^n(t|t_i) := (t_e^{(1)}, \dots, t_e^{(n)})$ where $t_e^{(1)} < t_e^{(2)} < \dots < t_e^{(n)}$ denotes a history of past events at times $t_e^{(1)}, \dots, t_e^{(n)}$ strictly prior to t and observed by arbitrageur \hat{t}_i . Let $t_e^{(n-1)} \in H(t|t_i)$ be the time of the most recent public event observed by arbitrageur \hat{t}_i . Let t_a be either $t_e^{(n-1)}$ or t_i if arbitrageur \hat{t}_i has not observed any public event. Denote arbitrageur \hat{t}_i 's posterior distribution over t_0 at $t = t_a$ by $\Phi(\cdot|t_i, \tilde{B}^c(t_a), H(t_a|t_i))$. In other words, trader \hat{t}_i conditions upon her awareness date, the fact that the bubble did not burst prior to t_a and the history of observed past public events if $t_a = t_e^{(n-1)}$. The event that the bubble did not burst prior to t_a , $\tilde{B}^c(t_a)$, depends on the realization of t_0 and also on the sequence of public events. To simplify notation we summarize $\{t_i, \tilde{B}^c(t_a), H(t_a|t_i)\}$ by $\mathcal{I}_{t_i}(t_a)$.

Definition 5 (i) The function $T_e(t_i) = \inf \{t | \sigma(t, t_i) > 0, H^{n=0}(t|t_i) = \emptyset\}$ denotes the first instant at which arbitrageur \hat{t}_i sells out any of her shares if she did not observe a public event so far.

(ii) The function $T_e^*(t_0) = \min \{T_e(t_0 + \eta\kappa), t_0 + \bar{\tau}\}$ specifies the date at which the bubble bursts absent the occurrence of a public event, which is observed by trader $t_0 + \eta\kappa$.

The analysis for $T(\cdot)$ and $T^*(\cdot)$ of Section 4 applies to $T_e(\cdot)$ and $T_e^*(\cdot)$ in this section. In particular, $T_e^*(\cdot)$ is strictly increasing and continuous. The arguments are analogous and are not repeated here. Since each arbitrageur incurs positive transactions costs c when selling his shares, and the bubble bursts with zero probability at each instant, she has to leave the market at least for an interval. Arbitrageur \hat{t}_i leaves the market at least until arbitrageur $\widehat{t_i + \varepsilon}$ exits the market. Recall that the occurrence of

public events alone does not justify the loss in appreciation from exiting over a certain interval, since $\theta\bar{\beta} < g - r$. Hence, a generalized version of Proposition 1 applies *between* public events, leading to what we have termed interim-trigger-strategies. Furthermore, the distribution function that the bubble will burst in the absence of further public events is given by $\Pi(t|\mathcal{I}_i(t_a)) = \Phi(T_e^{*-1}(t)|\mathcal{I}_i(t_a))$.

Since in this section arbitrageurs re-enter the market in equilibrium, the existence of transactions costs are potentially incurred repeatedly. Keeping track of all transactions costs complicates the analysis in uninteresting ways. In what follows we exploit the strategic restrictions obtained above (i.e. interim-trigger-strategies), while setting transactions costs to zero.

In contrast to the previous sections, there is no hope of finding a unique equilibrium in this generalized setting, even though we were able to restrict the possible class of equilibrium strategies. For example, the equilibrium behavior of the previous section would be exactly replicated if all arbitrageurs were to simply ignore all public events, and there are potentially numerous intermediate levels of responsiveness to public signals. We focus on *responsive equilibria*.

Definition 6 *A ‘responsive equilibrium’ is a trading equilibrium, where each arbitrageur believes that all other traders will synchronize (sell out) at each public event.*

This raises the bar for our analysis since bubbles are harder to sustain in responsive equilibrium. This equilibrium may be viewed as the other extreme to the equilibrium in which all arbitrageurs ignore all public events.²¹

We establish that there exists a unique responsive equilibrium. It is symmetric in the sense that as long as an arbitrageur does not observe a public event, each arbitrageur rides the bubble for a fixed number of τ^{**} periods after becoming aware of it.

Proposition 5 *There exists a unique responsive equilibrium. In this equilibrium, each arbitrageur \hat{t}_i always sells out at the instances of public events $t_e \geq t_i + \tau_e$. Furthermore, she leaves the market for all $t \geq t_i + \tau^{**}$ except in the event that the last attack failed in which case she re-enters the market for the interval $t \in (t_e, t_e + \tau^{**} - \tau_e)$ unless a new public event occurs in the interim.*

²¹One problem is that there are potentially millions of sunspots which might serve as a potential coordination device. The decision of which sunspot to synchronize on entails another coordination problem. This leads us to the interesting theoretical discussion of what constitutes the “publicness” of a public event.

Absent the arrival of a public event the bubble will burst at $t_0 + \eta\kappa + \tau^{**}$. The arrival of public events might prepone the bursting date. After a failed sell out at the latest public event t_e , even traders who started selling out prior to t_e , that is, traders with $t_i < t_e - \tau^{**}$, buy back shares. Sell-outs only resume at $t_e + \tau^{**} - \tau_e$. Market sentiment “bounces back” after a failed attack. However, even in the event of failed attacks and re-entry by arbitrageurs into the market, the bubble bursts no later than $t_0 + \eta\kappa + \tau^{**}$.

Before proving Proposition 5 let us derive the sell-out condition for the case with unanticipated public events.

Lemma 9 (Generalized Sell-out Condition) *Suppose arbitrageur \hat{t}_i has not observed a failed sell out. Then she sells out at t satisfying*

$$h(t|\mathcal{I}_{t_i})\beta(t - T^{*-1}(t)) + \theta \int_{t_0 \leq t - \tau_e - \eta\kappa} (\alpha_{t_0,t}\beta(t - t_0)) d\Phi(t_0|\mathcal{I}_{t_i}) = g - r.$$

For a proof see Appendix A.1. Intuitively, the first term reflects the possibility that the bubble might burst in the next instant due to endogenous selling pressure. From trader \hat{t}_i 's viewpoint this occurs at the hazard rate $h(t|t_i) = \frac{\pi(t|\mathcal{I}_{t_i})}{1 - \Pi(t|\mathcal{I}_{t_i})}$. In this case, the relative size of the bubble is $\beta(t - T^{*-1}(t))$. Arbitrageur \hat{t}_i also has to take into account the possibility that an unexpected public announcement might occur at t with a density of θ . The sunspot is followed by a crash if $t_0 \leq t - \eta\kappa - \tau_e$. In this case the relative size of the bubble $\beta(t - t_0)$ depends on the realization of t_0 . Note that in the case of a public event, trader \hat{t}_i still has a chance to leave the market at the pre-crash price with probability $(1 - \alpha_{t_0,t})$. Hence, one has to multiply the advantage of attacking prior to the arrival of a public event by $\alpha_{t_0,t}$. Recall that $(1 - \alpha_{t_0,t})$ is the fraction of orders executed at the pre-crash price over all orders submitted immediately after a public event.

Proof of Proposition 5.

Step 1: Derive symmetric equilibrium τ^{**} .

1.1. Recall that $\tau^{**} = T_e(t_i) - t_i \forall t_i$ denotes the number of periods an arbitrageur waits to sell out if she does not observe any public event. Note that τ^{**} fully characterizes the symmetric responsive equilibrium outlined in Proposition 5. In any trading equilibrium all arbitrageur who sold out after the last public event, will re-enter the market after a failed attack and sell out again exactly when the ‘first’ trader, who did not participate at this common sell out, exits the market. This follows directly from the cut-off property. Hence, we can restrict our attention to the trading activity of traders who did not observe a public event yet. By describing their trading strategy we fully characterize any trading equilibrium.

1.2. There exists a unique symmetric trading equilibrium.

Note that each arbitrageur's \hat{t}_i posterior about t_0 at $T_e(t_i)$ exactly coincides with the posterior she had at $T(t_i)$ in a setting without public events. That is,

$\Phi(t_0|t_i, B^c(T_e(t_i)), H^n(T_e(t_i)|t_i)) = \Phi(t_0|t_i, B^c(T(t_i)), H^n(T(t_i)|t_i))$. To see this, observe that in any symmetric equilibrium, the support of t_0 at $T_e(t_i)$ is also $[t_i - \eta\kappa, t_i]$ (as in the previous section). Furthermore, any unobserved public event at $t_e < t_i + \tau_e$ would not have led to a bursting of the bubble, since $t_e < t_0 + \eta\kappa + \tau_e$ for all possible $t_0 \in [t_i - \eta\kappa, t_i]$. Hence, public events do not serve to distinguish between t_0 in the latter interval. Since the posteriors are the same, so is the hazard rate $h(t|\mathcal{I}_{t_i})$ at the time that trader \hat{t}_i sells out in either setting.

1.3. Clearly, if $\tau_e \geq \widehat{\tau^{**}}$, then public events have no impact on the bursting of the bubble, since trader $t_0 + \eta\kappa$, does not observe a signal prior to $t_0 + \eta\kappa + \tau^{**}$. Hence in this case $\tau^{**} = \tau^*$. Suppose now that $\tau_e < \tau^*$.

Define $\varphi(\tau) = h(t_i + \tau|\mathcal{I}_{t_i})\beta(\tau + \eta\kappa) + \theta \int_{t_0 \leq t_i + \tau - \tau_e - \eta\kappa} (\alpha_{t_0, s} \beta(t_i + \tau - t_0)) d\Phi(t_0|\mathcal{I}_{t_i}) - (g - r)$. For equilibrium τ^{**} it is necessary that $\varphi(\tau^{**}) = 0$. We argue that τ^{**} such that $\varphi(\tau^{**}) = 0$ is (i) unique, (ii) exists, and finally that such a τ^{**} indeed defines an equilibrium. By Step 1.2 $h(t_i + \tau^{**}|\mathcal{I}_{t_i})$ is constant across equilibrium τ^{**} . However, $\beta(\cdot)$ is strictly increasing, $\Phi(t_0|\mathcal{I}_{t_i})$ is the same, and the upper bound of the integral is increasing in τ . Thus, $\varphi(\tau^{**})$ is strictly increasing in equilibrium τ^{**} . Uniqueness follows directly.

For $\tau < \tau_e$ the sell-out condition reduces to $h(t_i + \tau|\mathcal{I}_{t_i})\beta(\tau + \eta\kappa) - g - r = 0$. There does not exist a $\tau < \tau_e$ which solves this equation. This follows directly from the fact that we focus on $\tau_e < \tau^*$ and that τ^* solves the preceding equation uniquely (as established in Section 5). Hence, $\varphi(\tau) < 0$ for $\tau \leq \tau_e$. Moreover, $\lim_{\tau \nearrow \bar{\tau}} \varphi(\tau) \rightarrow \infty$, since $h(t_i + \tau|\mathcal{I}_{t_i}) \rightarrow \infty$. Continuity of $\varphi(\tau)$ and the intermediate value theorem imply existence of τ^{**} .

1.4. Immediately after a public event each arbitrageur who observes the public event is assumed to sell out in a 'responsive equilibrium.' Hence, from each arbitrageur's point of view, a bubble bursts with strictly positive probability at each t_e in a responsive equilibrium. Given these beliefs, and the fact that an instantaneous attack is costless for $c = 0$, it is indeed (strictly) optimal for an arbitrageur who observes the public event to sell out at this time.

1.5. To fully specify all relevant strategies, it only remains to consider continuation strategies after a failed sell-out attempt. After a failed attack, arbitrageurs learn that fewer than κ traders have observed the public event. That is, $t_0 > t_e - \tau_e - \eta\kappa$. Since all other arbitrageurs, who did not observe the public event only sell out at $t_i + \tau^{**}$, all arbitrageurs who observed the public event at t_e can rule out the possibility that the bubble bursts prior to $t_e + \tau_e^{**} - \tau_e$ provided that no new public event occurs. Note that the bubble will not burst for exogenous reasons prior to $t_e + \tau^{**} - \tau_e$, since the endogenous bursting time $t_0 + \eta\kappa + \tau^{**}$ occurs strictly before $t_0 + \eta\kappa + \tau^* < t_0 + \bar{\tau}$.

The bubble might only burst prior to $t_e + \tau^{**} - \tau_e$ if a new public event occurs. After $t = t_e + \tau_e^{**} - \tau_e$, the analysis coincides with a setting without a public event at t_e . At this point all arbitrageurs who had participated in the failed sell out and subsequently re-entered the market, exit again.

The next two steps of the proof (see Appendix A.5) establish that equilibria are necessarily symmetric. They are analogous to Steps 2 and 3 of the proof of Proposition 3.

■

This section develops the idea that it might be more important to focus on news events that other traders consider as possible price movers than to focus on fundamentals per se. As noted earlier, these themes appear prominently in Keynes (1936).

Our analysis also sheds some light on the fact that there are fads and fashions in information. For example, trade figures drove the market during the 1980's. In contrast, in the late 1990's Alan Greenspan's statements moved stock prices, while trade figures were ignored.

The analysis also demonstrates that an unanticipated increase in uncertainty might itself lead to significant price swings. An unanticipated increase in θ , the Poisson density with which public signals arrive reduces τ^{**} and consequently may trigger arbitrageurs to attack the bubble even absent any relevant news. The outcome of the 2000 US-presidential elections may be viewed in this light. Some people have argued that the *uncertainty* surrounding whether Gore or Bush won the election and when the issue would be resolved served as a "smokescreen" for the price correction in high-tech stocks that was anyway necessary.

6.2 Pre-scheduled News Events

While some news events are unexpected, many important public announcements are pre-scheduled. In this subsection, we consider news whose announcement date is known in advance but the content of the news itself is unknown. Examples of this type of news include quarterly and monthly announcements of macroeconomic data like unemployment figures, trade balance data, inflation numbers as well as company specific news like regular dividend and earning announcements. This subsection highlights that such news disclosures cannot serve as a synchronization device and hence, they do not lead to price changes beyond the fundamental content of the news. If they did, a crash would occur at a specific date with strictly positive probability. Each arbitrageur would then have an incentive to pre-empt this ostensible crash, which rules it out as an equilibrium outcome. Note that conditioning on anticipated public events is similar to conditioning on a specific date, say "Friday 13th of April 2001." We have argued in the preceding Section 5 that an anticipated crash on a specific date can not occur. Therefore, it is not surprising that the introduction of pre-scheduled public events does not alter the analysis. Note that the result also holds for announcements which

are expected to occur with positive probability. Nevertheless, one should be wary of interpreting this result too literally. A news announcement can also trigger some unexpected information revelation. These unanticipated events, which occur only with positive density, lead us back to the analysis of Subsection 6.1.

7 Price Cascades and Market Rebounds

Arguably, the most visible public events on Wall Street are large past price movements or breaks through psychological resistance lines. In this section, we allow random temporary price drops to occur, which are possibly due to mood changes by the behavioral momentum traders. This enables us to illustrate how a large price decline might either lead to a full blown crash or to a rebound. In the latter event the bubble is strengthened in the sense that all arbitrageurs are ‘in the market’ for some interval, including those who had previously exited prior to the price shock.

In our simple and highly stylized model, we assume that exogenous price drops occur with a Poisson density θ_p at the end of a random trading round t . Let $H_p^n := (t_p^{(1)}, \dots, t_p^{(n)})$ where $t_p^{(1)} < t_p^{(2)} < \dots < t_p^{(n)}$ denotes a history of past temporary price drops at times $t_p^{(1)}, \dots, t_p^{(n)}$. The price drop shakes momentum traders’ mood temporarily and they are only willing to take on shares if the price is less than or equal to $(1 - \gamma_p) p_t$. If the bubble does not burst in the “subsequent” trading round, momentum traders regain their confidence and are willing to sell and buy at a price of p_t until their absorption capacity κ is reached or another price drop occurs. Arbitrageurs who exit the market immediately after a price drop receive $(1 - \gamma_p) p_t$ per share. Should the synchronized attack after a price drop fail, arbitrageurs can only buy back their shares at a price of p_t . In other words, leaving the market even only for an instant is very costly despite the absence of transactions costs, which we set equal to zero. Hence, only traders who are sufficiently certain that the bubble will burst after the price drop will leave the stock market. More specifically, Proposition 6 shows that only traders who became aware of the mispricing more than $\tau_p(H_p^n)$ periods earlier, will choose to leave the market and attack the bubble after a price drop. Notice that, although $\tau_p(H_p^n)$ is derived endogenously for the subgame after a price drop, it serves the same role as τ_e in Subsection 6.1. Consequently, Proposition 6 has the same structure as Proposition 5. One difference between both sections is that the history of past price drops H_p^n is known to all arbitrageurs, which simplifies the analysis. Corresponding to the previous subsection, τ^{***} reports the length each arbitrageur rides the bubble after t_i if there was no price drop so far.

Proposition 6 *There exists an equilibrium $(\tau_p(H_p^n), \tau^{***})$ in which arbitrageur \hat{t}_i exits the market after a price drop at $t_p^{(n)}$ if $t_p^{(n)} \geq t_i + \tau_p(H_p^n)$. Furthermore, she is*

out of the market at all $t \geq t_i + \tau^{***}$ except in the event that the last attack failed in which case she re-enters the market for the interval $t \in \left(t_p^{(n)}, t_p^{(n)} + \tau^{***} - \tau_p(H_p^n) \right)$.

Proposition 6 shows that a price drop which is not followed by a crash leads to a rebound and temporarily strengthens the bubble. In this case all arbitrageurs can rule out that the bubble will burst within $\left(t_p^{(n)}, t_p^{(n)} + \tau^{***} - \tau_p(H_p^n) \right)$ for endogenous reasons. Within this time interval, the price grows at a rate of g with certainty modulo another exogenous price drop. Consequently, all arbitrageurs re-enter the market after a failed attack and buy back shares at a price of p_t , even if they have sold them an instant earlier for $(1 - \gamma^p) p_t$.

The structure of this equilibrium is similar to the unique responsive equilibrium of the preceding section. However, the critical value $\tau_p(H_p^n)$ is endogenous and depends on the whole history of past price drops H_p^n . For example, a failed attack at $t_p^{(n-1)}$ makes arbitrageurs more cautious about the prospects of mounting a successful attack after a price drop at $t_p^{(n)} > t_p^{(n-1)}$. If $t_p^{(n)}$ is close to $t_p^{(n-1)}$, arbitrageurs will not find it optimal to exit again at $t_p^{(n)}$ if a synchronized sell out at $t_p^{(n-1)}$ did not succeed. The failed attack at $t_p^{(n)}$ increases $\tau_p(H_p^n)$. The equilibrium described in Proposition 6 is also maximally responsive in that price drops are responded to ‘maximally’ in the chronological order in which they appear. However, the bubble is more likely to burst at $t_p^{(n)}$ in an equilibrium in which the earlier event at $t_p^{(n-1)}$ was not responded to than in the equilibrium where arbitrageurs sold out at $t_p^{(n-1)}$ and the attack failed. We do not have a uniformly most responsive equilibrium in this section but rather one which is responsive in chronological order to the maximum extent possible. In short, the above equilibrium is not necessarily the one, in which the bubble bursts earliest for any possible sequence of price drops. Note that the same issue also arises in a setting with unanticipated public events and strictly positive transactions costs $c > 0$. We abstracted from these effects in Section 6 by assuming $c = 0$.

Note that it is important that price drops can also occur prior to time t_0 . Otherwise, immediately after a price drop it is commonly known that a bubble exists and a backwards induction argument starting from $t_p + \bar{\tau}$ would lead to an immediate collapse of the bubble.

8 Conclusion

This paper argues that bubbles can persist even though all rational arbitrageurs know that the price is too high and they jointly have the ability to correct the mispricing. Though the bubble will ultimately burst, in the intermediate term, there can be a large and long-lasting departure from fundamental values. A central (and we believe,

realistic) assumption of our model is that there is a dispersion of opinion among rational arbitrageurs concerning the timing of the bubble. This assumption serves both as a general metaphor for differences of opinion, information and belief among traders, and, more literally, as a reduced-form modeling of the temporal expression of heterogeneities amongst traders. While it is well understood that appropriate departures from common knowledge will permit bubbles to persist, we believe that our particular formulation is both natural and parsimonious. The model provides a setting in which ‘overreaction’ and self-feeding price drops, leading to full-fledged crashes, will naturally arise. It also provides a framework which allows one to rationalize phenomena such as ‘resistance lines’ and fads in information gathering.

Finally we note here that many of the assumptions of our simple model may be viewed as being conducive to arbitrage. In particular, we assume that all professionals are in agreement that assets are overvalued, while arguably there are substantial differences in opinion even amongst professionals regarding the possibility that current valuations indeed reflect a new era of higher productivity growth, lower wages and inflation etc. Presumably incorporating these realistic complications would reinforce our conclusions.

A Appendix

A.1 Details of Section 4

Lemma 1 (No partial purchases or sell-outs) $\sigma(t, t_i) \in \{0, 1\} \forall t, t_i$.

Proof. Consider an equilibrium strategy $\sigma(\cdot, t_i)$ which involves a change in position at t^k from the preceding position adopted at $t^{k+1} < t^k$. (The initial position is represented by K .) Suppose that $x^k(t^k) \notin \{0, 1\}$ and $\sigma(\cdot, t_i)$ is optimal. The expected payoff from t^{k+1} onwards is given by $V^{k+1}(\cdot)$ plus the value of bond holdings. Notice that $V^{k+1}(\cdot)$, excluding the transaction costs $ce^{rt^{k+1}}$, is linear in x^k and furthermore must be strictly positive. Hence, for $x^k - x^{k+1} > (<)0$ the payoff is strictly dominated by $x^k = 1$ ($x^k = 0$). This contradicts the initial presumption that $\sigma(\cdot, t_i)$ is optimal. ■

Lemma 2 (Preemption) *In equilibrium, arbitrageur \hat{t}_i believes at time $T(t_i)$ that at most a mass κ of arbitrageurs became aware of the bubble prior to him. That is, $\underline{t_0^{supp}}(t_i) \geq t_i - \eta\kappa$.*

Proof. Suppose $\underline{t_0^{supp}}(t_i) < t_i - \eta\kappa$. Then with a strictly positive probability the aggregate selling pressure at $t = T(t_i)$ is $s_{t_0, t} > \kappa$ and with strictly positive probability traders who first sell out at $T(t_i)$ will only receive the post-crash price for their sales. This follows from Lemma 1 and Corollary 1, which states that all traders $t_j \leq t_i$ will also be fully out of the market at $T(t_i)$. This leads to an immediate contradiction

unless $T(t_i) = 0$. Let $\underline{t}_0^{\text{supp}}(t_i) = t_i - \eta\kappa - a$. For arbitrageur \hat{t}_i to believe that $a > 0$, it is necessarily the case that for all $t_j \in [t_i - a, t_i)$, arbitrageurs \hat{t}_j do not sell out prior to $T(t_i)$. However, this is not optimal for them since all of them have an incentive to preempt (attack slightly earlier) the possible crash at $T(t_i)$. ■

Lemma 3 (Bursting Time) *The function $T^*(\cdot)$ is strictly increasing.*

Proof. Recall $T^*(t_0) = \min\{T(t_0 + \eta\kappa), t_0 + \bar{\tau}\}$. Since $T(\cdot)$ is weakly increasing by the cut-off property, so is $T^*(\cdot)$. Now suppose that there exists $\bar{t}_0 > \underline{t}_0$ such that $T^*(\underline{t}_0) = T^*(\bar{t}_0)$. Clearly, $\underline{t}_0 + \bar{\tau} \geq T^*(\underline{t}_0)$. It follows that for all $t_0 \in (\underline{t}_0, \bar{t}_0]$, $T(t_0 + \eta\kappa) = T^*(t_0) = T^*(\underline{t}_0)$. Let $t_i = t_0 + \eta\kappa$. By Lemma 2 $\underline{t}_0^{\text{supp}}(t_i) \geq t_0$. However, $\underline{t}_0^{\text{supp}}(t_i) \leq \underline{t}_0 < t_0 \equiv t_i - \eta\kappa$. ■

Lemma 4 (Continuity of T^*) *The function $T^* : [0, \infty) \rightarrow [0, \infty)$ is continuous.*

Proof. (i) Since T^* is increasing, we only need to rule out upward jumps. Suppose, that there is an upward jump at t_0 . Then for small enough $\varepsilon > 0$ and $s \in (t_0 - \varepsilon, t_0)$ $T^*(s) = T(s + \eta\kappa)$. Let $\bar{T}^* = \lim_{s \searrow t_0} T^*(s)$ and $\underline{T}^* = \lim_{s \nearrow t_0} T^*(s)$, and suppose $\varepsilon < \bar{T}^* - \underline{T}^*$. Recall that transaction costs ce^{rt} is incurred for any change of position at t . This precludes an arbitrageur from changing position at t and t' if the probability of a crash between t and t' is small relative to c . Hence, for small enough $\varepsilon > 0$, type $t_i = t_0 + \eta\kappa - \varepsilon$ is strictly better off selling out at $\bar{T}^* - \varepsilon > \underline{T}^*$ than at $T(t_i) < \underline{T}^*$, a contradiction. An almost identical argument establishes the continuity of $T(\cdot)$. ■

Lemma 5 (Continuity of T) *The function $T : [0, \infty) \rightarrow [0, \infty)$ is continuous.*

Proof. This argument uses Corollary 3 and is almost identical in structure to the proof of Lemma 4. ■

Lemma 6 (Zero Probability) *For all $t_i > 0$, arbitrageur \hat{t}_i believes that the bubble bursts with probability zero at the instant $T(t_i)$. That is, $\Pr[T^{*-1}(T(t_i)) | t_i, B^c(T(t_i))] = 0$ for all $t_i > 0$.*

Proof. $\Phi[t_0 | t_i, B^c(T(t_i))] = \frac{\Phi(t_0) - \Phi(\underline{t}_0^{\text{supp}}(t_i))}{\Phi(t_i) - \Phi(\underline{t}_0^{\text{supp}}(t_i))}$. It is obvious that this conditional c.d.f. is well behaved. ■

A.2 Uniqueness proof of Proposition 2

Step 2: Uniqueness.

Let us suppose that there is another equilibrium. Consider $\underline{t}_j = \arg \min_{t_i} \{\tau(t_i)\}$.²² Since $\tau(t_i) \leq \tau^1$ for all t_i , it must be the case that $\tau(\underline{t}_j) < \tau^1$. Consider $\underline{t}_0^{\text{suppp}}(\underline{t}_j)$

(i) By the Pre-emption Lemma $\underline{t}_0^{\text{suppp}}(\underline{t}_j) \geq \underline{t}_j - \eta\kappa$.

(ii) If $\underline{t}_0^{\text{suppp}}(\underline{t}_j) > \underline{t}_j - \eta\kappa$ then arbitrageur \hat{t}_j does not delay selling out at $\underline{t}_j + \tau(\underline{t}_j)$ only out of fear of an *exogenous* bursting of the bubble. By the argument of the first part, $\tau(\underline{t}_j) = \tau^1$. This contradicts the initial assumption that $\tau(\underline{t}_j) < \tau^1$.

(iii) Finally, suppose $\underline{t}_0^{\text{suppp}}(\underline{t}_j) = \underline{t}_j - \eta\kappa$. In this case the hazard rate that the bubble bursts at the time when \hat{t}_j sells out is at most $\frac{\lambda}{1-e^{-\lambda\eta\kappa}}$ since $\underline{t}_j = \arg \min \{\tau(t_i)\}$. Since this is in turn less than $\frac{g-r}{\beta}$, the sell-out condition is violated.

A.3 Uniqueness proof of Proposition 3

Step 2: Bubble always bursts for endogenous reasons when $\frac{\lambda}{1-e^{-\lambda\eta\kappa}} > \frac{g-r}{\beta}$.

Let τ^1 solve $h(t_i + \tau^1 | t_i) = \frac{\lambda}{1-e^{-\lambda(\bar{\tau}-\tau^1)}} = \frac{g-r}{\beta}$. Each trader would exit the market at $t_i + \tau^1$ if they believed that the bubble would burst for exogenous reasons when it reached its maximum possible size $\bar{\beta}$. Under our assumptions $t_0 + \tau^1 + \eta\kappa < t_0 + \bar{\tau}$. Hence, the bubble does not always burst for endogenous reasons only if $\tau(t_i) > \tau^1$ for at least some t_i . Consider arbitrageur \hat{t}_j , where $t_j = \arg \max_{t_i} \{\tau(t_i)\}$. Such an equilibrium can be ruled out by looking at $\underline{t}_0^{\text{suppp}}(t_j)$.

(i) By the Pre-emption Lemma $\underline{t}_0^{\text{suppp}}(t_j) \geq t_j - \eta\kappa$.

(ii) If $\underline{t}_0^{\text{suppp}}(t_j) > t_j - \eta\kappa$ then arbitrageur \hat{t}_j 's hazard rate at $T(t_j) > t_j + \tau^1$ that the bubble will burst for exogenous reasons strictly exceeds $\frac{g-r}{\beta}$. Since this is also the case at $T(t_j) - \varepsilon$ for sufficiently small $\varepsilon > 0$, she has an incentive to sell out strictly prior to $T(t_j)$.

(iii) Finally, suppose $\underline{t}_0^{\text{suppp}}(t_j) = t_j - \eta\kappa$. Since $\tau(t_j) \geq \tau(t_i)$ for all t_i , the hazard rate that the bubble will burst at $T(t_j)$ is at least $\frac{\lambda}{1-e^{-\lambda\eta\kappa}}$. Since the bubble bursts at $t_0 + \eta\kappa + \tau(t_j) > t_0 + \eta\kappa + \tau^1$, the size of the bubble $\beta(\eta\kappa + \tau(t_j))$ exceeds $\beta(\eta\kappa + \tau^1)$. It follows that the first order conditions for arbitrageur \hat{t}_j are violated at $T(\hat{t}_j)$.

Step 3: Minimum and maximum of $\tau(t_i)$ coincide for $t_i \geq \eta\kappa$.

By Step 2 $T(t_0 + \eta\kappa) < t_0 + \bar{\tau} \forall t_0 \geq 0$. That is, the bubble only bursts for *endogenous* reasons. By the Pre-emption Lemma, $\underline{t}_0^{\text{suppp}}(t_i) \geq t_i - \eta\kappa$. Furthermore, $\underline{t}_0^{\text{suppp}}(t_i) > t_i - \eta\kappa$ can be ruled out since arbitrageur \hat{t}_i would be strictly better off selling out at some $\varepsilon > 0$ time after her ostensible equilibrium sell-out date $T(t_i)$. Recall that $T(t)$ is continuous (Lemma 5). Hence, given $\underline{t}_0^{\text{suppp}}(t_i) = t_i - \eta\kappa$, arbitrageur \hat{t}_i 's conditional

²²Whenever $\arg \min$ and $\arg \max$ are not defined, the corresponding arguments can be extended in terms of infimums and supremums respectively.

density of t_0 at $T(t_i)$ is $\phi(t_i - \eta\kappa | t_i, B^c(T(t_i))) = \frac{\lambda e^{\lambda\eta\kappa}}{e^{\lambda\eta\kappa} - 1}$ which is independent of t_i . Let $\underline{t}_i \in \arg \min \tau(t_i)$ and $\bar{t}_i \in \arg \max \tau(t_i)$ and suppose that $\max \tau(t_i) > \min \tau(t_i)$. By the continuity of $T(\cdot)$ shown in Lemma 5 and the definitions of \underline{t}_i and \bar{t}_i , it follows that $\Pi(T(\underline{t}_i) + \Delta | \underline{t}_i, B^c(T(\underline{t}_i))) < \Pi(T(\bar{t}_i) + \Delta | \bar{t}_i, B^c(T(\bar{t}_i)))$ for all $\Delta > 0$, and conversely for $\Delta < 0$. Consequently, $h(T(\underline{t}_i) | \underline{t}_i, B^c(T(\underline{t}_i))) < h(T(\bar{t}_i) | \bar{t}_i, B^c(T(\bar{t}_i)))$. However, $\beta(\eta\kappa + \tau(\underline{t}_i)) \leq \beta(\eta\kappa + \tau(\bar{t}_i))$. Thus, the sell out condition cannot be satisfied for both arbitrageurs \hat{t}_i and \bar{t}_i , a contradiction.

A.4 Proof of Lemma 9 in Section 6.1

The payoff of a strategy, where arbitrageur \hat{t}_i is fully invested until either t or until she observes her first public sunspot at $t_e \geq t_i + \tau_e$. That is, she holds her shares until $\min\{t, t_e\}$ and follows the optimal continuation strategy thereafter. It proves useful, to specify the payoffs of arbitrageur \hat{t}_i at the time t_i . Let $\Gamma(t)$ denote the cumulative distribution function that the bubble bursts due to public event prior to t . Let $\gamma(t)$ be its associated density.

The payoff of this strategy at the time $t_i + \tau_e$ can be written as the sum of four components:

$$\begin{aligned} & e^{-rt} p(t) (1 - \Pi(t|t_i)) (1 - \Gamma(t|t_i)) + \\ & + \int_{t_i}^t e^{-rs} (1 - \beta(s - T_e^{*-1}(s))) p(s) (1 - \Gamma(s)) \pi(s|t_i) ds + \\ & + \int_{t_i}^t e^{-rs} \gamma(s) (1 - \Pi(s|t_i)) p(s) \left[\int_{t_0 \leq s - \tau_e - \eta\kappa} (1 - \alpha_{t_0, s} \beta(s - t_0)) \frac{d\Phi(t_0|t_i)}{\Phi(s - \tau_e - \eta\kappa|t_i)} \right] ds + \\ & + V(t_i) \end{aligned}$$

First, if the bubble does not burst prior to t and the bubble does not burst due to a public event up to t , then the arbitrageur sells his shares at the price $p(t) = e^{gt}$. Note that this only occurs with probability $(1 - \Pi(t|t_i)) (1 - \Gamma(t))$.

The second payoff component deals with the case where bubble does not burst after a public event but it burst for endogenous reasons at s prior to t . In these cases, the arbitrageur only receives $(1 - \beta(s - T_e^{*-1}(s))) p(s)$ and hence the expected payoff is $\int_{t_i}^t e^{-rs} (1 - \beta(s - T_e^{*-1}(s))) p(s) (1 - \Gamma(s)) \pi(s|t_i) ds$. Note that $T_e^*(\cdot)$ is invertible, since all arbitrageurs employ ‘interim-trigger strategies’.

The third term, considers the case where the bubble bursts after a public event. In this case the arbitrageur receives $(1 - \alpha_{t_0, s}) p(s) + \alpha_{t_0, s} (1 - \beta(s - t_0)) p(s)$. Note that the selling pressure might strictly surpass κ at the instant after s and hence arbitrageur

only receives a convex combination between the pre-crash and post-crash price. For $s < t_i + \tau_e$, arbitrageur \hat{t}_i does not observe the public event. Since he cannot react to it, she always only receives the post-crash price $(1 - \beta(s - t_0))p(s)$. That is $\alpha_{t_0,s} = 1 \forall s < t_i + \tau_e$. Since the bubble only bursts for $s \geq t_0 + \tau_e + \eta\kappa$, we have to divide $d\Phi(t_0|\mathcal{I}_i)$ by $\Phi(s - \tau_e - \eta\kappa|\mathcal{I}_i)$.

Finally, the fourth component reflects the value of the *option* to re-enter the market and to ride the bubble after the first observed failed sell-out at s . $V(t_i) = \int_{t_i + \tau_e}^{\infty} \theta e^{-\theta s} v(s|t_i) [1 - \Pi(s|t_i)] [1 - \Phi(s - \tau_e - \eta\kappa|t_i)] ds + V^{s+}(t_i)$, where $v(s|t_i)$ is the expected maximum payoff of starting from a zero position (i.e. $x(s) = 0$) provided that the first public event occurred at s and the bubble did not burst prior to s . Note that this ‘option value’ only arises for the cases where the bubble still exists when a public event occurs and the bubble survives the sell out after the public event, that is if $t_0 > s - \tau_e - \eta\kappa$. $V^{s+}(t_i)$ captures the option value for the second observed failed attack onwards.

Differentiating with respect t yields a generalized sell-out condition.

$$\begin{aligned} & -re^{-rt}p(t)(1 - \Pi(t|t_i))(1 - \Gamma(t|t_i)) + e^{-rt}p'(t)(1 - \Pi(t|t_i))(1 - \Gamma(t|t_i)) \\ & - e^{-rt}p(t)\pi(t|t_i)(1 - \Gamma(t|t_i)) - e^{-rt}p(t)(1 - \Pi(t|t_i))\gamma(t|t_i) \\ & + e^{-rt}(1 - \beta(t - T_e^{*-1}(t)))p(t)(1 - \Gamma(t))\pi(t|t_i) \\ & + e^{-rt}\gamma(t)(1 - \Pi(t|t_i))p(t) \left[\int_{t_0 \leq t - \tau_e - \eta\kappa} (1 - \alpha_{t_0,t}\beta(t - t_0)) \frac{d\Phi(t_0|t_i)}{1 - \Phi(t - \tau_e - \eta\kappa|t_i)} \right]. \end{aligned}$$

Dividing it by $e^{-rt}p(t)(1 - \Pi(t|t_i))(1 - \Gamma(t|t_i))$ simplifies the FOC to

$$\begin{aligned} & -r + \frac{p'(t)}{p(t)} - \frac{\pi(t|t_i)}{(1 - \Pi(t|t_i))} - \frac{\gamma(t|t_i)}{(1 - \Gamma(t|t_i))} + (1 - \beta(t - T_e^{*-1}(t))) \frac{\pi(t|t_i)}{(1 - \Pi(t|t_i))} \\ & + \frac{\gamma(t)}{(1 - \Gamma(t|t_i))} \left\{ 1 - \int_{t_0 \leq t - \tau_e - \eta\kappa} \alpha_{t_0,t}\beta(t - t_0) \frac{d\Phi(t_0|t_i)}{1 - \Phi(t - \tau_e - \eta\kappa|t_i)} \right\} = 0 \end{aligned}$$

Lemma 8 follows by replacing $\frac{p'(t)}{p(t)}$ with g , $\frac{\pi(t|t_i)}{(1 - \Pi(t|t_i))}$ with $h(t|t_i)$ and $\frac{\gamma(t|t_i)}{(1 - \Gamma(t|t_i))} \frac{1}{1 - \Phi(t - \tau_e - \eta\kappa|t_i)}$ with θ .

A.5 Details of Proof of Proposition 5

As noted in the text, the following steps are analogous to Step 2 and 3 of the proof of Proposition 3.

Step 2: The bubble always bursts for endogenous reasons when $\frac{\lambda}{1 - e^{-\lambda\eta\kappa}} > \frac{g-r}{\beta}$. Step 2 of proof of Proposition 3 shows that the standard sell-out condition is violated for each trader who sells out only at $t_i + \tau^1$ even in the absence of public events. Since

the additional term $\theta \int_{t_0 \leq t_i + \tau - \tau_e - \eta\kappa} (\alpha_{t_0, s} \beta(t_i + \tau - t_0)) d\Phi(t_0 | \mathcal{I}_{t_i})$ in the general sell-out condition is always positive and increasing in τ , it is also violated in the generalized setting with public events.

Step 3: Minimum and maximum of $\tau(t_i)$ coincide for $t_i \geq \eta\kappa$.

By Step 2, $T_e(t_0 + \eta\kappa) < t_0 + \bar{\tau} \forall t_0 \geq 0$. That is, the bubble only bursts for *endogenous* reasons. By the Pre-emption Lemma, $\underline{t_0^{\text{supp}}}(t_i) \geq t_i - \eta\kappa$. Furthermore, $\underline{t_0^{\text{supp}}}(t_i) > t_i - \eta\kappa$ can be ruled out since arbitrageur \hat{t}_i would be strictly better off selling out at some $\varepsilon > 0$ time after her ostensible equilibrium sell-out date $T_e(t_i)$. Recall that $T_e(t)$ is continuous. Hence, arbitrageur \hat{t}_i 's conditional density of t_0 at $T_e(t_i)$ is $\phi(t_i - \eta\kappa | t_i, B^c(T_e(t_i))) = \frac{\lambda e^{\lambda\eta\kappa}}{e^{\lambda\eta\kappa} - 1}$ which is independent of t_i . Let $\underline{t}_i \in \arg \min \tau(t_i)$ and $\bar{t}_i \in \arg \max \tau(t_i)$ and suppose that $\max \tau(t_i) > \min \tau(t_i)$. By the continuity of $T_e(\cdot)$ and the definitions of \underline{t}_i and \bar{t}_i , it follows that $\Pi(T_e(\underline{t}_i) + \Delta | \underline{t}_i, B^c(T_e(\underline{t}_i))) < \Pi(T_e(\bar{t}_i) + \Delta | \bar{t}_i, B^c(T_e(\bar{t}_i)))$ for all $\Delta > 0$, and conversely for $\Delta < 0$. Consequently, $h(T_e(\underline{t}_i) | \underline{t}_i, B^c(T_e(\underline{t}_i))) < h(T_e(\bar{t}_i) | \bar{t}_i, B^c(T_e(\bar{t}_i)))$. However, $\beta(\eta\kappa + \tau(\underline{t}_i)) \leq \beta(\eta\kappa + \tau(\bar{t}_i))$. Furthermore, $\theta \int_{t_0 \leq T_e(\underline{t}_i) - \tau_e - \eta\kappa} (\alpha_{t_0, T_e(\underline{t}_i)} \beta(T_e(\underline{t}_i) - t_0)) d\Phi(t_0 | \mathcal{I}_{t_i}) < \theta \int_{t_0 \leq T_e(\bar{t}_i) - \tau_e - \eta\kappa} (\alpha_{t_0, T_e(\bar{t}_i)} \beta(T_e(\bar{t}_i) - t_0)) d\Phi(t_0 | \mathcal{I}_{t_i})$. Thus, the generalized sell-out condition cannot be satisfied for both arbitrageurs \hat{t}_i and \bar{t}_i , a contradiction.

A.6 Proof of Proposition 6

Lemma 7 *There exists a function $\tau_p(\cdot)$ such that for all $n = 1, 2, \dots$ and all histories of price drops H_p^n all traders \hat{t}_i who became aware of the mispricing prior to $t_p - \tau_p(H_p^n)$ leave the market and all other arbitrageurs stay in the stock market.*

Proof. Recall that τ^{***} denotes the time elapsed after which each arbitrageur leaves the market in the absence of a price shock. We proceed inductively, defining $\tau_p(H_p^k)$ given $\tau_p(H_p^l)$ where $l = 1, 2, \dots, k-1$. We are looking for the smallest τ_p such that $t_i = t_p^{(k)} - \tau_p$ is indifferent between exiting and staying in the market. At $t_p^{(k)}$, arbitrageurs with $t_i < t_p - \tau^{***}$ are already out of the market provided that $t_p^{(k)} > \max \left\{ t_p^{(k-j)} + \tau^{***} - \tau_p(H_p^{(k-j)}) \right\}_{j=1,2,\dots}$ (in which case traders returned to the market because of a previous failed sell-out attempt). Arbitrageur $t_i = t_p^{(k)} - \tau_p$ is indifferent between staying in the market or not if

$$\int_{t_0 \leq t_p - \tau_p(H_p^{(k)})} [-\gamma_p + \beta(t_p - t_0)] (1 - \tilde{\alpha}_{t, t_0}) d\Phi(t_0 | t_i = t_p^{(k)} - \tau_p, t_p, H_p) +$$

$$\begin{aligned}
& + \int_{t_0 > t_p - \tau_p}^{(H_p^{(k)})} (-\gamma_p) d\Phi(t_0 | t_i = t_p^{(k)} - \tau_p, H_p(t_p^{(k)})) = 0. \\
& \int_{t_0 \leq t_p - \tau_p}^{(H_p^{(k)})} [\beta(t - t_0)] [1 - \tilde{\alpha}_{t, t_0}] d\Phi(t_0 | t_i = t_p^{(k)} - \tau_p, H_p(t_p^{(k)})) - \gamma_p = 0.
\end{aligned}$$

Let's define the LHS by $f_p^{(k)}(\tau_p | t_i = t_p^{(k)} - \tau_p)$. Note that if the bubble does not burst, then all arbitrageurs sell their shares at the price of $(1 - \gamma_p) p_t$. If the bubble does burst, only the first $\left[\kappa - \frac{1}{\eta} (t_p^{(k)} - t_0 - \tau^{***}) \right]$ orders (κ orders) are executed at $(1 - \gamma_p) p_t$ if $t_p^{(k)} \geq (<) \max_{j=1,2,\dots} \left\{ t_p^{(k-j)} + \tau^{***} - \tau_p \left(H_p^{(k-j)} \right) \right\}$. The term $(1 - \tilde{\alpha}_{t, t_0})$ reflects this fact. If $f_p^{(k)}(0 | t_i = t_p^{(k)} - \tau_p) \geq 0$, then $\tau_p \left(H_p^{(k)} \right) = 0$. If $f_p^{(k)}(0 | t_i = t_p^{(k)} - \tau_p) < 0$ look for $\min \tau_p$ such that $f_p^{(k)}(\tau_p | t_i = t_p^{(k)} - \tau_p) \geq 0$. If such a τ_p does not exist set $\tau_p = t_p$.

It remains to check that all arbitrageurs with $t_i < t_p^{(k)} - \tau_p \left(H_p^{(k)} \right)$ strictly prefer to leave the market and all traders with $t_i > t_p^{(k)} - \tau_p \left(H_p^{(k)} \right)$ prefer to remain in the market. By looking at trader \hat{t}_i 's distribution $\Phi(t_0 | \cdot)$, it is easy to check that the $f_p^{(k)}(\tau_p | t_i = t_p^{(k)} - \tau_p)$ is decreasing in t_i . Notice that each trader can rule out any $t_0 < t_i - \eta$. From the fact that the bubble did not burst before $t_p^{(k)}$ for endogenous reasons all arbitrageurs can rule out $t_0 < t_p^{(k)} - \tau^{***} - \eta\kappa$. Finally, since the bubble survived all sell-out attempts after previous price drops $t_0 \geq \max_{j=1,2,\dots} \left\{ t_p^{(k-j)} - \tau_p \left(H_p^{(k-j)} \right) \right\}$. For all traders with $t_i - \eta \leq \max \left\{ t_p^{(k)} - \tau^{***} - \eta\kappa, \max_{j=1,2,\dots} \left\{ t_p^{(k-j)} - \tau_p \left(H_p^{(k-j)} \right) \right\} \right\}$, the lower bound is the same, while the upper bound is given by t_i . Hence, for arbitrageurs with lower t_i the density on t_0 for which $t_0 \leq t_p - \tau_p \left(H_p^{(k)} \right)$ is higher. In addition, the (conditional) exponential distribution has the nice property that the relative likelihood of possible states in the support across arbitrageurs is the same. The distribution $\Phi(\cdot)$ for arbitrageurs whose support is $[t_i - \eta, t_i]$ is totally symmetric except that it start at different t_i . Hence, $f_p^{(k)}(\tau_p | t_i = t_p^{(k)} - \tau_p)$ is also decreasing in t_i in this case. ■

Proof of Proposition 6.

After establishing the critical value $\tau_p \left(H_p^k \right)$ in Lemma 8, the remaining proof of Proposition 6 is analogous to the proof of Proposition 5. There are three differences: all arbitrageurs observe each price drop, the first orders after the price drop are only executed at a price of $(1 - \gamma^p) p_t$ instead of the pre-crash price p_t , and all $\tau_p \left(H_p^k \right)$ are history dependent. Note that $\tau_p \left(H_p^k \right)$ is known to all traders in equilibrium since all arbitrageurs can observe the past price process. ■

References

- ALLEN, F., AND G. GORTON (1993): “Churning Bubbles,” *Review of Economic Studies*, 60, 813–836.
- ALLEN, F., S. MORRIS, AND A. POSTLEWAITE (1993): “Finite Bubbles with Short Sale Constraints and Asymmetric Information,” *Journal of Economic Theory*, 61, 206–229.
- BARBERIS, N., AND R. THALER (forthcoming): “A Survey of Behavioral Finance,” in *Handbook of the Economics of Finance*, ed. by G. Constantinides, M. Harris, and R. Stulz. Elsevier Science B.V., Amsterdam.
- BLANCHARD, O. J., AND M. W. WATSON (1982): “Bubbles, Rational Expectations, and Financial Markets,” in *Crisis in the Economic and Financial Structure*, ed. by P. Wachtel, pp. 295–315. Lexington, Lexington, MA.
- BRUNNERMEIER, M. K. (2001): *Asset Pricing under Asymmetric Information: Bubbles, Crashes, Technical Analysis and Herding*. Oxford University Press, Oxford.
- BURDZY, K., D. M. FRANKEL, AND A. PAUZNER (2001): “Fast Equilibrium Selection by Rational Players Living in a Changing World,” *Econometrica*, 69(1), 163–189.
- CARLSSON, H., AND E. VAN DAMME (1993): “Global Games and Equilibrium Selection,” *Econometrica*, 61(5), 989–1018.
- CHAMLEY, C. (1999): “Coordinating Regime Shifts,” *Quarterly Journal of Economics*, 114(3), 869–905.
- DANIEL, K., D. HIRSHLEIFER, AND A. SUBRAHMANYAM (1998): “Investor Psychology and Security Market Under- and Overreaction,” *Journal of Finance*, 53(6), 1839–1885.
- DELONG, J. B., A. SHLEIFER, L. H. SUMMERS, AND R. J. WALDMANN (1990a): “Noise Trader Risk in Financial Markets,” *Journal of Political Economy*, 98(4), 703–738.
- (1990b): “Positive Feedback Investment Strategies and Destabilizing Rational Speculation,” *Journal of Finance*, 45, 379–395.
- FAMA, E. F. (1965): “The Behavior of Stock-Market Prices,” *Journal of Business*, 38, 34–105.
- FRANKEL, D., AND A. PAUZNER (2000): “Resolving Indeterminacy in Dynamic Settings: The Role of Shocks,” *Quarterly Journal of Economics*, 115(1), 283–304.

- FROOT, K. A., D. S. SCHARFSTEIN, AND J. C. STEIN (1992): “Herd on the Street: Informational Inefficiencies in a Market with Short-Term Speculation,” *Journal of Finance*, 47, 1461–1484.
- HALPERN, J. Y., AND Y. MOSES (1990): “Knowledge and Common Knowledge in a Distributed Environment,” *Journal of the Association of Computing Machinery*, 37, 549–587.
- HIRSHLEIFER, D. (2001): “Investor Psychology,” *Journal of Finance*, 56(4), 1533–1597.
- KEYNES, J. M. (1936): *The General Theory of Employment, Interest and Money*. Macmillan, London.
- KINDLEBERGER, C. P. (1978): *Manias, Panics and Crashes: A History of Financial Crises*. Macmillan, London, 3rd edition, 1996.
- MATSUYAMA, K. (1991): “Increasing Returns, Industrialization, and Indeterminacy of Equilibrium,” *Quarterly Journal of Economics*, 106(2), 617–650.
- MORRIS, S. (1995): “Co-operation and Timing,” *CARESS Working paper 95-05*.
- MORRIS, S., AND H. SHIN (1998): “Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks,” *American Economic Review*, 88(3), 587–596.
- OBSTFELD, M. (1996): “Models of Currency Crises with Self-Fulfilling Features,” *European Economic Review*, 40, 1037–1047.
- ODEAN, T. (1998): “Volume, Volatility, Price, and Profit When All Traders Are Above Average,” *Journal of Finance*, 53(6), 1887–1934.
- SANTOS, M. S., AND M. WOODFORD (1997): “Rational Asset Pricing Bubbles,” *Econometrica*, 65(1), 19–57.
- SHILLER, R. J. (2000): *Irrational Exuberance*. Princeton University Press, Princeton, New Jersey.
- SHLEIFER, A. (2000): *Inefficient Markets - An Introduction to Behavioral Finance*. Oxford University Press, Oxford.
- SHLEIFER, A., AND R. W. VISHNY (1997): “The Limits of Arbitrage,” *Journal of Finance*, 52(1), 35–55.
- THALER, R. H. (1991): *Quasi Rational Economics*. Russell Sage Foundation, New York, NY.

TIROLE, J. (1982): "On the Possibility of Speculation under Rational Expectations,"
Econometrica, 50, 1163–1182.