RATIONAL FINITE BUBBLES

by

Franklin Allen
Gary Gorton

(41-88)

RODNEY L. WHITE CENTER FOR FINANCIAL RESEARCH
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367

RATIONAL FINITE BUBBLES*

Franklin Allen and Gary Gorton

The Wharton School
University of Pennsylvania
Philadelphia, PA 19104

Latest Revision: November 1988

Abstract

There has been a long-running debate about whether stock market prices are determined by fundamentals. To date no consensus has been reached. An important issue in this debate concerns the circumstances in which deviations from fundamentals are consistent with rational behavior. A continuous-time example where there are a finite number of rational traders with finite wealth is presented. It is shown that a finitely-lived security can trade above its fundamental.

---

*We are grateful to Andy Abel, Colin Camerer, Joe Haubrich, Alan Kraus, Max Maksimovic, Steve O'Connell, Gordon Sick, Mark Weinstein, Randy Wright, Josef Zechner, Steve Zeldes and participants at seminars at Cornell University, University of British Columbia, University of North Carolina-Duke, the University of Pittsburgh, University of Pennsylvania, University of Rochester, and the 1988 Johnson Symposium at the University of Wisconsin for many helpful suggestions. The comments of our discussant at this Symposium, Chester Spatt, were particularly helpful. This research began as joint work with Peter Knez to whom we remain indebted. The authors alone are responsible for errors and omissions. Financial support was provided to the first author by NSF grant no. SES 8813719 and to the second author by NSF grant no. SES 8618130.

1. Introduction

There has been a long and continuing debate on the determinants of stock prices. One view is that these prices reflect economic fundamentals; that is, a firm's stock price equals the present discounted value of its dividends. Another view is that stock prices are "bubbles" and consistently deviate from their fundamentals. As an empirical matter, there is currently no consensus on which of these views is correct.

Historically, the possibility that stock prices are "bubbles" was raised by a number of extreme incidents. Perhaps the most well-known of these is the South Sea Bubble. During the first six months of 1720 the stock price of the British South Sea Company rose by 642 percent; during the last six months of 1720 the price fell back to its original value. A similar rise and sudden decline occurred in the stock price of John Law's Mississippi Land Company in France. Both episodes were reminiscent of the Dutch Tulip Mania in the previous century, and were precursors of the stock market crashes of subsequent centuries of which October 1929 and October 1987 are perhaps the most famous. While it is by no means clear that these events constitute evidence of "bubbles", they were important in that many of them lead to regulation. For example, the South Sea Bubble caused the British Parliament to pass the South Sea Act which effectively eliminated the stock market as a source of funds for over a century. In the United States, the Great Crash of 1929 lead to the creation of the SEC and the introduction of numerous regulations, many of which are still in force.

More recently, the results of Shiller (1981), Grossman and Shiller (1981), and Leroy and Porter (1981), among others, suggest that stock prices deviate from market fundamentals. There is again no wide agreement on the validity of these studies; those who have challenged the methodology adopted include Flavin (1983), Kleidon (1986a,b) and Marsh and Merton (1986). West (1988) provides a more complete survey of this and related controversies. Price paths that deviate from fundamentals have also been observed in experimental settings (see Smith, Suchanek and Williams (1988)).

In addition to the empirical debate about the determinants of stock prices, a growing theoretical literature has begun to address the question of how asset prices can deviate from market fundamentals. Camerer (1987) gives a full survey of this literature. In infinite horizon models, rational bubbles have appeared as explanations for the existence of fiat money starting with Samuelson (1958). Important contributions were subsequently made by Wallace (1980), Flood and Garber (1980), Blanchard (1979), Blanchard and Watson (1982) and Tirole (1985), among others. Although these theories can explain a number of features of "bubbles" they are not entirely satisfactory explanations of the phenomena the empirical literature has been concerned with. Some of these models require that prices grow slower than the expected growth rate of the aggregate wealth of the economy. There is no explanation of how bubbles get started or of why they crash. Starting and stopping are taken as exogenous. Diba and Grossman (1988) have argued there is no possibility that price bubbles can crash and restart. Also these theories cannot address the question of whether finitely-lived security prices can deviate from fundamentals.

The major result for finite-horizon models is a negative one. Tirole (1982) argues that in a discrete-time finite-horizon setting stock prices cannot deviate from fundamentals unless traders are irrational or myopic. He makes three important points in ruling out finite bubbles. First, he points out that with a finite horizon the bubble would never get started because it would "unravel." To see this let the final date in the economy be T. Then at the date T - 1 an agent would not buy the asset at a price above the discounted value of its payoff at T because he would incur a loss if he did so. Therefore, the bubble cannot exist at T - 1. Similarly, by backward induction it follows that a bubble cannot exist at any point in time. Secondly, with a finite horizon traders cannot be induced to hold the stock by a price path that goes to infinity because there is finite wealth. Consequently, there must be a date at which the (real) price path necessary to support the bubble exceeds the total available wealth in the economy. At that date the bubble will crash, but then at the date before that no other trader will buy the asset. Again by backward induction the bubble cannot get started. Finally, without insurance motives for trading not all of the finite number of traders can rationally expect to benefit since they know that the bubble is a zero-sum game. If traders are risk averse, some must be strictly worse off since they bear risk and not everybody can have a positive expected return.

Tirole's (1982) results exemplify the difficulties of constructing theories which are based on conventional assumptions and which are consistent with bubbles. These difficulties have lead some authors to abandon the traditional neoclassical assumption of rational behavior. One example is Shiller (1984) who models stock prices as being subject to "fads". Another is DeLong, Shleifer, Summers, and Waldman (1987) who assume that some traders continue to hold beliefs even after it becomes clear these are rejected by the data. These irrational traders are consistently overly optimistic (or overly pessimistic) and take larger positions than they would do if they were rational. This means they bear more risk than is optimal but their wealth is not driven to zero. They therefore persistently cause stock prices to deviate from their fundamental. (See Camerer (1987) for other examples.)

The model presented below takes a different approach. We assume all agents are rational but they populate an imperfect world. In particular, there is an agency problem arising from an information asymmetry. In the corporate finance literature the analysis of agency relationships is commonplace and their implications for firms' investment decisions are well known. It is widely accepted that asymmetric information can lead to firms making inefficient investment decisions. Despite the fact that in the United States and many other countries, a majority of the wealth held in stocks is invested indirectly through financial intermediaries, the implications of agency relationships for asset pricing has not been fully investigated. It is argued below that one of the manifestations of asymmetric information in this context is that asset prices can deviate from their fundamental values and be subject to bubbles.

We assume there are two types of people that can obtain the qualifications necessary to become a portfolio manager. The first group can each identify a certain number of undervalued firms. The supply of these firms relative to the total number that can be identified by this group is such that their prices are not bid up. Thus markets are not strong-form efficient but this is not inconsistent with rationality. The second group is unable to identify undervalued firms. Lenders cannot observe which type of portfolio manager they are entrusting their wealth to.

In Section 2 we assume that the portfolio managers, who have no wealth of their own, receive a proportion of the profits that they make so their payoff has the form of a call option; this is later shown to be an optimal contract. We focus on the decisions of the second group that cannot identify undervalued securities. It is shown that these traders are willing to speculate in the sense of Harrison and Kreps (1978), that is, they are "willing to pay more for [the security] than they would pay if obliged to hold it [to the horizon]" This is because of the fact that there is an asymmetry in their incentives. If they lose the money entrusted to them they obtain nothing no matter how badly they do. However, if they do well they keep a proportion of what they make. They are therefore prepared to purchase securities which are trading above their fundamental provided there is some chance of a capital gain even though they know that there is a good chance they will lose their investors' money when the bubble crashes.

The key issue is therefore whether or not traders perceive there to be some chance of a capital gain at all points in time. This depends on what they conjecture about the strategies of other traders. We consider a very stylized structure which makes traders' conjectures about other traders' actions very simple. In particular, we assume at the outset that traders leave the market when they "die" and that their "death" times are correlated in a particular way. This rationale for exiting from the market, and the correlation structure of these "death times," are clearly not meant to be taken literally but rather are devices for streamlining the model in order to focus on the theoretical issue of the existence of bubbles in a finite world. The main point is to develop a simple structure under which the logical sequence of conjectures traders go through will not lead to unraveling but to traders deciding rationally to speculate. Having developed this structure we go on to show how the model can be extended to the case where a trader's exit from the market arises from an endogenous decision rather than being due to an exogenous event.

In Section 3 we consider both groups of portfolio managers and demonstrate that the contract assumed in Section 2 is an equilibrium contract. It is not worthwhile for lenders to knowingly entrust their wealth to portfolio managers that speculate. However, they cannot tell them apart from the good portfolio managers that can identify undervalued firms. Therefore in equilibrium the good managers subsidize the ones that speculate and lenders earn their opportunity cost.

Tirole (1982) argued that bubbles could not occur in standard finite horizon models unless traders were myopic or irrational in some other way. The example presented in Sections 2 and 3 shows that bubbles can occur in finite horizon models when traders are rational. The reason for this difference in results is that our example is in a different class from those considered by Tirole. An important issue concerns the robustness of our example and its implications for theories of asset pricing. Section 4 discusses the critical elements of the example that lead to bubbles.

2. Speculative bubbles

This section considers a stylized model of a stock market in which there are three traders. We assume that these traders have no wealth of their own but instead manage other people's wealth for them. They receive a proportion of any profits they make; this contract is shown to be optimal in Section 3 below. Subsection (i) outlines the basic model. Subsection (ii) considers how bubbles arise when traders leave the market because of some exogenous stochastic event which we term "death". The timing of a trader's own death is unknown to that trader until the instant before it occurs. In Subsection (iii) we consider an example where death times are known from the beginning. Finally, Subsection (iv) shows how this can be extended to the case where exit from the market does not arise as a result of "death" but from an exogenous decision of traders. They decide to leave when the expected gain from holding the share is no longer sufficient to compensate them for the risk of continuing to hold it.

(i) The basic model

The following assumptions detail the basic model.

(A1) There are three traders called Persons 1, 2 and 3.

(A2) The model lasts for one continuous period, beginning at $t = 0$ and ending at $t = 1$. Trades can occur at any time between 0 and 1.

(A3) The traders consume just before they die which occurs somewhere between 0 and 1.

(A4) The agents' utility is an increasing function of consumption. They can be either risk neutral or risk averse.

(A5) Person 1 dies at date t_1 which is drawn from a uniform distribution on $[0, 1)$. Person 2 dies at t_2 where

$$t_2 = t_1 + \frac{1}{3}(1 - t_1). \quad (1)$$

Person 3 dies at t_3 where

$$t_3 = t_1 + \frac{2}{3}(1 - t_1). \quad (2)$$

(A6) Agents learn their death times just in time to allow them to trade and consume before they die. Death is private information.

(A7) There exists a firm with a known and certain payoff which for simplicity we normalize to zero. In other words the fundamental is zero. The firm issues one indivisible share. This share cannot be short sold.

(A8) Person 1 is always endowed with the share. Person 1 knows his identity.

(A9) Persons 2 and 3 are not endowed with any shares. They do not know their identities (i.e., whether they are going to die last) and assign equal weight to each of the two possibilities.

(A10) Persons 2 and 3 have no wealth of their own. However, they are able to invest other people's wealth. They are to be thought of as portfolio managers. They have a fixed amount B (= 1 in illustrations) they invest. The amount π they repay to investors if the amount they have at the end is y, is:

$$\pi = B + \alpha(y - B) \qquad \text{for } y \geq B \tag{3a}$$

$$= y \qquad \text{for } y < B \tag{3b}$$

where $0 \leq \alpha \leq 1$. (In illustrations it is assumed that $\alpha = 0.95$.) In effect, the payoff the portfolio managers receive is a call option.

The accounting system is such that they cannot simply consume the money they borrow. They can only consume the fee that they are paid for managing the portfolio. It follows from (3) that this is a proportion $1 - \alpha$ of the profits that they make if these are positive and nothing if they are negative.

(A11) The identity of the owner of the share is private information throughout.

(A12) Trade occurs in the following way. All traders have the same expectations about prices $p(t)^e$ at which trades will occur at time t. When a person decides to sell the share he seeks out a buyer. He locates one or the other of the traders that remain in the market with equal probability. If he finds another trader then trade occurs at $p(t)^e$. If the seller cannot find a buyer this becomes public information and the price of the share falls to zero. In illustrations it is assumed that

$$p(t)^e = t \qquad \text{for } t \in [0, 1). \tag{4}$$

(A13) When a trade occurs only parties to the trade observe the transaction.

(A14) All agents know the structure of the model and the distributions of the random variables but not particular realisations of random variables they do not observe.

(ii) Unknown death times

The share considered has a fundamental of zero. It is clear that an equilibrium with $p(t)^e = 0$ exists where the share price reflects this fundamental. The question that we address is the following. Do there exist other price paths such that a rational agent is prepared to buy the share at a strictly positive price even though he knows the final payoff to the share is zero? Our first result is:

Proposition 1

When death times are unknown there exists a set of self-fulfilling beliefs such that two trades will always occur at a strictly positive price between date 0 and date 1 provided:

(i) $0 < p(t)^e < B$ for all $t \in [0, 1)$;

(ii) $p'(t)^e > 0$ for all $t \in [0, 1)$.

To see why this holds first consider a numerical example where $t_1 = 0.1$ so that from (1) and (2) $t_2 = 0.4$ and $t_3 = 0.7$. As mentioned in the previous subsection we also assume $p(t)^e = t$, $B = 1$ and $\alpha = 0.95$. For ease of exposition we first describe a possible sequence of events without analyzing the traders' decisions. This sequence of events is illustrated in Figure 1. We then consider a set of beliefs and show that these support the decisions in Figure 1. Finally, we generalize the example and show that the beliefs are self-fulfilling.

At $t = 0$ Person 1 is endowed with the share. At $t = 0.1$ he finds out he's going to die and searches for a buyer which is Person 2 or 3 with probability 0.5. For concreteness we assume he finds Person 2 who buys the share at a price of 0.1. At $t = 0.4$, just before he dies, Person 2 searches for a buyer and finds Person 3 who buys the share at a price of 0.4. He makes a profit of $0.4 - 0.1 = 0.3$ and after repaying his investors consumes $(0.05)(0.3) = 0.015$. At $t = 0.6$ Person 3 searches for a buyer, but finds none. The bubble bursts and the price of the stock falls from 0.6 to zero. Finally, at $t = 0.7$ Person 3 dies. At the time of his death, he has 0.6 remaining and so is only able to return this amount to his investors. He consumes nothing.

Consider the following set of beliefs. Given that $p'(t)^e > 0$, all agents believe that:

(a) if there is a prospective buyer alive he will be prepared to buy the share when approached.

Agents who do not know their own identity believe that:

(b) any agent offering to sell at a date in the interval $0 \leq t < 1/3$ is Person 1 selling at t_1 with probability 1; and

(c) any agent offering to sell at a date in the interval $1/3 \leq t < 1$ is Person 1 selling at t_1 with probability 0.4 or Person 2 or 3 selling at t_2 with probability 0.6.

We demonstrate below that these beliefs support the sequence of actions in Figure 1 as an equilibrium and that they are self-fulfilling.

First consider Person 1's decision. He knows from the structure of the model that for $0 \leq t \leq t_1$ Persons 2 and 3 will be alive. Thus from (a), Person 1 believes he can sell the share at any time until his death. Since the share price is increasing through time, it is optimal for Person 1 to hold the share until he has to sell it at his death time t_1. Thus at $t = 0.1$, Person 1 will search for a buyer. There is a 0.5 probability he will find Person 2 and a 0.5 probability he will find Person 3. For concreteness we suppose that Person 2 is found.

Next consider Person 2's decision. From (b), he believes that the seller is Person 1. He can put himself in the place of Person 1 and by doing so deduce that Person 1's optimal strategy is to sell at his death time. He therefore knows that $t_1 = 0.1$ which implies that $t_2 = 0.4$. This means that he should not wait past $t = 0.4$ to sell the share since if he survives that date he will be the sole remaining trader; until that date there will definitely be another buyer. Since the price is increasing, he should sell at $t = 0.4$.

Person 2 finds the remaining trader, Person 3, at $t = 0.4$. Consider Person 3's decision. Since he was not endowed with the share he knows he is Person 2 or 3, but does not know which. Since he is approached at $t = 0.4$ he does not know whether the seller is Person 1, 2, or 3. It follows from (c) that he believes there are two possibilities. There is a 0.4 probability that the seller is Person 1 in which case $t_1 = 0.4$. This implies that $t_2 = 0.6$ in which case from (c) the share could be sold at any date up to this point. We refer to this first possibility as state S to indicate the share can be sold again. He also believes there is a 0.6 probability that the seller is Person 2 or 3. In this case Person 1 must have sold it at $t_1 = 0.1$ and there will be no one for the trader to resell it to. We refer to this second possibility as state N to indicate that no resale is possible.

The payment schedule in (3) implies Person 3 cannot lose from buying the share and he can gain if he manages to resell it at a higher price. Since he attaches a 0.4 probability to there being another trader who he can resell the share to at a higher price, he is strictly better off purchasing the share.

What is the optimal time for him to try to sell the share? If state N is the true state, then there is no other trader to sell the share to. This possibility therefore has no effect on his optimal selling time. If state S is the true state, then $t_1 = 0.4$ and $t_2 = 0.6$. Hence, since price is rising his optimal action is to search for a buyer at $t = 0.6$. In fact in this example there is no other buyer to be found, so at $t = 0.6$ he realises that he is Person 3 and the bubble crashes. At $t = 0.7$ Person 3 dies.

So far we have considered the case where $t_1 = 0.1$. It can be seen that for $0 \leq t_1 \leq 1/3$ the analysis is the same because only Person 1 can die in this interval. For $1/3 \leq t_1 \leq 1$ it can be seen that the beliefs (a) again make it optimal for Person 1 to sell at his death time. The difference here is that the identity of the seller in the first transaction will be unknown. The buyer's decision is then the same as Person 3's at $t = 0.4$ above; he cannot distinguish between states S and N and assigns probabilities of 0.4 and 0.6 to these, respectively. Any other transactions in the interval $1/3 \leq t \leq 1$ also have this feature so that the analysis of other possible cases is similar to that of the illustration.

Why are the beliefs (a), that when found a prospective buyer will always purchase the share, correct in equilibrium? First, consider somebody who is approached after $t = 1/3$. A prospective purchaser will be better off buying provided he believes that there is some probability that he can resell the share. This depends on whether there is some probability he can locate a prospective buyer and this conjectured buyer believes that he can resell the share, and so on. From the point of view of any new buyer there is always a 0.4 probability of another willing buyer later. This chance of state S is independent of time. At any point a prospective buyer cannot distinguish between the seller being Person 1 or the seller being Person 2 or 3 and hence whether or not another buyer remains. As a buyer goes through the logical sequence of conjectures concerning whether he will be able to resell, he knows for certain that the share cannot be resold more than once. However, the person that he might sell to will think there is a 0.4 probability he will be able to resell and so on; as far as each buyer in the sequence is concerned there is always a possibility that the share can be resold once. This is true for an infinite sequence of conjectured buyers. No matter how close to $t = 1$ a sale were to occur, (1) and (2) together with $p'(t)^e > 0$ imply that there is always a 0.4 possibility of reselling the share at a profit so unravelling does not occur. For anybody approached before $t = 1/3$ the analysis is similar except there is a probability of 1 they can locate another willing buyer. Thus beliefs (a) are correct in equilibrium.

Why are beliefs (b) and (c) correct in equilibrium? It was argued above that Person 1 always sells at his death time t_1 and the person he sells to always sells at t_2. The unconditional distribution of t_1 is uniform on $[0, 1)$ with density 1 and the unconditional distribution of t_2 implied by (1) is uniform on $[1/3, 1)$ with density 3/2 as shown in Figure 2. Hence, the beliefs (b) that for $0 \leq t_1 \leq 1/3$, anybody selling the stock is Person 1 with probability 1 are correct. For $1/3 \leq t_1 \leq 1$ the probability the seller is Person 1 (i.e. state S) is $1/(1 + 3/2) = 0.4$ and the probability the seller is Persons 2 or 3 (i.e. state N) is $(3/2)/(1 + 3/2) = 0.6$. Beliefs (c) are, therefore, also correct in equilibrium.

These arguments show that provided the expected price is always below B so that traders have enough resources to buy the share and provided the price path is rising, there will always be two trades at a strictly positive price. Thus Proposition 1 is demonstrated.

Why do our results differ from those of Tirole (1983)? His first argument is that in a discrete-time finite-horizon model a bubble would never get started because it would unravel. If an asset's payoff at date T is known to be zero then at date T - 1 nobody will buy it at a positive price. Similarly at date T - 2 and so on so that the asset is always worthless. In our model time is continuous so that although there is a final date t = 1 there is no date corresponding to T - 1; no matter how close to t = 1 it is always possible to resell the share before the final date. The unravelling argument is not applicable.

Tirole's second argument is that with a finite horizon the price path cannot go to infinity because there is finite wealth. If the price path did go to infinity the amount needed to purchase the share would exceed the total wealth available in the economy. Again by backward induction the bubble cannot get started. In our model the price path does not go to infinity. The reason that this is not necessary to support the equilibrium is the correlation structure of death times. No matter how close to t = 1 a trade occurs the probability of finding a subsequent buyer is 0.4. It is always optimal for the trader to hold the stock until the conjectured t_2. Without some correlation structure of this type the chance of finding a buyer would fall towards zero and the price path would need to rise to infinity to induce the trader not to sell.

His final argument for bubbles not existing is that without insurance motives for trading not all of the finite number of traders can expect to be better off ex ante since they know that the bubble is a zero-sum game. If they are risk averse some must be strictly worse off. In our model all the traders participating in the bubble are strictly better off ex ante. The reason is that they are investing with other people's money and their reward structure is such that they do not care about the magnitude of any losses they incur. The people who bear the losses ex-post are the investors lending them the money. They are willing to lend because the traders are pooled with portfolio managers that can identify profitable investment opportunities and they effectively subsidize these losses by paying a higher interest rate than they would have to in the absence of the bad portfolio managers. This aspect of the model is explained below in Section 3.

The purpose of most of the assumptions used is to simplify the analysis. For example, having Person 1 endowed with the share limits the number of cases that need to be considered. It would also be possible to have the share randomly endowed. In that case neither Persons 1, 2, or 3 knows their identity and the number of possible states of the world each agent must consider is significantly increased. However, the results do not change substantively in this case.

The focus of the analysis on bad portfolio managers whose payoff is effectively a call option is also to reduce the number of possible states that need to be considered. It would be possible to include ordinary traders who simply invest their own wealth. They will be prepared to participate in the bubble in its early stages since for $0 \leq t < 1/3$ they will be able to find a buyer.

The assumption that outlines the way in which trade occurs is an important one. Its role is essentially similar to that of the Walrasian auctioneer and price-taking in standard competitive models since it allows strategic aspects of traders' behavior to be ignored. Its purpose is again to simplify the nature of the conjectures that people make about what could have happened in the past.

The proposition indicates that any price path which is monotone increasing is an equilibrium. In addition the fundamental is of course also an equilibrium. This multiplicity of equilibria is similar to that which arises in infinite-horizon overlapping generations models. As in these cases one way of describing which equilibrium occurs is to associate each equilibrium with the outcome of an exogenous random event or "sunspot".

(iii) Bubbles with known death times

The analysis above has the feature that agents' decisions to leave the market are exogenous. As discussed in the introduction, "death" is not meant to be taken literally but rather is meant to represent any event that causes the trader to leave the market. For example, death could correspond to the timing of liquidity needs. We next develop examples where the decision to leave is endogenous. As a first step in this direction, we start by assuming that the model is the same as above except that each person knows his own death time from the start. We again show that bubbles can exist. In the next subsection we show how the model can be reinterpreted so that decisions are endogenous.

We replace assumptions (A4), (A5) and (A6) with:

(A4') The agents' utility is an increasing function of consumption. They are risk neutral.

(A5') Person 1 dies at date t_1 which is drawn from a uniform distribution on $[0, 1)$. Person 2 dies at t_2 where

$$t_2 = t_1 + \beta(1 - t_1), \tag{5}$$

where β is the unobservable realisation of a random variable distributed uniformly on $(0, 0.5)$. Person 3 dies at t_3 where

$$t_3 = t_1 + 2\beta(1 - t_1). \qquad (6)$$

(A6') Agents learn their death times at $t = 0$. Death times are private information.

We again consider whether a rational agent would strictly prefer to buy the share even though he knows its payoff is zero. We demonstrate the following:

Proposition 2

When death times are known, there exists a set of self-fulfilling beliefs such that at least one trade will always occur at a strictly positive price between date 0 and date 1 provided:

(i) $0 < p(t)^e < B$ for all $t \in [0, 1)$;

(ii) $p'(t)^e > 0$ for all $t \in [0, 1)$;

(iii) $p''(t)^e \geq 0$.

To see why this holds first consider a numerical example with $t_1 = 0.2$, $\beta = 0.25$ and the other parameters as before. From (5) and (6) these values imply that $t_2 = 0.4$ and $t_3 = 0.6$. As before, we first describe a possible sequence of events without analyzing the traders' decisions. This sequence of events is illustrated in Figure 3. We then consider the beliefs that support these decisions as equilibrium outcomes and show they are correct in equilibrium.

At $t = 0$ Person 1 is endowed with the share. At $t = 0.2$ he searches for a buyer which is Person 2 or 3 with probability 0.5. For concreteness we assume he finds Person 2 who buys the share at a price of 0.2. At $t = 0.2$

Person 2 looks for somebody to sell the share to and finds Person 3. Person 2 sells the share to Person 3 for 0.4, making a profit of 0.2. The seller is thus able to consume (0.05)(0.2) = 0.01 before he dies. At $t = 0.6$, Person 3 looks for somebody to sell to but discovers that he is the only remaining trader so that he must be Person 3. The bubble bursts and the stock price falls to zero. He is only able to repay 0.6 and consumes nothing.

Consider the following set of beliefs. Given that $p'(t)^e > 0$, all agents believe that:

(a') if there is a prospective buyer alive he will be prepared to buy the share when approached.

Agents who do not know their own identity believe that:

(b') any agent offering to sell at date t in the interval $0 \leq t \leq 0.5$ is Person 1 selling at t_1 with probability $1/[1 - \ln(1 - t)]$, Person 2 selling at t_2 with probability $- \ln(1 - t)/[1 - \ln(1 - t)]$ and Person 3 selling at t_3 with probability 0;

(c') any agent offering to sell at date t in the interval $0.5 < t < 1$ is Person 1 selling at t_1 with probability $1/(1 + \ln 2)$, Person 2 selling at t_2 with probability $\ln 2/(1 + \ln 2)$ and Person 3 selling at t_3 with probability 0.

We demonstrate below that these beliefs are self-fulfilling and support the sequence of actions in Figure 3 as an equilibrium.

First consider Person 1's decision. By the same logic as before, Person 1 will sell at $t_1 = 0.2$ just before he dies. There is a 0.5 probability he will sell to Person 2 and a 0.5 probability he will sell to Person 3. For concreteness we assume he finds Person 2.

The buyer knows his death time is $t = 0.4$ and that he could be Person 2 or 3 since he was not endowed with the share. From (b') he believes the seller is Person 1 with probability 1/(1 - ln 0.8) and Person 2 with probability - ln 0.8/(1 - ln 0.8). If the seller is Person 1 then it is equally likely the buyer is Person 2 or 3 given that Person 1 is equally likely to find Persons 2 and 3; if the seller is Person 2 then the buyer is Person 3. As far as the buyer is concerned there are thus three possibilities which are described in Table 1.

Table 1

State	Seller's Identity	Buyer's Identity	Probability
S1	Person 1	Person 2	0.5/(1 - ln 0.8)
S2	Person 1	Person 3	0.5/(1 - ln 0.8)
N	Person 2	Person 3	ln 0.8/(1 - ln 0.8)

For each of these states the buyer can use (5), (6), his own death time of 0.4 together with his belief that the seller is selling at the seller's death time to deduce that β, t_1, t_2 and t_3 have the values shown in Table 2.

Table 2

State	β	t_1	t_2	t_3
S1	0.25	0.2	0.4	0.6
S2	0.125	0.2	0.3	0.4
N	0.2	0	0.2	0.4

What is the buyer's optimal selling strategy given these beliefs? In state N he is unable to sell the share so his payoff is 0 no matter what he does. Thus his decision only depends on states S1 and S2. For any sale at t such that $0.2 < t \le 0.3$, the seller expects a payoff of $p(t)^e$ if the true state is either S1 or S2. For any sale at t such that $0.3 < t \leq 0.4$ the seller expects a payoff of $p(t)^e$ if state S1 holds and 0 if state S2 holds. Given $p'(t)^e > 0$ and the fact that states S1 and S2 are equally likely conditional on the seller being Person 1, it follows that it is optimal for the buyer to sell at his death time of t = 0.4 if

$$p(0.3)^e - p(0.2)^e \leq 0.5\ [p(0.4)^e - p(0.2)^e] \qquad (7)$$

and at t = 0.3 otherwise. (It is assumed that when indifferent the person sells at his death time). In the example, $p(t)^e = t$, so this is the boundary case where (7) is satisfied with equality. It can readily be seen given risk neutrality that the buyer will always sell at his death time provided $p''(t)^e \geq 0$. However, if $p''(t)^e < 0$ then the buyer will sell at the earlier time of 0.3. In this case it is possible to derive a result similar to Proposition 2 where the buyer believes that if the seller is not Person 1 he is selling at a time half way between his death time and the time at which he bought.

To summarize, Person 2's optimal strategy is to buy at t = 0.2 and try to sell at t = 0.4; he is strictly better off doing this than not buying since the probability he can sell at a positive profit at t = 0.2 is 0.5/(1 - ln 0.8). Hence at t = 0.4 Person 2 searches for a buyer and finds Person 3.

Now the buyer knows that his death time is t = 0.6, but does not know whether he is Person 2 or Person 3. His beliefs are as in Table 1 except t = 0.4 rather than 0.2 so 0.6 replaces 0.8. Similarly to Table 2 he can deduce that β, t_1, t_2 and t_3 have the values given in Table 3 in each of the three states.

Table 3

State	β	t_1	t_2	t_3
S1	0.333	0.4	0.6	0.8
S2	0.167	0.4	0.5	0.6
N	0.25	0.2	0.4	0.6

What is the buyer's optimal strategy if he buys? As before only states S1 and S2 are relevant for his decision. Comparing as in (7), it follows that his optimal strategy is to sell at his death time. At $t = 0.6$ he looks for somebody to sell to but discovers that he is the only remaining trader so that he must be Person 3. The bubble bursts, the stock price falls to zero and he dies.

So far we have considered the case where $t_1 = 0.2$, and $\beta = 0.25$. It can be seen that other parameter values will lead to a similar analysis. Whenever a prospective buyer is approached for the first time he cannot distinguish between states S1, S2 and N and makes a decision similar to Person 2's and Person 3's decisions above. Given risk neutrality and $p''(t)^e \geq 0$ it is always optimal for a person to hold the share until his death time.

Why are the beliefs (a') correct in equilibrium? The argument is the same as when death times are unknown; from the point of view of any new buyer approached for the first time there is always a chance of another willing buyer later and this is true for an infinite sequence of conjectures.

Why are beliefs (b') and (c') correct in equilibrium? Using the facts that t_1 is uniformly distributed on [0, 1), β is uniformly distributed on (0, 0.5) and (5), it can be shown that the unconditional density function of t_2 is:

$$f_2(t) = -2 \ln (1 - t) \qquad \text{for } 0 \le t \le 0.5; \tag{8a}$$

$$= -2 \ln 2 \qquad \text{for } 0.5 < t < 1. \tag{8b}$$

Figure 4 illustrates this together with the unconditional density function of t_1.

It can be seen that the probability, ρ, that a death occurring at a particular time is that of Person 1 relative to that of Person 2 is simply the ratio of the density functions $f_1(t)/f_2(t)$ at that point. Within the interval $0 \le t \le 0.5$:

$$\rho = \frac{1}{1 - 2 \ln (1 - t)} . \tag{9}$$

Given the assumption that $p''(t)^e \le 0$ and risk neutrality then, as argued above, holders of the share always sell at their death time. This implies that sales occur in situations S1, S2 and N. Unlike the case where death times are unknown, which was analyzed in the previous subsection, Person 3 never sells to Person 2; this is because if Person 3 is sold the share by Person 1 he holds it until his death time at t_3. Hence, in all the cases where a death occurs at t_1 a sale occurs, but in only half the cases where a death occurs at t_2 does a sale occur. Within the interval $0 \le t \le 0.5$ the probability that a sale is by Person 1 is:

$$\frac{\rho}{\rho + 0.5(1 - \rho)} = \frac{1}{1 - \ln (1 - t)} . \tag{10}$$

Since Person 3 never sells, it follows that the probability Person 2 is the seller is $-\ln (1 - t)/[1 - \ln (1 - t)]$. Hence the beliefs (b') are correct in equilibrium. The argument that the beliefs (c') are correct is similar except that $\ln 2$ replaces $-\ln (1 - t)$ in (9) and (10).

In the analysis above it was assumed that the agents were risk neutral.

If the traders are risk averse the only substantial difference is that the boundary case where sales by Persons 2 and 3 occur at death times rather than half-way between the purchase time and the death time is no longer $p''(t)^e \geq 0$ but requires $p''(t)^e$ to be larger.

(iv) Bubbles with standard utility functions

So far it has been assumed that traders are motivated to leave the market exogenously, namely, they die. Moreover, we assumed that these "death times" were correlated in a particular way. In this subsection we provide another version of the theory. The time at which the security will be redeemed is uncertain. There is also assumed to be a level of wealth for each trader beyond which the marginal utility of consumption is so low that the rate at which their wealth is increasing is insufficient to compensate them for the risk of the share being redeemed at its fundamental so they sell and leave the market.

We again illustrate the theory using an example with the standard parameter values. Investors' consume at the end of the period when they receive their compensation for being portfolio managers. We replace (A4) with the following.

(A4'') For trader i:

$$u(C) = C \qquad \text{for } C < C_i^* \tag{11a}$$

$$= C_i^* \qquad \text{for } C \geq C_i^*. \tag{11b}$$

We also replace (A5) and (A6) with:

(A5'') For Person 1 the critical level of wealth C_1^* is drawn from a uniform distribution on [0, 0.05). For Person 2

$$C_2^* = \beta(0.05 - C_1^*) \tag{12}$$

where β is the the unobservable realisation of a random variable distributed uniformly on (0, 0.5). For Person 3 it is also the case that

$$C_3^* = \beta(0.05 - C_1^*). \qquad (13)$$

(A6'') Agents learn their C_i^* at t = 0. The realisations are private information.

The following assumption is also added.

(A15) There is some chance that the security will be retired early. The retirement time t_R is drawn from a uniform distribution on (0, 1] with density $f_R(t) = 0.1$ and mass of 0.9 on $t_R = 1$. When the security is retired the person holding it at that time receives the fundamental which is still normalized to zero. Denoting the distribution function of t_R by $F_R(t)$ it follows that the expected utility of holding the share until time t is:

$$Eu(t) = u[C(t)][1 - F_R(t)] + u(0)F_R(t) \qquad (14)$$

It can readily be seen that for C in the relevant range $0 \le C \le 0.05$ a trader will sell the share when his wealth reaches C_i^*. This is because the utility from holding the share (assuming he is not the last person) is increasing when wealth is below C_i^* and decreasing when it is above:

$$\frac{dEU}{dt} = u'(C)\,\frac{dC}{dt}\,[1 - F_R(t)] - \{u[C(t)] - u(0)\}f_R(t) \qquad (15)$$

where $C = (1 - \alpha)[p(t)^e - p_B]$ and p_B is the purchase price of the share. Therefore, $\frac{dC}{dt} = (1 - \alpha)\,p'(t)^e$. Recall that $p'(t)^e = 1$ in our example. Consequently in this case,

$$\frac{dEU}{dt} = [1 - F_R(t)] - 0.1\,C > 0 \quad \text{for} \quad C < C_i^* \qquad (16a)$$

$$\frac{dEU}{dt} = -\ 0.1\ C_i^* \ < 0 \qquad \text{for } C > C_i^* . \qquad (16b)$$

Thus when a trader i's wealth reaches C_i^* the effect is the same as dying in terms of his behavior: he sells and leaves the market. Person 1 sells at $t_1 = C_1^*/0.05$ which is distributed uniformly on [0, 1). The person he sells it to, who we shall define to be Person 2, sells it to Person 3 when his wealth reaches C_2^* which occurs at $t_2 = [C_1^* + \beta(1 - C_1^*)]/0.05$. The distribution of t_2 is thus the same as in Subsection 2(iii). Finally, Person 3 attempts to sell it when his wealth reaches C_3^* which occurs at $t_3 = [C_1^* + 2\beta(1 - C_1^*]/0.05$. The distribution of this is again the same as in Subsection 2(iii).

Since the trader knows his C_i^* initially, the problem can be analyzed as in Subsection 2(iii) above where the trader knows his death time. It can be seen from a comparison of (A5') and (A5'') that the formal arguments are similar since the structure of the C_i^*'s (i = 1, 2, 3) induces the same ordering of exits as was previously assumed by the structure of death times. The main differnce is that (by definition) Person 2 always receives the share after Person 1. Hence two trades always occur and the probability that a person selling the share is Person 1 as opposed to Person 2 is simply the ratio $f_1(t)/f_2(t)$.

We summarize with:

<u>Proposition 3</u>

Given traders' utility functions are of the form (11), there exists a set of self-fulfilling beliefs such that two trades will always occur at a strictly positive price between date 0 and date 1 provided:

(i) $0 < p(t)^e < B$ for all $t \in [0, 1)$;

(ii) $p'(t)^e > 0$ for all $t \in [0, 1)$.

In the case presented there is a dramatic change in traders marginal utilities of consumption. In general this is not necessary. All that is required is that there is some critical consumption level such that the marginal utility of consumption is low enough that it is no longer worthwhile holding onto the security because there is a chance the security will be redeemed. Hence, in principal, any standard utility function, u(.), with a declining marginal utility of consumption can be consistent with bubbles provided the marginal utility of consumption falls to a low enough level.

The assumption concerning the possibility that the security will be retired ensures that it is strictly optimal for the agents to sell the stock when they reach their critical consumption level. One interpretation of this possibility of retirement is bankruptcy of the firm.

An important feature of the case with standard utility functions is that the correlated structure of death times that was assumed in the previous subsections is no longer critical. In the case considered, Persons 2 and 3 have identical utility functions but this is not essential. The critical consumption levels determine the period of time the traders hold the share; the ordering of times at which the traders leave is determined by the order in which they receive the share. The main thing that is important is that traders cannot identify whether or not they are the last person who is prepared to buy the share; it must always be possible that the Person selling the share is the one that was endowed with it so that one other person remains to sell it to. Provided they always attach a positive probability to being able to resell they are strictly better off buying the share and bubbles can exist.

3. The entire stock market

In Section 2 we considered the three traders who trade the stock which experiences the bubble. The three people who trade this stock are strictly better off in expected utility terms from doing this compared to not doing anything even when they are risk averse. Person 1 is endowed with the stock and is able to sell at a positive price. The traders who are not endowed with the stock, Persons 2 and 3, are also strictly better off. The reason for this is that the money they invest is not their own. They manage other people's money and keep a share of any of the profits they make. If they are unsuccessful they repay less than they were given to manage and are no worse off than if they had not managed people's wealth. This implies, of course, that the lenders cannot make money or break even by lending to these portfolio managers alone. Why then would anybody be willing to lend to them? In this section we consider a more complete model of the stock market with asymmetric information, where it is optimal for people to lend to portfolio managers using the contract assumed in Section 2.

We suppose that there are three classes of securities in the market: (a) securities which are correctly priced; (b) securities which are underpriced; and (c) the "bubble" securities discussed in the previous section. A value-weighted portfolio of all securities yields an expected return δ. Any ordinary investor can simply invest on his own and obtain this return. Hence this is their opportunity cost. In order to be willing to lend to portfolio managers an ordinary investor must expect a return of at least δ.

There are two types of people who acquire the necessary qualifications to become a portfolio manager. There are good portfolio managers, denoted by the subscript g, who can identify securities which are undervalued. The amount of stock they can identify as being undervalued costs B. The second type of person that attains the qualifications necessary to be a portfolio manager cannot identify undervalued securities. They can only identify the bubble security, and find it optimal to speculate, as described in Section 2. They are denoted by the subscript s. The lender cannot observe the type of the portfolio manager. In this case it may be possible for type s portfolio managers to obtain funds to speculate with even though in a full information world they would not be able to do so.

We replace (A4) with (A4') so that all agents are risk neutral and add the following assumptions to the basic model to consider this argument.

(A16) There is a group of risk neutral lenders who are prepared to lend as much as investment firms require provided that on average their expected return is equal to their opportunity cost, δ, which for simplicity is taken to be zero.

(A17) The sequence of events when portfolio managers are hired by investment firms is the following.

(i) The investment firms offer jobs specifying the contracts for the employees.

(ii) The people qualified to be portfolio managers decide which positions to apply for.

(iii) The investment firms choose which applications to accept.

(A18) The investment firms operate in a competitive industry and so make zero expected profits.

(A19) For ease of exposition we assume that the bad portfolio managers can identify a bubble stock with the price path and distribution of returns considered in Section 2(ii) (a similar analysis can be made for the distributions of returns in the other sections). They each have a probability of 1/3 of being Persons 1, 2 and 3. This implies that the probability distribution of their final gross return y (i.e., including the money they borrow initially), given that B = 1, is distributed as illustrated by the solid line in Figure 5. Person 1 makes a profit which is uniformly distributed between 0 and 1 so their gross profit is uniformly distributed between 1 and 2. Person 2 and Person 3's profit depends on whether or not they are found by Person 1 when he decides to sell the stock. If they are found they make a profit which is uniformly distributed between 0 and 1/3 so that their gross return is uniformly distributed between 1 and 4/3. If they are not found by Person 1 when he sells, they make a loss which is uniformly distributed between 1/3 and 1 so their gross return is uniformly distributed between 0 and 2/3.

(A20) The good portfolio managers can each identify undervalued securities which cost B = 1 (but no more than this). A portfolio of these securities has a stochastic return y which is distributed by the dotted line in Figure 5.

(A21) It is not possible for a lender to observe whether the portfolio manager invests his money in a profitable project or whether it is used for speculation. However, the final value of y is observable.

(A22) The good portfolio managers represent a proportion γ and the portfolio managers who speculate represent a proportion $1 - \gamma$ of those who manage portfolios.

(A23) The total amount of undervalued securities that the good portfolio managers can identify is less than the total amount that exists. Hence the price of these securities is not bid up and the portfolio managers earn a rent from their talent. This means that markets are not strong-form efficient since prices do not reflect all privately available information. However, this is fully consistent with rationality.

(A24) The parameter values are such that the portfolio managers who cannot identify the undervalued stocks are better off speculating than investing in all securities.

The sequence of events outlined in (A17) means that any portfolio manager that chooses a contract at stage (ii) which identifies him as bad will have his application denied and will be unable to manage other people's money. (See Hellwig (1986) for an analysis of the importance of the sequencing of events in a similar context.) This is because bad portfolio managers on average make a loss from speculating for the people whose money they borrow. The implication of this is that bad portfolio managers will behave in exactly the same way as good ones during the application process no matter what contracts the firm offers. This means that any contract which is attractive to good portfolio managers will also attract bad portfolio managers in the same proportions as they exist in the population. Since the investment firms earn zero profits and must earn a return equal to investors' opportunity cost of δ to attract lenders, it follows that the optimal payment schedule must satisfy the following program.

$$\max_{\pi(y)} \quad E_g[y - \pi(y)] \tag{17}$$

subject to

$$\gamma E_g \pi(y) + (1 - \gamma) E_s \pi(y) \geq B(1 + \delta), \tag{18}$$

where E_g denotes the expectation operator with respect to the good portfolio managers' distribution of returns and E_s denotes the expectation operator with respect to the bad portfolio managers that speculate.

It is possible to show the following.

Proposition 4

When $\delta = 0$, $\gamma = 0.095$, $B = 1$ and the portfolio managers are risk neutral, the contract with the linear repayment schedule:

$$\pi^*(y) = 1 + 0.95(y - 1) \qquad \text{for } y \geq 1 \tag{19a}$$

$$= y \qquad \text{for } y < 1. \tag{19b}$$

with a value of α such that lenders earn their opportunity cost is an optimal contract.

To see this consider the first part of the schedule specified in (19a). The good portfolio managers only produce outputs in this region. The expected revenue received from them is given by

$$E_g\pi(y) = \int_1^{4/3} 2\pi(y)dy + \int_{4/3}^{2} (1/2)\pi(y)dy \tag{20}$$

For $y \geq 1$ the expected revenue received from the bad portfolio managers is

$$E_s^1\pi(y) = \frac{2}{3} E_g\pi(y), \tag{21}$$

where the superscript 1 refers to the expectations taken over the range $y \geq 1$. Hence, no matter what the form of the payment schedule $\pi(y)$ the amount of revenue raised from the bad portfolio managers is always 2/3 the amount raised from the good managers; altering the form of the payment schedule for $y \geq 1$ does not enable any more to be extracted from the bad group. It follows that the first part of the schedule in the proposition is optimal.

The second part of the schedule for $y < 1$ is also optimal because only bad portfolio managers produce outputs which fall in this region. The good portfolio managers' utility is unaffected by the form of the payment schedule in this region and lowering the payment below y can only reduce the revenue raised from the bad portfolio managers. Thus, the second part of the payment schedule is optimal.

Since $\delta = 0$ the bad portfolio manager is clearly better off speculating than investing in all securities. The fact that $\delta = 0$ also means it is not worthwhile changing the payment schedule so that the bad portfolio managers choose to invest in all securities since the expected return on these is the same as investing in the bubble security. Hence, the proposition is demonstrated.

In order to derive the proposition it was assumed that all agents are risk neutral. If agents are risk averse then the form of the optimal contract will not be the same as that in Proposition 4; risk sharing will become a factor. Nevertheless, the characteristics of contract will usually be similar. It will be optimal to extract revenue from the bad portfolio managers by penalizing poor performance and rewarding good performance. Although the optimal contract may not have the exact form of a call option, it may often provide incentives for bad managers to speculate and go for large risky payoffs even when this is associated with poor average returns.

4. Robustness and implications

This essay has addressed a theoretical question, namely, can a security trade above its fundamental when there are a finite number of traders with a finite amount of weath, and there is a finite horizon? We have shown that their exists a class of models different from those considered by Tirole (1982) where rational behavior is consistent with security price bubbles. The bubbles can grow at any rate, at least for short periods. If one imagines repetitions of the model, nothing rules out bubbles starting again after they have crashed in the previous period. They can occur on finitely lived securities. Clearly similar bubbles could occur with infinite-lived securities. Perhaps most importantly, the model explains the setting in which bubbles can arise and shows when and how they end.

We have demonstrated the existence of bubbles by considering a specific example. For tractability the assumptions made were very specific. This was necessary to enable a set of self-fulfilling beliefs that ensure existence of equilibrium to be identified. An important question concerns the robustness of this example. In other words how general is the class of models in which bubbles can arise? There are four elements of the example that appear crucial to the result:

1. At any point in time there must be an infinite number of trading possibilities before the horizon.
2. Agents must be unable to deduce whether or not they are the last person in the market.
3. Some owners of wealth invest indirectly so that investment decisions are made by portfolio managers and there is an agency relationship.
4. Markets are inefficient so that there exists a group of portfolio managers that makes an above normal rate of return which allows the losses of the bad portfolio managers to be covered.

We discuss each of these points in turn and then make some final comments.

It is clear that bubbles of the type considered in this paper require an infinite number of trading possibilities. However, the example we presented is not isomorphic to an infinite-horizon overlapping generations model; there are a finite number of agents in our model whereas in an overlapping generations model there is an infinite number of agents. Nevertheless, it is possible to reinterpret the model here as an infinite-horizon model. Tirole (1982) showed that in infinite-horizon models with a finite number of agents bubbles cannot exist because of the finite wealth argument and the zero sum game argument. In our model these arguments are not relevant.

The second point concerns what information agents have. What is critical for our result is that agents have an identification problem. In particular, they must not be able to deduce whether or not they are the last person. In the example presented, adding one piece of information allows traders to determine whether they are the last person. However, this does not mean that the result is not robust since, as we showed when analyzing the model with unknown death times, whenever additional information is provided to traders, adding noise recreates the setting in which the logical sequence of conjectures a rational trader goes through will lead to a bubble.

The third point was the necessity of an agency relationship between investors and the people that make investment decisions on their behalf. Given the existence of this type of agency relationship, it is necessary that the bad managers always pool with the good managers. If the contracts are designed to penalize bad performance and reward superior performance, there will always be a tendency for the bad managers to have an incentive to speculate. These types of agency problems are well understood in corporate finance contexts, though their role in theories of asset pricing is not. In the example above the bubble can be thought of as a manifestation of the inefficiency resulting from the agency relationship.

Finally, it is necessary that securities markets are not strong-form efficient. There are a number of ways in which markets may not be strong-form efficient. We modelled this inefficiency in a particularly simple way by assuming that the supply of good portfolio managers was insufficient to bid up the prices of the undervalued stocks. All that is really required is some form of inefficiency where one group can outperform another. For example, a version of the Grossman and Stiglitz (1980) model where a group of traders has a comparative advantage at gathering information, will lead to similar results.

Any arguments concerning the generality of the example presented are clearly only speculative. The important issue for future research is to identify more precisely how general the class of models where bubbles exist is.

References

Blanchard, Olivier (1979), "Speculative Bubbles, Crashes and Rational Expectations," Economic Letters 3, 387-89.

__________ and Mark Watson (1982), "Bubbles, Rational Expectations, and Financial Markets," in P. Wachtel (ed.), Crises in the Economic and Financial Structure, (Lexington Books; Lexington, Mass.).

Camerer, Colin (1987), "Bubbles and Fads in Asset Prices: A Review of Theory and Evidence," The Wharton School, University of Pennsylvania, mimeo.

DeLong, J. Bradford, Andrei Shleifer, Lawrence Summers, and Robert Waldmann, "The Economic Consequences of Noise Traders," Manuscript, Harvard University, 1988.

Diba, Behzad and Herschel Grossman (1988), "Rational Inflationary Bubbles," Journal of Monetary Economics 21(1), 35-46.

Flavin, Marjorie (1983), "Excess Volatility in Financial Markets: A Reassessment of the Empirical Evidence," Journal of Political Economy 91, 929-56.

Flood, Robert and Peter Garber (1980), "Market Fundamentals Versus Price-Level Bubbles: The First Tests," Journal of Political Economy 88, 745-70.

Grossman, Sanford and Robert Shiller (1981), "The Determinants of the Variability of Stock Market Prices," American Economic Review 71, 222-27.

Grossman, Sanford and Joseph Stiglitz (1980), "On the Impossibility of Informationally Efficient Markets," American Economic Review 70, 393-408.

Harrison, J. Michael and David Kreps (1978), "Speculative Behavior in a Stock Market with Heterogeneous Expectations," Quarterly Journal of Economics 92, 323-36.

Hellwig, Martin, (1986), "Recent Developments in the Theory of Insurance," unpublished seminar notes.

Kleidon, Alan (1986a), "Bias in Small Sample Tests of Stock Price Rationality," Journal of Business 59, 237-61.

__________ (1986b), "Variance Bounds Tests and Stock Price Valuation Models," Journal of Political Economy 94, 953-1001.

Marsh, Terry and Robert C. Merton (1986), "Dividend Variability and Variance Bounds Tests for Rationality of Stock Prices," American Economic Review 76, 483-98.

Shiller, Robert (1981), "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," American Economic Review 71, 421-36.

__________ (1984), "Stock Prices and Social Dynamics," Brookings Papers on Economic Activity 2, 457-98.

Smith, Vernon, Gerry Suchanek and Arlington Williams (1988), "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," Econometrica 56, 1119-1151.

Tirole, Jean (1982), "On the Possibility of Speculation Under Rational Expectations," Econometrica 50, 1163-81.

_________ (1985), "Asset Bubbles and Overlapping Generations," Econometrica 53, 1499-1528.

Wallace, Neil (1980), "The Overlapping Generations Model of Fiat Money," in J. Karaken and N. Wallace (Eds.), Models of Monetary Economies, (Federal Reserve Bank of Minneapolis; Minneapolis).

West, Kenneth (1988), "Bubbles, Fads and Stock Price Volatility Tests: A Partial Evaluation," Journal of Finance 43, 639-56.

FIGURE 1

Example with Unknown Death Times

TRUE VALUES OF DEATH TIMES: $t_1 = 0.1$; $t_2 = 0.4$; $t_3 = 0.7$.

SEQUENCE OF EVENTS

$t =$ 0	0.1	0.4	0.6	0.7
PERSON 1 endowed with share	PERSON 1 finds out he's going to die. He searches for a buyer. Suppose he finds PERSON 2 who buys.	PERSON 2 searches for a buyer and finds PERSON 3 who buys. PERSON 2 dies.	PERSON 3 searches for a buyer. None is found. The bubble bursts.	PERSON 3 dies.

FIGURE 2

Distribution of Unknown Death Times

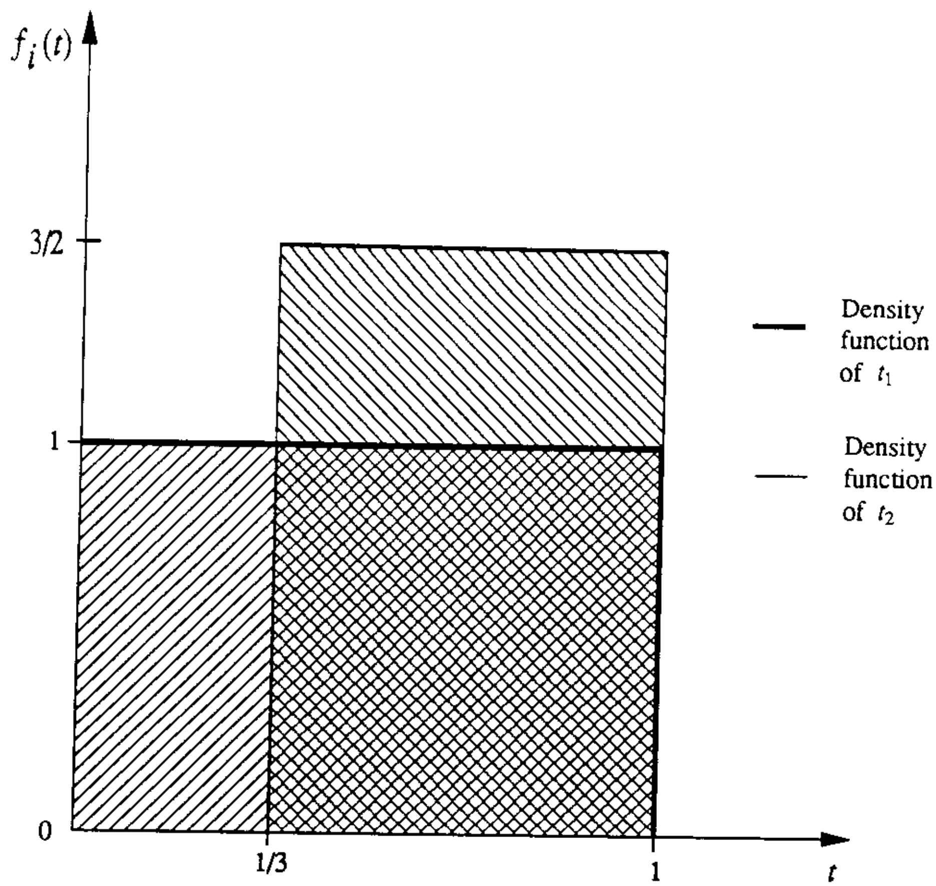

FIGURE 3

Example with Known Death Times

TRUE VALUES OF DEATH TIMES: $t_1 = 0.2$; $\beta = 0.25$; $t_2 = 0.4$; $t_3 = 0.6$.

SEQUENCE OF EVENTS

$t =$ 0	0.2	0.4	0.6
PERSON 1 endowed with share	PERSON 1 searches for a buyer. Suppose he finds PERSON 2 who buys.	PERSON 2 searches for a buyer and finds PERSON 3 who buys. PERSON 2 dies.	PERSON 3 searches for a buyer. None is found. The bubble bursts. PERSON 3 dies.

FIGURE 4

Distribution of Known Death Times with Noise

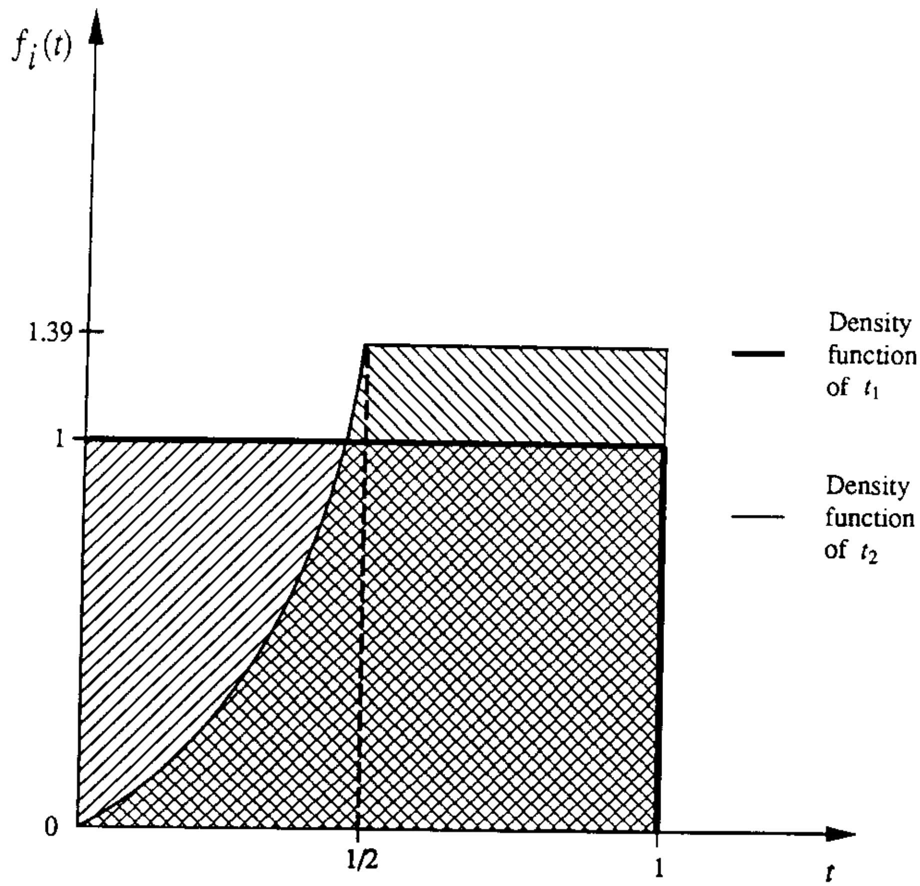

FIGURE 5

Probability Distributions of Returns by Agent Type