Finite Bubbles, Infinite Bubbles, and Crypto Assets*

Luis Araujo
Michigan State University and
São Paulo School of Economics - FGV

Antonio Doblas-Madrid
Michigan State University

April 25, 2026

Abstract

This paper develops micro-foundations in a bubble riding model in the style of Abreu and Brunnermeier (2003; AB). Providing an explicit specification of the model's primitives, especially in the post-crash phase yields insights that help us bridge distinct classes of models of bubbles. The literature is often divided between infinitely lived "rational bubbles" and finite "greater fool" bubbles driven by asymmetric information and lack of common knowledge. We extend the Doblas-Madrid (2012) model—itself a rational, market-clearing version of AB—by explicitly modeling dividends, which depend on technology, from prices, which depend on budget constraints and optimizing behavior. Our analysis, while parsimonious, delivers sharp results. If endowments grow faster than dividends, the model lies in the domain of infinite bubbles; if dividends grow faster than endowments, the model becomes effectively finite, as is typical in the "greater fool" strand of literature. The AB case emerges as a knife-edge in which the two growth rates coincide, clarifying the model's position in the literature. We also discuss how the different kinds of models can be applied to analyze zero dividend assets such as cryptocurrencies.

Keywords: Bubbles, Speculation, Cryptocurrencies
JEL Codes G01, G12, G14

*The authors would like to thank Yu Awaya, Gadi Barlevy, John Conlon, Feng Liu, and participants in the Greater Fool Bubbles Seminar series, Midwest Macroeconomics Meetings, Midwest Theory Meetings, and Meetings of the Southwest Financial Association for useful comments and suggestions. All errors are our own.

1 Introduction

The Great Recession renewed scholarly and policy interest in asset price bubbles. Recently, 'bubble talk' has resurfaced, especially as equity markets have soared and given the continued proliferation of crypto assets. Given their volatility and lack of dividends, these new assets certainly pique the curiosity of economists interested in monetary and bubble theory.

A well known model of speculative bubbles is the one developed by Abreu and Brunnermeier (2003; AB henceforth). It is often credited with capturing the dynamics described in narratives of manias and panics à la Kindleberger (1978) and Minsky (1986). While arising in disparate settings—from Tulipmania to roaring 20s, dotcom, subprime, crypto, etc.—boom-bust episodes nevertheless tend to unfold through strikingly similar stages. An initial impulse sets off a wave-like process that is at first self-reinforcing but later turns self-destructive as the gap between price and fundamental value grows. To many, this popular narrative rings true, which is why it is interesting to scrutinize the theory more closely, comparing it to established rational models and recognize the links with other strands of theory. This seems like an impossible task starting from the behavioral AB version. Doblas-Madrid (2012; DM henceforth) bridged some of the gap between AB and standard equilibrium models by building a rational, market clearing framework with infinite horizon. Still, before taking comparisons any further, more work is needed to bring micro foundations to a level comparable with the rational models of reference. Specifically, in this paper we open the 'black box' of the post-crash price, and explore the implications of different specifications of the dividend process.

The AB/DM model provides only a very vague description of what happens after a bubble bursts. The price is simply supposed to fall by a certain percentage that is an increasing function of bubble duration, and there is a maximum exogenous bubble duration. However, these are exogenous ingredients of the model, not derived from dividends or maturity dates. The model is built around t_0, defined as the random time when an initially fundamental boom goes too far and agents begin to become aware of this fact. Signals and beliefs are also functions of prices in the periods after t_0. This raises two issues. First, in the context of bubbles—especially finite ones—a vague definition of the final stage leaves the ever present question of backward induction unanswered. Second, intertwining primitives with endogenous variables subjects comparative statics to a Lucas critique. We cannot, for example, disentangle the price growth rate—a function of endowments, constraints, and the equilibrium—from the distribution of dividends, an entirely different variable conceptually.[1]

[1] In Doblas-Madrid (2016) the post-crash payment of dividends and resolution of fundamental uncertainty are modeled, but in a special finite-horizon case with lack of common knowledge, and still subject to the same Lucas critique.

To address these issues we build a more complete model, which borrows some structure from DM, but makes explicit assumptions about the dividend process that are independent from endowments. To clear the fog about what is exogenous or endogenous, we specify probability distributions over dividends d_t and maturity (i.e., dividend payout) dates t_m, defining fundamental value simply as the present value of dividends. The distribution of maturity dates ranges from a minimum $\underline{t_m}$ to infinity, and d_t is a function of t_m parameterized by a growth rate D. Private signals are defined around t_m. That is, if the realization of the maturity date is t_m, there are N different types with private signals $\nu_m \in \{t_m, ..., t_m + N - 1\}$. The bubble starting time t_0 is now an endogenous outcome that emerges in equilibrium.

Having disentangled dividends from prices, we can now analyze the model without confounding concepts. Our first benchmark is the special case where the growth rate of dividends D equals the growth rate of endowments G. This is the only parametrization consistent with AB's assumption that the difference between a bubble's starting time t_0 and its exogenous bursting time—the maturity date t_m in our setting—is independent of t_0. Since we assume rational agents, our initial, benchmark analysis effectively replicates the baseline case in DM. In this case, the distance between t_m and t_0 is always the same and equal to $\overline{\tau}$. Similarly, given an agent's signal ν_m about t_m it is straightforward to construct an equivalent signal $\nu_0 = \nu_m - \overline{\tau}$ about t_0. An equilibrium can equivalently be described as riding the bubble for τ^* periods after the ν_0 signal or preempting the exogenous crash by selling ς^* periods before the maturity date signal ν_m. These equilibria are symmetric, in the sense that all types have the same bubble riding—and, equivalently, preemption times—relative to their signal.

When dividends grow faster than endowments, i.e., when $D > G$, the model is effectively finite. Bubbles can only grow and burst before a certain date $\overline{t_m}$ that is pinned down by parameters, and thus, is common knowledge. For realizations of t_m above $\overline{t_m}$, fundamental value is so large that the price never catches up with it. Even if agents remain—preference shocks permitting—fully invested in the booming market until the maturity date, this date arrives and a dividend is paid that exceeds the price. This makes the model effectively finite, since bubbles may only arise for realizations of t_m below $\overline{t_m}$. The lower the realization of t_m, the bigger the scope for bubbles. We characterize maximum bubble duration for those low realizations of t_m. Types below, but close to $\overline{t_m}$, have 'less room' to ride bubbles. Therefore, a symmetric bubble riding equilibrium in the style of AB or DM does not adequately characterize the set of possible bubbly equilibria. By contrast, a symmetric crash preemption equilibrium anchored on the signal ν_m, where agents sell at time $\nu_m - \varsigma^*$ more effectively characterizes the set of possible bubbly equilibria. With the distance between t_m and t_0 shrinking as t_m increases, a constant crash preemption time implies a shrinking bubble riding time. With all the relevant action happening below $\overline{t_m}$, the model is effectively finite.

While the particulars of the environments differ, conceptually our model in this scenario becomes a finite 'greater fool' bubble model. In this strand of literature, we find the series of models starting with Allen, Morris, and Postlewaite (1993) and further developed by Conlon (2004, 2015), Liu and Conlon (2018), Liu, White, and Conlon (2023), and Liu (2026). These papers show how rational, finite, asymmetric information bubbles can arise with the minimum number of agents and states of the world. Another prominent contribution in this strand of literature is Awaya, Iwasaki, and Watanabe (2022) who, in a parsimonious setting, show how greater fool bubbles arise as a unique equilibrium with a chain of middlemen. As in our model, in this strand of literature, finite bubbles arise because of a combination of asymmetric information and gains from trade.

When $D < G$, there is no state of the world in which fundamental values exceed prices. Finite truncation is impossible and the maturity date must have infinite support to preclude backward induction. Sequential rationality requires iterating the argument that buyers believe they may find others to whom they can resell the asset at a profit. This logic must iterate to infinity since there is no state of the world, no higher order belief, where a final buyer would acquire the asset for its fundamental value with no intent to resell. If dividends grow less slowly than endowments, signals about ν_0 are 'compressed'. The limit case where dividends are zero is particularly interesting, since the existence of a bubble becomes common knowledge. If signals about t_0 are compressed to the point of being the same for all types, uncertainty about the start of the bubble vanishes, and equilibria with symmetric τ^* cease to exist, since the bursting date becomes predictable. Equilibria with symmetric ς^* still exist—implying a bubble riding time $\tau^* = \nu_m - \varsigma^*$ that is asymmetric and unbounded—as long as the bubble grows fast enough to compensate for a total loss in the event of a crash. As an infinite bubble with common knowledge of overvaluation, this type of equilibrium no longer resembles those in 'greater fool' bubbles, which are finite and based on lack of common knowledge of overvaluation. The bubble is still speculative in nature, but its infinite duration and common knowledge of its existence makes it more similar to what much of the literature defines as a 'rational bubble' following the well known contributions by, among others, Samuelson (1958), Tirole (1985), Santos and Woodford (1997), Martin and Ventura (2012), Hirano and Yanagawa (2017), and Miao and Wang (2018).

The zero dividend case gives us a framework to think about crypto assets, which due to their volatility and lack of dividends, seem like the quintessential bubbles. In our model, a crypto asset that experiences a dramatic boom–bust cycle and subsequently disappears naturally corresponds to the $D = 0$ case, i.e., zero fundamental value. Our model, however, can also be reinterpreted to analyze other crypto assets—like, famously, Bitcoin—which alternate periods of boom and bust with phases where prices, while still volatile, drifts sideways without clear trend. To see how our model can be adapted to such cases, consider the case in which dividends are zero but f_t¿0 because the asset continues to serve as a store of value after the crash. Even without dividends, the risky asset can become a perfect substitute for the safe asset, provided its price is expected to grow at rate R in perpetuity. Under this interpretation, the distribution of signals and the meaning of t_0 may remain unchanged, while t_m becomes the date at which the 'steady state value' f_t is publicly revealed rather than a dividend payout date.

The paper is organized as follows. In the next section we present the environment. In Section 3, we describe bubbly equilibria in the baseline case. In Section 4, we consider the effectively finite model. Section 5 considers infinite bubbles and crypto assets. Section 6 concludes.

2 Model

2.1 Environment

Time is discrete and infinite with periods labeled $t \in \mathbb{N}_0$. There is a unit mass continuum of agents $i \in [0, 1]$, who are risk neutral and may be hit by preference shocks. Every period, with probability $\theta \in (0, 1)$ agent i is impatient and has discount factor of zero. With probability $1 - \theta$, the agent is patient and has discount factor $\delta \in (0, 1)$. There is a safe asset with exogenous gross return $R > 0$ and a risky asset that exists in unit supply. In period t, all agents collect an endowment of $e_t = G^t$ units of the safe asset. We assume that $\delta R \geq 1$ and $G > R$.[2] Agent i's holdings of the safe and risky assets at the beginning of period t are denoted by $b_{i,t-1}$ and $h_{i,t-1}$, respectively. Initial holdings are given by $b_{i,-1} = 0$ and $h_{i,-1} = 1$ for all i.

The risky asset pays a single dividend d_t at maturity date t_m. To capture AB's idea that fundamental value comes from dividends paid in a relatively distant future, we assume that t_m cannot occur earlier than a lower bound $\underline{t_m} > 0$. Since the dividend is observable, at this point, the fundamental value of the risky asset becomes public, if it has not been endogenously revealed before.[3] The maturity date is random and distributed geometrically with success probability parameter $1 - \lambda \in (0, 1)$. Specifically, the cdf of t_m is given by

$$\Phi(t_m) = 1 - \lambda^{t_m - \underline{t_m} + 1}, \text{ for all } t_m \in \{\underline{t_m}, \underline{t_m} + 1, ...\}, \tag{1}$$

[2] If $\delta R < 1$ formulas would have to be changed to reflect the fact that agents would never save using the safe asset, but in essence, all our results would still obtain.

[3] For simplicity, we assume a single dividend disbursement date t_m and that information about fundamental value is not exogenously revealed earlier. It would be equivalent to assume that dividends of equal present value are paid at one or more dates after t_m, as long as that present value became public knowledge at t_m.

with corresponding pmf $\varphi(t_m) = (1-\lambda)\,\lambda^{t_m - \underline{t_m}}$. Assuming that $\gamma > 0$, we define dividends as a function of t_m as follows:

$$d_t = \begin{cases} \gamma D^{t_m - \underline{t_m}} R^{\underline{t_m}} & \text{if } t = t_m \text{ and } D > 0 \\ 0 & \text{if } t \neq t_m \text{ or } D = 0. \end{cases} \tag{2}$$

Thus, at time t_m, the risky asset pays a dividend of $\gamma D^{t_m - \underline{t_m}}$ multiplied by $R^{\underline{t_m}}$ so that the lower bound $\underline{t_m}$ has no effect on present value. Given a realization of t_m, the fundamental value f_t grows at the risk-free rate before t_m, then falls to zero when the dividend is paid. Hence,

$$f_t = \begin{cases} \gamma \left(\frac{D}{R}\right)^{t_m - \underline{t_m}} R^t & \text{if } t \leq t_m \\ 0 & \text{if } t > t_m. \end{cases} . \tag{3}$$

The unconditional distribution of $t_m \in \{\underline{t_m}, \underline{t_m}+1, ...\}$ is common knowledge, but there is asymmetric information about the realization of t_m. We assume that t_m is realized at the start of time 0, and each agent observes a signal $\nu_m \in \{t_m, ..., t_m + N - 1\}$. These signals define N different types, evenly distributed in the sense that there is a mass $1/N$ of agents of each. Furthermore, an agent observing signal ν_m can also discard values of t_m below $\nu_m - (N-1)$ and above ν_m. Thus, the conditional pmf $\varphi(t_m|\nu_m)$ is given by

$$\varphi(t_m|\nu_m) = \begin{cases} \frac{\varphi(t_m)}{\varphi(\max\{\underline{t_m}, \nu_m - (N-1)\}) + \cdots + \varphi(\nu_m)} & \text{if } \max\left\{\underline{t_m}, \nu_m - (N-1)\right\} \leq t_m \leq \nu_m \\ 0 & \text{otherwise.} \end{cases} . \tag{4}$$

Signals order agents along a line, without them knowing where it starts—except for the special case where t_m is $\underline{t_m}$ in which case agents of type $\nu_m = \underline{t_m}$ know they are first in line. This dimension of uncertainty is crucial to generate bubbles in all AB-style models. Agents do not know if they are more or less optimistic relative to others, and it is because of this uncertainty that they cannot coordinate on a common expected bursting date. Thus, the well-known backward induction argument that defuses finite bubbles in symmetric information environments does not apply, since there is no common expected bursting date to iterate backward from.

After receiving endowments and learning the realization of their preference shocks, agents have a chance to buy or sell risky shares at a Shapley-Shubik trading post.[4] For all $t \leq t_m$, the market clearing protocol is as follows:

Step 1: Agents submit market orders In a first stage, agents submit orders to buy or sell. To buy, they submit $m_{i,t} \in [0, b_{i,t-1} + e_t]$ units of the safe asset which they wish to spend buying risky shares. To sell, they submit $s_{i,t} \in [0, h_{i,t-1}]$ risky shares for sale.

[4] Araujo and Doblas-Madrid (2026) develop a version of the model with multidimensional uncertainty with Walrasian market clearing.

Agents submit orders before observing others' orders or the price, which will be observable later. In other words, only *market orders* are available.

Step 2: The market clears Orders are combined and the price is given as the ratio of the total amount of safe assets bid M_t to the total mass of risky shares offered for sale S_t and is thus given by

$$p_t = \frac{M_t}{S_t},$$

Agent i leaves the market with

$$h_{i,t} = h_{i,t-1} - s_{i,t} + \frac{m_{i,t}}{p_t} \tag{5}$$

risky shares. Safe balances are given by the initial amount $b_{i,t-1} + e_t$ plus any proceeds from selling risky shares $p_t s_{i,t}$ minus $m_{i,t}$, which is the amount of the safe asset spent buying risky shares. After the asset market closes, but only at $t = t_m$, risky shares pay a dividend $d_t h_{i,t}$. At the end of the period, agent i chooses to consume or save subject to

$$c_{i,t} + \frac{b_{i,t}}{R} = b_{i,t-1} + e_t + p_t s_{i,t} - m_{i,t} + d_t h_{i,t}, \tag{6}$$

and non-negativity constraints for both consumption $c_{i,t} \geq 0$ and end-of-period safe balances $b_{i,t} \geq 0$.

After period t_m the risky asset is worthless and stops trading, with agents simply consuming if impatient and adding to risk-free savings if patient.

To recap, agent i starts period t with $(b_{i,t-1}, h_{i,t-1})$, collects endowment e_t and learns if they are patient or impatient. The agent may then trade the risky for the safe asset in a Shapley-Shubik-type market. After trading, if $t = t_m$, the agent collects dividend $d_t h_{i,t}$. At the end of the period, agent i may consume or save whatever amount of the safe asset she holds after trading and, possibly, collecting dividends.

2.2 From the maturity date t_m to the overvaluation date t_0

At time 0, a boom begins. Patient investors, of which there is a mass $1 - \theta$, buy as many shares of the risky asset as they can. Unable to borrow against future endowments, they spend their entire current endowment G^t buying new risky shares to add to their existing holdings. At this stage, the only sellers are impatient investors, who liquidate their entire portfolios. Since preference shocks are i.i.d., the total mass of impatient agents and the total mass of shares for sale both equal θ. Hence, while zero types are selling, the booming price p_t^0 is given by

$$p_t^0 = \frac{1-\theta}{\theta} G^t. \tag{7}$$

The boom is initially justified by fundamentals. For some time, prices are suppressed by wealth, and thus the price p_t^0 is below f_t, as endowments gradually grow, allowing prices to catch up to the present value of dividends. This fundamental part ends at time t_0. If the boom continues past this point, the market becomes overvalued, starting the transition into the bubbly part of the boom. Period t_0 is defined as

$$t_0 \equiv \min \left\{ t | t \in \mathbb{N}_0 \text{ and } p_t^0 > f_t \right\}. \tag{8}$$

To avoid inessential complications, we assume that dividends are large enough to make the fundamental part of the boom non-degenerate. That is, we assume that the time-0 price p_0^b does not already surpass the lowest possible value of f_0, corresponding to the realization $t_m = \underline{t_m}$. Thus, we restrict attention to parameters such that

$$\gamma \geq \frac{1-\theta}{\theta}, \tag{9}$$

which avoids 'corner solutions' that set t_0 to zero because it cannot take negative values. Given this, we substitute p_t^b and f_t for their values and solve for t_0 to find

$$t_0(t_m) = \left\lceil \frac{(t_m - \underline{t_m}) \log(D/R) + \log(\gamma\theta/(1-\theta))}{\log(G/R)} \right\rceil, \tag{10}$$

where the ceiling function $\lceil \cdot \rceil$ rounds its argument up to the nearest integer. This formula also converts signals ν_m about t_m into signals ν_0 about t_0. Simply substitute ν_m in the right-hand side instead of t_m to obtain the corresponding value of $\nu_0 = \nu_0(\nu_m)$.

We define the minimum possible value of t_0, as $\underline{t_0} = t_0(\underline{t_m})$. Since assumption (9) guarantees that $\underline{t_0}$ is nonnegative, the probability distribution of t_0 is given by

$$\Phi_0(t_0(t_m)) = \Phi_m(t_m) \text{ for all } t_m \in \{\underline{t_m}, \underline{t_m}+1, ...\},$$

with corresponding pmf

$$\varphi_0(t_0) = \Phi_0(t_0) - \Phi_0(t_0 - 1) \text{ for all } t_0 > \underline{t_0}, \tag{11}$$

and $\varphi_0(\underline{t_0}) = \Phi_0(\underline{t_0})$.

3 Type Independent Equilibria

In this section, we analyze the baseline case where $D = G$. This allows us to simplify $t_0(t_m)$ from (10) to

$$t_0(t_m) = t_m \underbrace{- \underline{t_m} + \left\lceil \frac{\log(\gamma\theta/(1-\theta))}{\log(G/R)} \right\rceil}_{=\overline{\tau}}. \tag{12}$$

Note that the difference between t_m and t_0 is simply a constant $\overline{\tau}$, which in the special case where $\gamma = (1-\theta)/\theta$, the constant $\overline{\tau}$ further simplifies to $\underline{t_m}$. In other words, for any realization of the maturity date t_m, the first overvaluation period t_0 is simply $\overline{\tau}$ periods before. Similarly, $\nu_0 = \nu_m - \overline{\tau}$ for any signal ν_m.

Note that, while AB do not define dividends or fundamental value, they do assume that the market's overvaluation starts at time t_0 and that, if it does not burst endogenously before, the bubble bursts at time $t_0 + \overline{\tau}$ for exogenous reasons. Given (10), the only parametrization of our model consistent with this assumption is $D = G$. Figure 1 depicts the relationship between t_m and t_0.

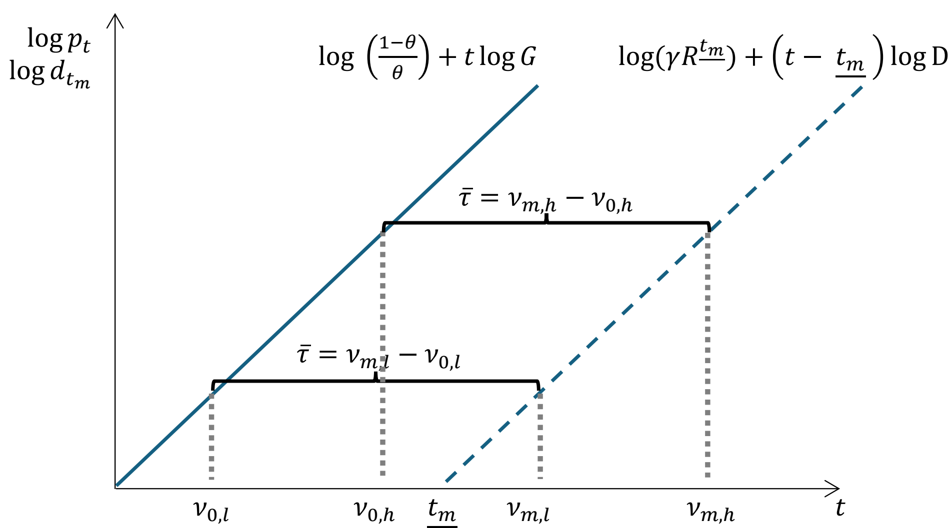

Figure 1: In our baseline case $D = G$, and thus the difference between t_m and $t_0(t_m)$ is always constant and equal to $\overline{\tau}$, which in this example equals $\underline{t_m}$. AB's assumption that, for any overvaluation starting time $t_0 \geq 0$, the bubble bursts for exogenous reasons at $t_0 + \overline{\tau}$ is consistent with this case, although in our model $\overline{\tau}$ is derived endogenously given prices and dividends.

3.1 Equivalence of bubble riding and crash preemption

In this section, we follow AB and DM constructing bubbles in which agents' strategies anchored in the overvaluation time t_0, or more precisely, their signal ν_0. That is, we analyze bubbly equilibria where agents' planned selling dates equal their signal ν_0 about plus a number of bubble riding periods τ.

While this is intuitive, and consistent with AB and DM, it is important to note, planting a seed for future sections, that riding the bubble for τ periods is equivalent to preempting the exogenous crash, i.e., the maturity date, by ς periods. That is, the planned selling time $\nu_0+\tau$ can alternatively be written as the signal about maturity ν_m minus a crash preemption time ς, so that

$$\nu_0 + \tau = \nu_m - \varsigma,$$

where the fact that $\nu_m - \nu_0 = \overline{\tau}$ implies that

$$\tau + \varsigma = \overline{\tau}.$$

Keeping this in mind, we now proceed to describing agents' strategies in terms of their signal ν_0 and bubble riding time τ.

Given asymmetric information about fundamentals, different agents observe different singals about when the overvaluation starts. Thus, the boom phase is subdivided into three parts:[5]

1. A *fundamental part* from 0 to $t_0 - 1$, where prices are justified by fundamentals in reality, as well as in the estimation of all agents.

2. A *semi-bubble* from t_0 to $t_0 + N - 2$, where the market is overvalued, but some agents are aware of this and others are not.

3. A *bubble* from period $t_0 + N - 1$ until the end of the boom, where all agents are aware that prices exceed fundamental value. That is, overvaluation is mutual knowledge.[6]

Given the above distinction between a semi-bubble and a bubble, it is necessary that $\overline{\tau} \geq N - 1$ so that bubbles are not impossible by assumption.

A boom can end with an endogenous or an exogenous crash. In an endogenous crash, some patient agents sell at time $t < t_m$. For example, suppose that all agents follow a

[5]This terminology follows Allen, Morris, and Postlewaite (1993) and followup papers, most recently, Liu, White, and Conlon (2023).

[6]Such a bubble satisfies Allen et al.'s (1993) definition of a *strong* bubble as a situation where the price exceeds the highest possible fundamental value of the most optimistic agent. That is, for a given t_m, the highest possible signal is $\nu_m = t_m + N - 1$.

symmetric rule: sell τ periods after their overvaluation signal ν_0, i.e, ride the bubble for τ periods and then sell. If $\tau \leq \overline{\tau}$, then a mass $1/N$ of type-t_0 agents sells at time $t_0 + \tau$. This increases the mass of sellers—and equivalently reduces the mass of buyers—by $(1\theta)/N$. Then, the price is given by

$$p_t^1 = \frac{(1-\theta)(1-1/N)}{\theta + (1-\theta)/N} G^t, \tag{13}$$

where p_t^1 indicates that one type is selling.

Because the market clearing protocol requires agents to submit orders before seeing the price, they cannot respond to the information conveyed by these sales until the following period. By then, all types know that the bubble has burst. Consequently, the price falls to fundamental value at time $t+1$, which becomes common knowledge after the crash. After the crash, from $t+1$ until t_m, the asset trades at fundamental value. At t_m, after trading, the dividend is paid and price falls to zero.

An exogenous crash occurs if, for all $t \leq t_m$, including t_m, the price p_t equals p_t^0 and the dividend is paid at t_m.

We assume that the sales by one type convey information but are otherwise too small to have a sizable effect on prices, that is, we assume that N is large enough so that p_t^1 is not far below p_t^0. Certainly, we assume that

$$p_t^1 > p_{t-1}^0 R. \tag{14}$$

3.2 Equilibrium

Since impatient agents always sell their assets, setting $s_{i,t} = h_{i,t-1}$ and $m_{i,t} = 0$ and consuming $c_{i,t} = b_{i,t-1} + e_t + p_t s_{i,t}$, we will henceforth focus on the choices of patient agents, who never consume, and refer to them simply as agents.

An equilibrium consists of asset market orders $(m_{i,t}, s_{i,t})$ for all i and all $t \leq t_m$ that maximize expected utility given the price history up to time $t-1$ and the signal ν_0. Equilibrium is found via a guess and verify procedure.

We guess a candidate equilibrium where agents with overvaluation signal at time ν_0 plan to ride the bubble for τ^* periods and sell at $\nu_0 + \tau^*$. Specifically, their orders while $t \leq t_m$ are given by

I. During the boom, i.e., as long as $p_r = p_r(\theta, 0)$ for all $r < t$, orders are given by:

$$(m_{i,t}, s_{i,t}) = \begin{cases} (b_{i,t-1} + G^t, 0) & \text{if } t < \nu_0 + \tau^* \\ (0, h_{i,t-1}) & \text{if } t = \nu_0 + \tau^*. \end{cases} \tag{15}$$

II. Post-crash, i.e., if for any $r < t$, $p_r < p_r(\theta, 0)$, orders are given by:

$$(m_{i,t}, s_{i,t}) = \left(\frac{\theta}{1-\theta} f_t, 0 \right). \tag{16}$$

Since the risky asset becomes worthless after the dividend is paid, after t_m the price falls to zero and the risky asset no longer trades.

If all agents follow (15) and (16), the booming price $p_t(\theta, 0)$ is observed up to—and including—period $t_0 + \tau^* - 1$. When agents submit asset market orders in period $t_0 + \tau^*$ they do so before observing the price p_t. Only when the market clears, all agents observe $p_t(\theta, 1)$, understand that sales have begun, and infer that t_0 must be equal to $t - \tau^*$.

If $\tau^* = 0$, agents sell as soon as they realize the market is overvalued. In that case, there is no bubble riding behavior at all. If $\tau^* > 0$, agents continue knowing that the market is overvalued, giving rise to a semi-bubble if $\tau^* < N - 1$ and a bubble if $\tau^* \geq N - 1$.

Now consider how agents learn in periods leading up to the crash. The signal ν_0 reveals to agents of that type that t_0 must be at least $\nu_0 - (N - 1)$ and at most ν_0. Thus, before period $\nu_0 - (N - 1) + \tau^*$ they can 'safely' ride the bubble without crash risk. At time $\nu_0 - (N - 1) + \tau^*$ there starts to be a possibility, if $t_0 = \nu_0 - (N - 1)$, that the bubble will burst. At first the crash probability is only $1/(1 + \lambda + \cdots + \lambda^{N-1})$, as there are N possible values of t_0 and the crash only occurs if t_0 is the first value. If the price continues to grow at the rate G, agents learn that t_0 is not $\nu_0 - (N - 1)$ and they discard this value from the support of t_0. Going forward, every period that the price growth rate does not slow, agents discard one more value from the low end of the support of t_0 and the crash probability rises. The crash risk is thus highest just one period before type-ν_0 agents are supposed to sell. At time $\nu_0 + \tau^* - 1$ the support of t_0 is just $\{\nu_0 - 1, \nu_0\}$ and the probability of a crash is $1/(1 + \lambda)$.

To understand if bubbly equilibria exist or not, we must examine conditions under which type-ν_0 agents are willing to buy and sell according to (15) before the crash.

Verifying that type-ν_0 agents are willing to sell at time $\nu_0 + \tau^*$ is straightforward. Given the definition of t_0, and assumption (14) for any $\tau^* \geq 0$ the price $p_{t_0}(\theta, 1)$ exceeds fundamental value. Moreover, an individual agent of type ν_0 understands that other agents of her same type are selling and that the price will reveal those sales. Every type-ν_0 agent therefore finds it optimal to sell at time $\nu_0 + \tau^*$.

Verifying the optimality of buying is more involved, as an individual type-ν_0 agent may be tempted to sell preemptively, especially at time $\nu_0 + \tau^* - 1$ to avoid the crash. Let's consider the sell-or-wait tradeoff faced by an agent at this juncture. At this time, t_0 equals $\nu_0 - 1$ with probability $1/(1+\lambda)$ and ν_0 with the remaining probability. Thus, the expected return from selling preemptively is a weighted average of $p_{\nu_0+\tau^*-1}(\theta, 1)$ and $p_{\nu_0+\tau^*-1}(\theta, 0)$ given by

$$\mathbf{E}\left[p_{\nu_0+\tau^*-1} | t_0 \in \{\nu_0 - 1, \nu_0\}\right] = \\ = \left(\frac{1}{1+\lambda} \frac{(1-\theta)(1-1/N)}{\theta + (1-\theta)/N} + \frac{\lambda}{1+\lambda} \frac{1-\theta}{\theta} \right) G^{\nu_0+\tau^*-1}.$$

The alternative to selling preemptively and investing the proceeds at the rate R is to follow (15) and ride the bubble for one more period. Following this strategy implies getting caught in the crash if $t_0 = \nu_0 - 1$ versus being able to sell at a higher price $p^1_{\nu_0+\tau^*}$ if $t_0 = \nu_0$. In expectation, we have

$$\mathbf{E}\left[p_{\nu_0+\tau^*} | t_0 \in \{\nu_0 - 1, \nu_0\}\right] = \\ \frac{1}{1+\lambda} \frac{1-\theta}{\theta} \left(\frac{G}{R}\right)^{\nu_0 - 2} R^{\nu_0+\tau^*} + \frac{\lambda}{1+\lambda} \frac{(1-\theta)(1-1/N)}{\theta + (1-\theta)/N} G^{\nu_0+\tau^*}.$$

Focusing on the case where N is large enough that the difference between p^1_t and p^b_t is small, we have that waiting is optimal as long as

$$RG^{\nu_0+\tau^*-1} \leq \frac{1}{1+\lambda} \left(\frac{G}{R}\right)^{\nu_0-2} R^{\nu_0+\tau^*} + \frac{\lambda}{1+\lambda} G^{\nu_0+\tau^*},$$

which we can rewrite as

$$1 \leq \frac{1}{1+\lambda} \left(\frac{G}{R}\right)^{-(\tau^*+1)} + \frac{\lambda}{1+\lambda} \frac{G}{R}. \tag{17}$$

This condition simply states that, relative to selling, buying yields crash losses with probability $1/(1+\lambda)$ and one more period of growth with the remaining probability. From here, note that if G/R is large enough, buying is optimal even if 100 percent of the investment is lost in the crash, i.e., even as $\tau^* \to \infty$. In this case, bubble duration is only limited by the exogenous upper bound $\overline{\tau}$. Otherwise, the need to recover some of the funds invested in the event of a crash imposes an endogenous upper bound on τ^*. The following Proposition summarizes these results.

Proposition 1. *1.* ***Exogenous upper bound.*** *If $G/R \geq (1+\lambda)/\lambda$, there exists an equilibrium for every $\tau^* \in \{0, 1, ..., \overline{\tau}\}$. Since $\overline{\tau} \geq N - 1$, bubbly equilibria exist.*

2. ***Endogenous upper bound.*** *If $G/R < (1+\lambda)/\lambda$, there is—in addition to $\overline{\tau}$—an endogenous upper bound $\tau^*_{\max}$ given by*

$$\tau^*_{\max} = -\left\lfloor \frac{\log(1+\lambda-\lambda G/R)}{\log G - \log R,} \right\rfloor - 1, \tag{18}$$

where the floor function $\lfloor\cdot\rfloor$ rounds its argument down to the nearest integer. Bubbly equilibria exist if $\tau^*_{\max} \geq N-1$.

Proof. The result has been derived before the statement of the proposition. Type-ν_0 agents are willing to sell at time $\nu_0 + \tau^*$ because other agents of her same type are selling and that the price will reveal those sales. To verify that agents are willing to buy at all times before $\nu_0 + \tau^*$, agents are most tempted to deviate by selling preemptively one period before, at which point their decision is governed by from (17), which determines $\tau^*_{\max}$. □

It is important to note that every equilibrium with τ^* bubble riding periods is also an equilibrium that can be written in terms of ς^* crash preemption periods, where

$$\varsigma^* = \overline{\tau} - \tau^*. \tag{19}$$

In this baseline case where $D = G$, the difference between t_0 and t_m is independent of the realization of t_m. Thus, since $\overline{\tau}$ is independent of t_m, a symmetric equilibrium for a given type-independent bubble riding time τ^* is also a symmetric crash preemption equilibrium with a type-independent crash preemption time ς^*.

While less intuitive and less similar to AB and DM, the crash preemption formulation has the advantage of being more suitable to analyze environments where D is disentangled from G. As explained in the introduction, these are two conceptually different parameters. One relates to the productivity of the technology. The other refers to the resources agents can bring to the market. To be able to conduct comparative statics, we must be able to differentiate the two.

In the following two sections, we consider the cases where $D > G$ and $G < D$. In both, equilibria with symmetric ς^* exist and allow us to preserve some continuity with respect to the baseline knife-edge case of $D = G$ presented in this section, which corresponds to AB and DM. The corresponding bubble riding equilibria, however, are no longer symmetric. If $D > G$, the bubble riding time τ^* must be a decreasing function of t_m; and if $D < G$ it must be an increasing (and unbounded) function. We will consider each case in turn.

4 Effectively Finite Environments

In the previous section, as in AB and DM, although bubbles are finite, the model's horizon is infinite. The distributions of t_m and t_0 both have infinite supports. Thus, agents who are riding the bubble buy in hopes of reselling higher to others, who also plan to resell higher, and so forth ad infinitum. Here, we consider the case where $D > G$, and show that, as far as the possibility of bubbles is concerned, the model is effectively finite. The chain of higher order beliefs about later types eventually ends in a state of the world where a buyer genuinely believes in the asset's fundamental value and buys it to hold it, not to flip it. Because of this, the model in this case looks a lot like the finite greater fool models initially proposed by Allen et. al (1993), and further developed by Conlon (2004, 2015), Liu, White and Conlon (2022), and Liu (2026), as well as Awaya et. al (2022), and others. Because of its effectively finite horizon, this case is also reminiscent of Doblas-Madrid (2016).

Bubbles can still arise for low realizations of t_m, in a way similar to the previous sections, where overvaluation is mutual knowledge but not common knowledge, and impatience shocks give rise to gains from trade.

A key implication of the assumption that $D > G$ is that the distance between t_m and t_0 decreases as t_m increases. Abstracting from rounding to integers, note that the expression inside the $\lceil \cdot \rceil$ bracket in (10) is linear in t_m with slope $\log(D/R)/\log(G/R)$. Thus, there exists a date $\overline{t_m}$ after which bubbles are no longer possible. In fact, if dividends grow faster as a function of the maturity date t_m than the price grows as a function of time, there exists—for bubble formation purposes—an effectively terminal date $\overline{t_m}$ given by the time when dividends first surpass the price. Specifically, his cutoff is defined by

$$\frac{1-\theta}{\theta}G^t = \gamma D^{t-\underline{t_m}} R^{\underline{t_m}}. \tag{20}$$

Solving and rounding up to the nearest integer, we have

$$\overline{t_m} = \left\lceil \frac{\log(1-\theta) - \log(\gamma\theta) + \underline{t_m}\,(\log D - \log R)}{\log D - \log G} \right\rceil. \tag{21}$$

From the point of view of agents with signals $\nu_m \geq \overline{t_m} + N - 1$, bubbles are simply impossible. An agent of type $\nu_m = \overline{t_m} + N - 1$ knows that, even if they are 'last in line' the price can never surpass fundamental value. Endowments simply aren't large enough for any t_m that is possible given ν_m.

For all types such that $\nu_m \geq \overline{t_m}$, it is the case that $\nu_0 > \nu_m$. In these cases, the maturity date arrives before the overvaluation starts, and thus there can be no 'bubble riding'. Moreover, these types would not follow a symmetric equilibrium selling at $\nu_m - \varsigma$. To see why, note that in any such equilibrium, the first type $\nu_m = t_m$ would know, at time

$t_m - \varsigma$ that they are 'first in line' given previous prices. However, they would not sell at this time, since they would know with certainty that the price would be below the present value of the dividend.

The types that are relevant for bubble formation are $\nu_m \in \{\underline{t_m}, \underline{t_m}+1, \ldots, \overline{t_m}-1\}$. Similarly, the time periods that are relevant for bubble formation span from 0 to $\overline{t_m}-1$, making the model effectively finite.

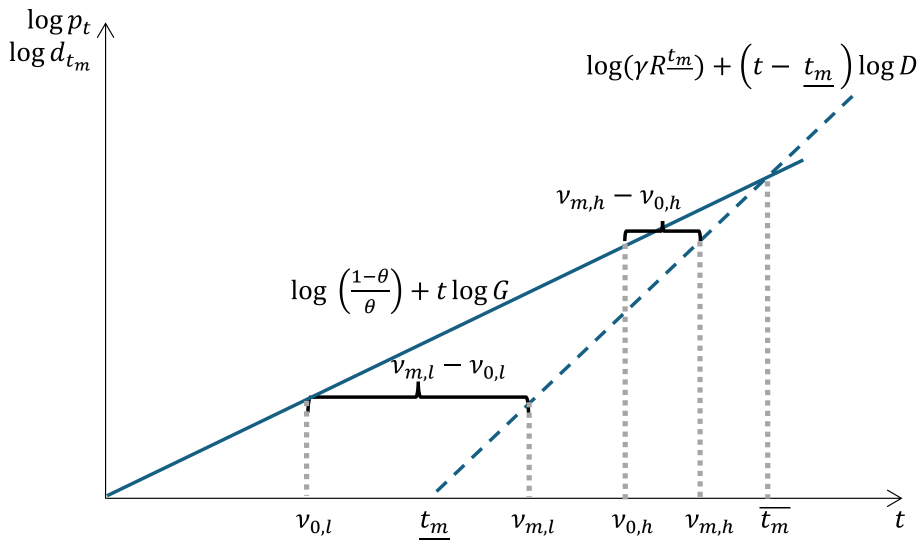

Figure 2: When $D > G$, the model is effectively finite, since there is a period $\overline{t_m}$ after which bubbles simply are not possible. Even in this effectively finite setting, backward induction fails, as agents do not know when the bubble may burst, and even if all agents know that assets are overvalued, there is, in a higher order sense, a 'reachable' state of the world $\overline{t_m}$ with buyers willing to buy the asset for its fundamental value, not to resell at a higher price.

4.1 Bubble riding vs. crash preemption

Whenever $D > G$, the difference $t_m - t_0$ instead of a constant $\overline{\tau}$, is decreasing in t_m. As depicted in Figure 2, the difference largest for $t_m = \underline{t_m}$ and—before rounding—-decreases linearly, vanishing when t_m reaches $\overline{t_m}$. In other words, since dividends grow faster than prices, it takes longer than one period for prices to rise by a factor of D. In fact, the distance between $t_0(t_m)$ and $t_0(t_m + 1)$ is—before rounding—greater than 1, as it is amplified by a factor $\log(D/R)/\log(G/R)$. Thus, signals ν_0 are further apart from each other than signals ν_m and, thus take longer than N periods to arrive.

For every type, we can write the selling time as ν_0 plus the bubble riding time τ or ν_m minus the crash preemption time ς. That is,

$$\nu_0 + \tau = \nu_m - \varsigma. \tag{22}$$

With a shrinking difference between t_m and t_0, τ and ς cannot both be type-independent, i.e., symmetric across types. Let us consider two possibilities. One is a constant, i.e., type-independent bubble riding time τ and therefore a crash preemption time ς that is decreasing in ν_m. The second, more fruitful specification, is centered around ν_m with a constant ς, which implies a bubble riding time τ that is a decreasing function of the signal ν_0, in a way similar—but not identical—to Doblas-Madrid (2016).

Let us first examine the possibility of symmetric bubble riding times, i.e., a strategy that focuses on selling at time $\nu_0 + \tau^*$ where $\nu_0 = t_0(\nu_m)$ for all $\nu_m < \overline{t_m}$. This kind of strategy is very limited in its ability to generate bubbles. Consider, in fact, the highest relevant type $\nu_m = \overline{t_m} - 1$, for whom the ν_0 signal is $t_0(\overline{t_m} - 1)$. This type cannot possibly ride the bubble longer than $\overline{t_m} - 1 - t_0(\overline{t_m} - 1)$ periods, since otherwise the maturity date will burst the bubble for exogenous reasons when the dividend is paid. With later types imposing a short limit on τ^*, symmetric bubble riding time is a poor way to characterize the economy's ability to generate bubbles. Simply put, early types have 'more room' to ride bubbles.

Strategies based on symmetric crash preemption allow us to characterize the possibility of bubble formation much more effectively. Suppose that agents of type $\nu_m \in \{\underline{t_m}, \underline{t_m} + 1, \ldots, \overline{t_m} - 1\}$ plan to sell at time $\nu_m - \varsigma^*$. Since the distance between ν_0 and ν_m is greatest for the lowest type, a type-independent ς^* implies longer bubble riding times for early types. Of particular interest is the case where $\varsigma^* = 0$. The longest bubble duration is then determined either by the distance between $\underline{t_m}$ and $t_0(\underline{t_m})$, or the maximum crash loss bearable by type $\underline{t_m} + 1$, the lowest type that does not know if they are first or second in line.

EXAMPLE: Suppose that $\gamma = 1/\theta - 1$, $R = 1$, $D = G^2$ and $\underline{t_m} = 100$. By (21), this implies $\overline{t_m} = 200$. Consider a symmetric crash preemption time $\varsigma^* = 0$, so that agents simply sell at ν_m. Table 1 shows, for every type ν_m in the relevant range, the implied overvaluation signals and bubble riding times if $\varsigma^* = 0$.

For this strategy profile to be an equilibrium, we need to verify that agents are willing to (i) sell at ν_m and (ii) buy before ν_m. Part (i) is straightforward. Given (14) and the fact that

ν_m	$\overbrace{100}^{=\underline{t_m}}$	101	102	...	196	197	198	$\overbrace{199}^{=\overline{t_m}-1}$
ν_0	0	2	4	...	192	194	196	198
$\nu_m - \varsigma^*$	100	101	102		196	197	198	199
$\tau^*(\nu_m)$	100	99	98	...	4	3	2	1

Table 1: Equilibrium with symmetric crash preemption time $\varsigma^* = 0$. This equilibrium is sustainable if type $\nu_m = \underline{t_m} + 1$ is willing to keep buying at $\nu_m - 1 = \underline{t_m}$ when these agents face a probability $1/(1+\lambda)$ of being 'second in line' and getting caught in the crash if not selling preemptively. The bubble riding time is decreasing in ν_m, longest for the lowest type.

$\tau^*(\nu_m) > 0$ for every type, all types are selling at a price higher than the dividend that will be paid later in the period. Part (ii) holds if $\tau^*(\underline{t_m}+1) \leq \tau^*_{\max}$, as given by (18). This implies that the type most inclined to sell—recalling that type-$\underline{t_m}$ agents face no crash risk since they know they are the lowest type—at the time when preemptive sales are most tempting, still chooses to buy. If $\tau^*(\underline{t_m}+1) > \tau^*_{\max}$, the equilibrium as stated can by modified by defining $\tau^*(\nu_m) = \min\{\nu_m - t_0(\nu_m), \tau^*_{\max\}}$ and thus $\varsigma^*(\nu_m) = \nu_m - t_0(\nu_m) - \tau^*(\nu_m)$. Another way of constructing a (partially) symmetric strategy would be to specify a ς^* large enough to be optimal for type $\nu_m = \underline{t_m} + 1$ and let all types sell at $\max\{\nu_0, \nu_m - \varsigma^*\}$. If $\varsigma^* > 0$, this specification does not generally maximize bubble duration. We summarize conditions for maximal bubble duration below.

Proposition 2. *Suppose that $D > G$. Then, the longest bubbles that can be sustained in equilibrium are such that type-ν_m agents, where $\nu_m \in \{\underline{t_m}, \underline{t_m}+1, \ldots, \overline{t_m}-1\}$ sell at time* $\min\{\nu_0(\nu_m) + \tau^*_{\max}, \nu_m\}$.

Proof. When one type is selling, all individuals of that type find it optimal to sell, since the price will reflect the sale, and since given (14), the current price is guaranteed to exceed the post-crash price. To verify that agents are willing to wait, we must verify $(\text{sell}_o r_w ait_1 d) fortype\nu_m = \underline{t_m} + 1$. If it holds, bubbles last until the maturity date. If not, type $\underline{t_m}+1$ sells at $\nu_0(\underline{t_m}+1) + \tau^*_{\max}$. Overall, each type ν_m rides the bubble for $\min\{\nu_m - \nu_0(\nu_m), \tau^*_{\max\}}$ periods. Types between $\overline{t_m}$ and $\overline{t_m} + N - 1$ face some crash risk, since it is possible that they are late in the line of investors. However, since earlier types $\nu_m < \overline{t_m}$ are willing to keep buying until ν_m despite higher crash probabilities, it follows that those later types are also willing to follow the equilibrium. □

The equilibrium described in this section resembles that in Doblas-Madrid (2016) in that the model is effectively finite, and bubble riding times are shorter for later types. In that model, however, agents to not have different signals about maturity dates, and therefore, several types must sell simultaneously, a fact that ends up having a key influence on bubble duration. By contrast, here, the existence of different beliefs about maturity dates, even in an effectively finite environment, allows for a more symmetric crash preemption equilibrium.

5 Infinite Bubbles: When Endowments Outgrow Dividends

We now consider the scenario where $D < G$. In this case, the expression inside the $\lceil\cdot\rceil$ bracket in (10) is linear in t_m with slope $\log(D/R)/\log(G/R)$ less than 1. As a result, the distance between t_m and $t_0(t_m)$ increases as t_m increases, as depicted in Figure 3. The implication is that is that signals ν_0 are more 'compressed' than signals ν_m. That is, for any given strategy profile with a constant crash preemption time ς^*, bubble duration $\tau^* = t_m - t_0(t_m) - \varsigma^*$ is strictly increasing in t_m. We thus proceed by considering equilibria with a constant crash preemption time ς^*.

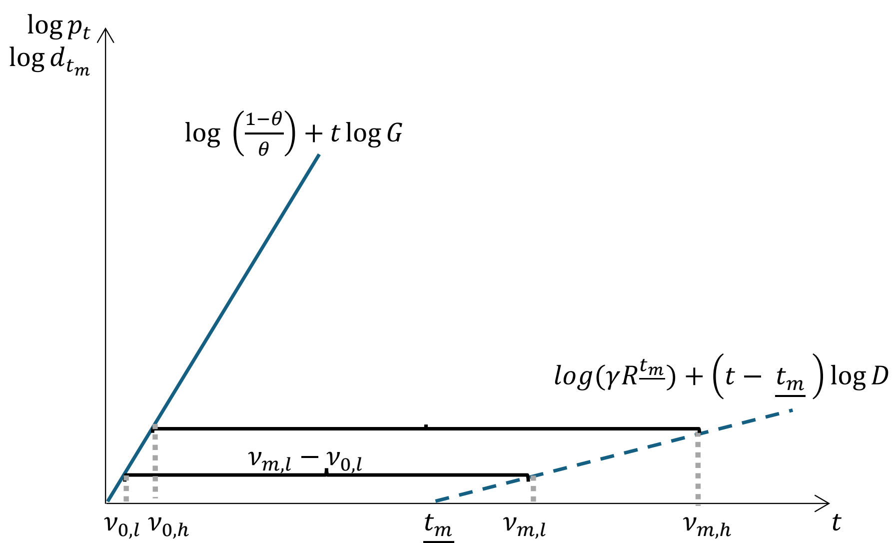

Figure 3: When $G > D$, signals about ν_0 are compressed, relative to signals ν_m. Thus, an equilibrium with a symmetric ς^* implies an equilibrium with a growing τ^*. In the limit as $D = 0$, there is no uncertainty about the overvaluation starting date, which makes it impossible to anchor strategies on ν_0, although anchoring strategies on ν_m remains possible.

The candidate equilibrium is the same as in (15), with the difference that $\nu_m - \varsigma^*$ replaces $\nu_0 + \tau^*$. The proof that there are no incentives to keep buying risky shares in period $\nu_m - \varsigma^*$ is straightforward, as each agent anticipates that all other agents of the same type are selling during the period so the price will revert to the fundamental price in period $\nu_m - \varsigma^* + 1$. Verifying the optimality of buying up to $\nu_m - \varsigma^*$ is more involved. However, it essentially follows the case where $D = G$, if we replace $\nu_0 + \tau^*$ with $\nu_m - \varsigma^*$, and we take into account that $D < G$. Formally, fix a realization $t_m \in \{\underline{t_m}, \underline{t_m} + 1, ...\}$, and consider an agent of type $\nu_m \in \{t_m, t_m + 1, ..., t_m + N - 1\}$. As before, it is sufficient to check the incentives to deviate and sell one period before the agent is called upon to sell. The expected price of a preemptive sale, $\mathbf{E}\left[p_{\nu_m - \varsigma^* - 1} | t_m \in \{\nu_m - 1, \nu_m\}\right]$, is

$$\left(\frac{1}{1+\lambda}\frac{(1-\theta)(1-1/N)}{\theta + (1-\theta)/N} + \frac{\lambda}{1+\lambda}\frac{1-\theta}{\theta}\right) G^{\nu_m - \varsigma^* - 1},$$

while the expected price of waiting to sell in $\nu_m - \varsigma^*$, $\mathbf{E}\left[p_{\nu_m - \varsigma^*} | t_m \in \{\nu_m - 1, \nu_m\}\right]$, is

$$\frac{1}{1+\lambda}\gamma\left(\frac{D}{R}\right)^{\nu_m - 1 - \underline{t_m}} R^{\nu_m - \varsigma^*} + \frac{\lambda}{1+\lambda}\frac{(1-\theta)(1-1/N)}{\theta + (1-\theta)/N} G^{\nu_m - \varsigma^*}.$$

Waiting is optimal for all $\nu_m \in \{t_m, t_m + 1, ..., t_m + N - 1\}$ if and only if

$$\gamma\left(\frac{D}{R}\right)^{\varsigma^* - \underline{t_m}} \left(\frac{D}{G}\right)^{\nu_m - 1 - \varsigma^*} + \frac{\lambda G}{R}\frac{(1-\theta)(1-1/N)}{\theta + (1-\theta)/N} \geq \frac{(1-\theta)(1-1/N)}{\theta + (1-\theta)/N} + \frac{\lambda(1-\theta)}{\theta}.$$

Observe that, as $t_m \to \infty$, $D < G$ implies that the first term in the left-hand side converges to zero. As a result, waiting is optimal irrespective of the realization of $t_m \in \{\underline{t_m}, \underline{t_m} + 1, ...\}$ if and only if

$$\frac{G}{R} \geq \frac{\frac{N-1}{1-\theta+N\theta} + \lambda\frac{1}{\theta}}{\lambda\frac{N-1}{1-\theta+N\theta}}.$$

If this inequality is satisfied, there exists a bubble equilibrium with a constant crash preemption time ς^* for all $\varsigma^* \in \left\{0, ..., \underline{t_m}\right\}$.

5.1 Crypto case $D = 0$.

Observe that, if $D \leq R$, the equilibrium exists despite the fact that it is common knowledge that the risky asset is valued above its fundamental value. In particular, the bubble equilibrium exists if $D = 0$ and the asset pays no dividends. This result is not aligned with AB, a model developed against the backdrop of the dotcom bubble, where firms investing in new technologies traded on the promise of future dividends. However, in recent years, booms and busts have occurred in markets for assets, such as *junk* cryptocurrencies and NFTs, which pay no dividends, are not used as media of exchange and, arguably, never will. This is the case of numerous cryptocurrencies, so-called *meme coins*, which appear to have no durable value. After being launched, their prices rise very quickly, only to fall to zero and disappear from the market. Thus, it is interesting, from a theoretical perspective, to examine the case of $D = 0$ in more detail. First, observe that finite, no-common-knowledge bubbles cannot arise without dividends. To see why, consider the case where $\gamma = 0$ and note that in this case it is common knowledge that f_t and t_0 are zero regardless of t_m. Thus, equilibria with symmetric τ^* cannot be sustained, since the precise bursting period would be known to all, activating backward induction. In the absence of dividends, it is thus necessary to use the maturity date t_m not as a payout date, but instead as a public signal or sunspot that serves as a coordination device. In fact, as shown above, uncertainty about the payout date allows us to build bubbles, as long as we define strategies not in terms of bubble riding time but in terms of crash preemption time. Although the preemption time ς^* is symmetric between types in this class of equilibria, the bubble riding time would vary between types without upper bound. Agents buy into the bubble hoping to sell to others who expect it to burst later than they do, with those buyers having the same reason for buying. This logic can (and is) iterated indefinitely, since t_m has no upper bound. Thus, bubbles on zero dividend assets are possible as long as there is no bound on bubble duration.

However, unlike junk cryptocurrencies, it can be argued that major cryptocurrencies, such as Bitcoin, may eventually be commonly used as digital payments, or simply be in demand as a store of value. In this sense, they behave similarly to an asset that pays a dividend in the future. For instance, Suppose that at t_m the (one and only) dividend is paid but agents continue to buy the asset at an arbitrary price growing at R in perpetuity simply because they still consider it a store of value. We could in this case redefine fundamental value as coming from dividends plus this residual perpetual store-of-value price.

6 Conclusion

In this paper, we analyze previously unexplored micro foundations in a model of speculative bubbles of the style of Abreu and Brunnermeier (2003; AB) and Doblas-Madrid (2012; DM), with specific focus on fundamental value and post-crash prices. Rather than treating post-crash values as an exogenous continuation payoff linked to the price process, we explicitly define dividends and asset maturity dates in the post-crash phase.

Disentangling prices from dividends yields a series of insights. Once dividends are considered independently from endowments, different regions of the parameter space create dynamics akin to different strands of the bubble literature, with the parametrization implied in AB emerging as a knife-edge case between the two. When dividends grow faster than endowments, speculative bubbles are finite and in the area of so-called greater fool bubbles based on lack of common knowledge and gains from trade. When dividends grow more slowly than endowments, the dynamics converge towards an infinite bubble model in which the role played by asymmetric information about dividends vanishes.

Our model also sheds new light on assets with zero dividends, such as cryptocurrencies. Standard greater-fool models fail in this environment because common knowledge of zero dividends eliminates speculation. We show that this failure arises from missing assumptions about the post-crash phase. If we reinterpret the maturity date, not as a dividend payment date but simply as a coordinating device that agents have asymmetric information about, infinite speculative bubbles are sustainable. This highlights the importance of rational—as opposed to purely behavioral—theories: bubbles persist not because agents repeatedly err, but because equilibrium trading remains possible under asymmetric information. More broadly, opening the black box of post-crash prices deepens our understanding of both traditional dividend-paying assets and modern zero-dividend assets, while sharpening the policy relevance of speculative dynamics.

Our focus on rational models of bubbles is motivated by the observation that, while behavioral factors may play a role in some episodes, rational theories make a stronger case for bubbles, allowing them to recur even if agents learn from mistakes, as opposed to 'tripping over the same stone' again and again. This idea is in line with experimental results by Asako et al. (2018). In their bubble game, when agents face symmetric information, and theory predict no bubbles, they observe that bubbles vanish after a few rounds of play. But with asymmetric information, as theory predicts, bubbles persist.

References

[1] Abreu, D. and M. Brunnermeier, 2003. Bubbles and crashes, Econometrica 71, 173–204.

[2] Allen, F., S. Morris and A. Postlewaite, 1993. Finite bubbles with short sales constraints and asymmetric information, Journal of Economic Theory 61, 209–229.

[3] Araujo, L. and A. Doblas-Madrid, 2026. Rational expectations fools' bubbles. Working paper, Michigan State University.

[4] Asako, Y. and K. Ueda, 2014. The boy who cried bubble, Economic Inquiry 52, 1137–1152.

[5] Awaya, Y., K. Iwasaki and M. Watanabe, 2022. Rational bubbles and middlemen, Theoretical Economics 17, 1559–1587.

[6] Barlevy, G., 2015. Bubbles and fools, Economic Perspectives 2Q, 54–76.

[7] Conlon, J., 2004. Simple finite horizon bubbles robust to higher order knowledge, Econometrica 72, 927–936.

[8] Conlon, J., 2015. Should central banks burst bubbles? Some microeconomics issues, Economic Journal 125, 141–161.

[9] Doblas-Madrid, A., 2012. A robust model of bubbles with multidimensional uncertainty, Econometrica 80, 1245–1293.

[10] Doblas-Madrid, A., 2016. A finite model of riding bubbles, Journal of Mathematical Economics 65, 154–162.

[11] Hirano, T. and N. Yanagawa, 2017. Asset bubbles, endogenous growth, and financial frictions, Review of Economic Studies 84, 406–443.

[12] Kindleberger, C. P., 1978. Manias, panics, and crashes: A history of financial crises, Basic Books.

[13] Liu, F., 2026. A simple higher-order rational Email-game bubble model and implications for momentum. Working paper.

[14] Liu, F. and J. Conlon, 2018. The simplest rational greater-fool bubble model, Journal of Economic Theory 175, 38–57.

[15] Liu, F., J. White and J. Conlon, 2023. A three-state rational greater-fool bubble model with intertemporal consumption smoothing, International Economic Review 64, 1565–1594.

[16] Martin, A. and J. Ventura, 2012. Economic growth with bubbles, American Economic Review 102(6), 3033–3058.

[17] Miao, J. and P. Wang, 2018. Asset bubbles and credit constraints, American Economic Review 108, 2590–2628.

[18] Minsky, H. P., 1986. Stabilizing an Unstable Economy, Yale University Press.

[19] Samuelson, P., 1958. An exact consumption-loan model of interest with or without the social contrivance of money, Journal of Political Economy 66(6), 467–482.

[20] Santos, M. and M. Woodford, 1997. Rational asset pricing bubbles, Econometrica 65, 19–57.

[21] Tirole, J., 1982. On the possibility of speculation under rational expectations, Econometrica 50, 1163–1181.

[22] Tirole, J., 1985. Asset bubbles and overlapping generations, Econometrica 53, 1499–1528.