

Ex ante versus interim rationality and the existence of bubbles[★]

Sugato Bhattacharyya¹ and Barton L. Lipman²

¹ Department of Finance, School of Business, University of Michigan, Ann Arbor, MI 48109, USA

² Department of Economics, Queen's University, Kingston, Ontario K7L 3N6, CANADA

Received: April 14, 1992; revised version August 24, 1994

Summary. Tirole (1982) is commonly interpreted as proving that bubbles are impossible with finitely many rational traders with common priors. We study a simple variation of his model in which bubbles can occur, even though traders have common priors and common knowledge that the asset has no fundamental value. In equilibrium, agents purchase the asset at successively higher prices until the bubble “bursts” and no subsequent trade occurs. Each trader’s initial wealth determines the last date at which he could possibly trade. The date at which the bubble bursts is a function of these finite “truncation dates” for the individual traders. Since initial wealth is private information, no trader knows when the bubble will burst. There are two key differences between our model and Tirole’s which enable us to construct equilibrium bubbles this way. First, Tirole requires ex ante optimality, while we only require every trader’s strategy to be optimal conditional on his information – i.e., interim optimal. As we argue in the text, this would seem to be the relevant definition of optimality. Second, Tirole considers competitive equilibria, while we analyze a simple bargaining game.

1. Introduction

It is very difficult to believe that the market’s perception of the fundamental value of the Dow Jones stocks plummeted so sharply on one day as to cause the stock market crash of 1987. Yet if a change in fundamentals did not cause the crash, must this be taken as evidence that traders are not rational? The pathbreaking work of Tirole (1982), Milgrom and Stokey (1982), and Sebenius and Geanakoplos (1983) is

[★] We thank Abhijit Banerjee, James Dow, Itzhak Gilboa, Debra Holt, Fallaw Sowell, Sanjay Srivastava, Stan Zin, seminar participants at Pennsylvania State University, Indian Statistical Institute at Calcutta, Carnegie Mellon, the University of Rochester, the University of Maryland, the London Business School, the University of British Columbia, and the Sixth World Congress, and Michael Woodford and an anonymous referee for helpful comments. Both authors were at Carnegie Mellon University when this research was begun. The second author also thank the Social Sciences and Humanities Research Council of Canada for financial support.

commonly interpreted as showing that if all traders are perfectly rational and have the same prior beliefs, then, even if they receive different information, assets must be priced according to fundamentals in equilibrium. In other words, bubbles – divergences in asset prices from fundamentals – are impossible with rational traders. These results would seem to force us to conclude either that the 1987 crash was simply a major one-day change in the market's perception of fundamentals or that it proved that traders are not rational (or have different priors). On a less dramatic note, it is difficult to reconcile the view that prices are determined by fundamentals with observed trading in "fundamental-less" assets such as stamps, coins, baseball cards, or wine. What market fundamentals make an undrinkable 1804 Chateau Lafite worth \$25,000?¹

Various authors have constructed models in which bubbles do occur. For example, it is well-known that bubbles can occur with myopic traders as in Blanchard (1979) or with infinitely many traders as in Tirole (1985), Weil (1987), and Jackson and Peck (1991). Allen and Gorton (1993) showed that bubbles can occur because of the preferences that agency problems induce for portfolio managers. Dow, Madrigal, and Werlang (1990) have shown that Milgrom and Stokey's no-trade results do not hold in general without expected utility preferences. Allen, Morris, and Postlewaite (1993) showed that in equilibrium, it can be true that all traders know the price of an asset will fall as long as this fact is not common knowledge. Finally, Neeman (1993) has shown that bubbles are possible if rationality is mutual knowledge, even of very high order, as long as it is not common knowledge.

Instead of following these approaches to constructing equilibrium bubbles, we consider a very simple model with perfectly rational (expected utility maximizing – in fact, risk neutral) traders who have the same prior beliefs. The model is the same as that of Tirole (1982), with two important changes discussed below. There are two goods, money and shares of an asset. It is common knowledge that this asset has no fundamental value at all. The traders are fully dynamically optimal. In the equilibria we construct, the agents trade the asset at successively higher prices until the bubble "bursts" and no subsequent trade occurs. We assume that cash changes hands with each trade and that it is not possible to borrow cash. Consequently, each trader's initial cash holdings determine for him a last possible date at which he could trade. In equilibrium, the date at which the bubble bursts is a function of these finite "truncation dates." Since no trader knows everyone's truncation date, none knows when the bubble will burst. The agents are willing to trade because no one knows in advance exactly when the bubble will burst.

There are six differences between our model and Tirole's (1982) model, two important and four unimportant ones. To deal with the irrelevant points first, unlike Tirole, we assume that traders do not discount future returns – i.e., the discount factor is 1. It would complicate the notation but not substantively change any results

¹ See the wine column in the 12/15/91 *New York Times* which also contains the following: "Wine collectors love to tell the story about the traders who keep selling each other the same consignment of canned herring. Finally, one of them opens a can. 'Why, this is terrible stuff,' he tells the chap who just sold it to him. The other fellow smiles. 'Herring,' he says, 'is for trading, not for eating.'"

to allow for discounting. Second, as mentioned above, we assume that it is common knowledge that the asset is worthless. Again, it would only complicate the notation to allow for a dividend stream, at least if we assume that traders are symmetrically informed about dividends. Third, again to avoid irrelevant complications, we assume that there are only two traders. With more traders, one could obtain very similar results by “grouping” the traders into two groups and constructing equilibria in which traders in a given group behave identically. Finally, while Tirole does consider the implications of short sales constraints, he does not impose the cash constraint we consider. As we discuss in the conclusion, however, one can generate very similar kinds of bubbles without this constraint.

The first important difference between our model and Tirole’s is that he requires ex ante optimality, while we require interim optimality. In our model, each trader is optimal conditional on his initial wealth – that is, each is interim optimal. (Each is also sequentially rational in the sense of Kreps and Wilson (1982). Note that sequential rationality only refers to optimality at an information set – not any sort of ex ante optimality.) Since each trader knows his initial wealth before any decisions must be made, this would seem to be the relevant definition of optimality. However, unconditional expected profits are not defined, so that one cannot say whether the traders are ex ante optimal.² (They are ex ante optimal in the weak sense that there is no strategy which is ex ante better – but there is no strategy which is ex ante worse either.) While there are strong philosophical arguments in favor of the consistency of beliefs imposed by common priors (see Aumann (1987), for example), we see no such arguments in favor of the consistency imposed by existence of the ex ante expectation. Furthermore, as we show, requiring the ex ante expectation to exist, by itself, rules out a very large class of intuitively plausible bubbles.

The second important difference is the modeling of price determination: Tirole uses competitive equilibria, while we analyze a simple bargaining game. As we explain in Section 3, our approach to constructing equilibrium bubbles does not work with competitive equilibrium, not because of rationality considerations, but because of price-taking. It is easy to extend our approach to demand submission games as we explain in Section 3. We conjecture that our approach could be adapted to other models with imperfectly competitive price determination, such as Kyle (1985) or Glosten and Milgrom (1985), to generate bubbles.

This paper is organized as follows. The model is presented in Section 2. In Section 3, we analyze the model and show that a very large class of bubbles are possible. We then explain the intuition for the difference between our result and Tirole’s in the context of a simple example. We also explain the connection between our result and the “envelope switching problem” discussed by Nalebuff (1989) and others. In Section 4, we show that a much larger class of bubbles are possible if we also add “sunspots” to the model. (This addition is consistent with Tirole’s assumptions.) We offer some concluding remarks in Section 5. All proofs are contained in the Appendix.

² As we explain in Section 3, the fact that expectations are undefined is related to Nalebuff’s (1989) demonstration that expected payoffs in the “envelope switching problem” must be infinite to generate the paradox.

2. The model

There are two traders, 1 and 2, both risk neutral. The two goods in the economy are shares of an asset and “money” or “wealth.” Trader i 's initial endowment is e_i shares of the asset and $w_i \geq 0$ units of money. For simplicity, we assume that $e_1 = 1$ and $e_2 = 0$. The endowments of the asset are common knowledge, but the initial endowments of wealth are not. In particular, there are distribution functions F_i which gives the common prior of the traders over the initial wealth of trader i . (Thus, in particular, we assume that the initial wealths are independently distributed.) Each trader learns his own initial wealth endowment only. We assume that

$$F_i(w) = 0, \forall w < 0, \forall i \tag{A.1}$$

and

$$\lim_{w \rightarrow \infty} F_i(w) = 1, \forall i. \tag{A.2}$$

In other words, the initial wealth endowments are nonnegative and finite with probability 1. We do not require the wealths to be bounded – in fact, later we will show that the equilibria we construct require unbounded support. Let $F = (F_1, F_2)$.

The model can be interpreted differently. In particular, what we refer to as “initial wealth” can be replaced by a stream of income that the trader receives in each period from some other investments. Under this interpretation, it is the per period payoffs from these other investments or the number of periods in which payoffs accrue which is private information (Rationalizing certain stochastic processes may require per period returns to vary over time with this interpretation, however.) Since there is an infinite horizon, it does not seem unreasonable to treat the number of periods in which other investments pay returns as unbounded and hence the total returns as unbounded.³

It is common knowledge that the asset pays no dividends ever. Trade may occur at dates, $t = 1, 2, \dots$. One can interpret the model as taking place in continuous but finite time (as in Allen and Gorton (1993)) but where trade can only take place at countably many points in time. Traders do not discount future returns. To be more precise, a trader's payoff given some finite or infinite sequence of trades is just given by the limit (or \liminf if the limit does not exist⁴) of his accumulated wealth over the sequence. That is, if the sequence of i 's holdings of the asset is $\{x_t^i\}$ and the sequence of prices is $\{p_t\}$, then i 's payoff is

$$\liminf \sum_{t=1}^T p_t(x_{t-1}^i - x_t^i)$$

³ Even with this interpretation, the unboundedness may seem unrealistic – after all, shouldn't each trader recognize that the other trader does not and never will have more dollars than, say, the number of protons in the known universe (believed to be about 10^{126})? On the other hand, as argued by Rubinstein (1991) for example, the key issue is not what the real world is like but how it is viewed by the agents of the model. If neither agent thinks about a finite upper bound for the possible wealth level of the other agent, then the assumption of unbounded supports is appropriate.

⁴ Using the \liminf guarantees that the traders are “pessimistic.” It is easy to construct infinite bubbles if both traders use the \limsup since each would view exchanging a dollar back and forth forever favorably. With the \liminf , this problem does not arise.

where $x_0^i = e_i$. In all the equilibria we analyze, every sequence of trades with positive probability is finite, so the limit always exists. The preferences are common knowledge.

We also assume that the exchange agreed upon occurs at the date of agreement. Hence when a trade takes place, the traders' wealths adjust accordingly. We assume that there is no outside source for loans, so that trades where one player gives up more than the amount of wealth he has at that date are not feasible.

The sequence of trades and prices is determined by bargaining. As we explain in more detail later, we obtain identical results with a wide variety of bargaining rules as well as with a demand submission game similar to that studied by Kyle (1989), Jackson (1991), and Peck, Shell, and Spear (1989). The bargaining game we focus on is basically the Nash (1953) demand game. At each date t , each trader submits an offer, y_t^i . An offer is a triple (p, x^1, x^2) , where p is the price and x^i the net position for trader i which is being proposed. Naturally, we require an offer (p, x^1, x^2) to satisfy $x^1 + x^2 = 1$. Let Y denote the set of such offers.

We also assume that a trader cannot offer to take more of the asset than he can actually afford to purchase. In particular, if the sequence of past prices and i 's positions is given by p_1, \dots, p_{t-1} and x_1^i, \dots, x_{t-1}^i , then the offer i makes in period t , (p, x^1, x^2) , must satisfy:

$$p[x^i - x_{t-1}^i] \leq w_i - \sum_{k=1}^{t-1} p_k[x_k^i - x_{k-1}^i] \tag{A.3}$$

where $x_0^i = e_i$. Note that x^i may be negative, so offers involving short sales are allowed.⁵ If the offers are identical – that is, $y_t^1 = y_t^2 = (p, x^1, x^2)$ – then trade takes place at price p and positions adjust to those proposed. In this case, $p_t = p$, trader i 's new position is $x_t^i = x^i$, and trader i 's new wealth level is

$$w_i - \sum_{k=1}^t p_k[x_k^i - x_{k-1}^i].$$

Our results do not depend on what happens if the offers are “compatible” in some sense but not identical.⁶ For simplicity, we assume that no trade occurs, so that $x_t^i = x_{t-1}^i$. As a notational convention, we set $p_t = 0$ in this event.⁷

A *history for trader i* summarizes everything i knows at a particular point in the game. Hence it lists his initial wealth, the offers he has submitted, his holdings of the asset at each past date, and the sequence of past prices. (Of course, some of this information is redundant.) More formally, for $t \geq 1$,

$$h_t^i = (w_i, y_1^i, \dots, y_{t-1}^i, x_1^i, \dots, x_{t-1}^i, p_1, \dots, p_{t-1}) \in \mathbf{R}_+ \times Y^{t-1} \times \mathbf{R}^{t-1} \times \mathbf{R}_+^{t-1}$$

⁵ As is normally assumed, we assume short sales do not require the posting of any margin.

⁶ For example, if the two traders propose the same positions but different prices with the net buyer offering to pay more than the net seller asks, we could assume that they trade at the average of the two offered prices.

⁷ What is really relevant in our analysis is the total amount of money exchanged at each point. If there is no trade, then no money exchanges hands and so the price is irrelevant to us.

is a t length history for i if the following two conditions hold. First, for all $k < t$, either $y_k^i = (p_k, x_k^i, 1 - x_k^i)$ or $p_k = 0$ and $x_k^i = x_{k-1}^i$. Second, for all $k < t$, if $y_k^i = (p, x^1, x^2)$, then

$$p[x^i - x_{k-1}^i] \leq w_i - \sum_{j=1}^{k-1} p_j[x_j^i - x_{j-1}^i].$$

In other words, such a sequence is a history if it is feasible and consistent with the way the market operates. Let H_t^i denote the set of t length histories for i . Note that we are assuming that trader i is the only one who observes the offers he makes or his position in the asset over time, while prices are observed by both traders. Of course, when there are only two traders, each can infer the holdings of the other. Also, when trade occurs, each can infer the offer of the other. Even with more than two traders, though, our results would not be changed if all trades and offers were commonly observed or if some summary statistics of trades (such as volume) were commonly observed. Given a history h_t^i , let $Y(h_t^i)$ denote the set of feasible offers for i – that is, it is the set of functions y_t^i satisfying (A.3).

A strategy for i is a sequence of functions $\sigma_t^i: H_t^i \rightarrow Y, t = 1, 2, \dots$, where $\sigma_t^i(h_t^i) \in Y(h_t^i)$ is the offer trader i proposes at date t as a function of everything he has observed up to this point. Let $\sigma^i = (\sigma_1^i, \sigma_2^i, \dots)$ and $\sigma = (\sigma^1, \sigma^2)$.

Each trader has beliefs about the complete history of the game. Formally, let $H_t = H_t^1 \times H_t^2$ and let $\Delta(H_t)$ denote the set of probability distributions defined on the Borel sets of H_t . A belief for trader i is a sequence of functions $\delta^i = (\delta_1^i, \delta_2^i, \dots)$ satisfying the following conditions. First, $\delta_t^i: H_t^i \rightarrow \Delta(H_t)$. That is, $\delta_t^i(h_t^i)$ is a probability distribution over H_t as a function of i 's history. Second, if $\hat{h}_t^i = (\hat{h}_t^1, \hat{h}_t^2)$ is in the support of $\delta_t^i(h_t^i)$, then $\hat{h}_t^i = h_t^i$. Third, if $h_t = (h_t^1, h_t^2)$ is in the support of $\delta_t^i(h_t^i)$, then the sequence of prices in h_t^i exactly matches the sequence of prices in h_t^j . Fourth, if (h_t^1, h_t^2) is in the support of $\delta_t^i(h_t^i)$, then the sequence of offers in h_t^j exactly matches the sequence of offers in h_t^i for every period in which the price was not zero. Finally, if (h_t^1, h_t^2) is in the support of $\delta_t^i(h_t^i)$, then for every $k \leq t - 1, 1 - x_k^j = x_k^i$. In other words, $\delta_t^i(h_t^i)$ can only give positive probability to histories for the other trader which the consistent with the prices i has observed, the trades i has made, and market clearing. In effect, then, i 's belief about the complete history of the game is a belief regarding j 's initial wealth and the sequence of offers submitted by j in periods of no trade. Let $\delta = (\delta^1, \delta^2)$.

Any strategies σ together with the distributions on initial wealths $F = (F_1, F_2)$ generates a stochastic process on wealth, prices, and trades. This, in turn, generates a stochastic process on private histories for each trader. Let $\psi_i(\sigma, F)$ denote the induced distribution on $H_t = H_t^1 \times H_t^2$. Let $\mathcal{H}_t^i(\sigma, F)$ denote the support of ψ_i , projected onto H_t^i . An equilibrium is a pair $(\bar{\sigma}, \bar{\delta})$ satisfying the following two conditions. First, for every $t \geq 1$, for every $h_t^i, \bar{\sigma}_t^i(h_t^i)$ is a best reply to $\bar{\sigma}^j (j \neq i)$ given $\bar{\delta}_t^j(h_t^j)$. Second, for every $h_t^i \in \mathcal{H}_t^i(\bar{\sigma}, F), \bar{\delta}_t^i(h_t^i)$ is generated from $\psi_i(\bar{\sigma}, F)$ by Bayes' Rule. Notice that the first condition requires strategies to be optimal conditional on each history. That is, there is no requirement that the strategies be optimal unconditionally – i.e., optimal on the basis of ex ante expectations. The importance of this point will be seen in the next section.

3. Equilibrium bubbles

While our model is very similar to Tirole's, our results are dramatically different. Bubbles are impossible in his model in the sense that the only equilibrium price process has $p_t = 0$ for all t . Here the opposite is true: not only are other processes supportable in equilibrium, but a plethora of bubbles are supportable. To state this precisely, we first describe the set of price paths we will study.

We say that a stochastic process on prices $\{\tilde{p}_t\}$ is *nice* if it satisfies the following:

- (1) For all t , $\text{supp}(\tilde{p}_t)$ is a finite subset of \mathbf{R}_+ , with $\text{supp}(\tilde{p}_1) \neq \{0\}$.
- (2) $\{\tilde{p}_t\}$ is a submartingale. That is, for all $(p_1, \dots, p_{t-1}) \in \text{supp}(\tilde{p}_1, \dots, \tilde{p}_{t-1})$,

$$E[\tilde{p}_t | p_1, \dots, p_{t-1}] \geq p_{t-1}.$$

- (3) $\text{Pr}[\tilde{p}_{t+1} = 0 | p_t = 0] = 1$.
- (4) There exists $\varepsilon > 0$ such that for all $(p_1, \dots, p_{t-1}) \in \text{supp}(\tilde{p}_1, \dots, \tilde{p}_{t-1})$ with $p_{t-1} > 0$,

$$\text{Pr}[\tilde{p}_t = 0 | p_1, \dots, p_{t-1}] \geq \varepsilon.$$

Define the stochastic process $\{\tilde{\omega}_t\}$ from $\{\tilde{p}_t\}$ by

$$\tilde{\omega}_t = \sum_{j=0}^{t-1} (-1)^j \tilde{p}_{t-j}.$$

For each t , let $\bar{\omega}_t$ ($\underline{\omega}_t$) denote the maximum (minimum) point in the support of $\tilde{\omega}_t$ conditional on $p_t > 0$. Then

- (5) $\underline{\omega}_{t+1} > \bar{\omega}_{t-1}$ for all t where $\bar{\omega}_0$ and $\bar{\omega}_{-1}$ are defined to be 0).

We will say that $\{\tilde{p}_t\}$ is *very nice* if it is nice and

- (1') For all t , $\text{supp}(\tilde{p}_t) = \{0, p_t^*\}$ where $p_t^* > 0$.

In this section, we will focus on very nice stochastic processes and so will use condition (1') instead of (1). Either condition obviously simplifies the analysis substantially. In particular, (1') will enable us to construct equilibria with a very intuitive structure. In the next section, we show how we can construct equilibria with a similar structure to rationalize nice stochastic processes if we incorporate sunspots into the model. We rule out the case where $\text{supp}(\tilde{p}_1) = \{0\}$ because, together with (3), this case corresponds to the "bubble-less" (and trivially supportable) process $\tilde{p}_t = 0$ with probability one for all t .

Condition (3) says that pricing according to fundamentals is an absorbing state. Put differently, if the bubble bursts, it stays burst. If we were to strengthen (2) to require that $\{\tilde{p}_t\}$ be a martingale, then (1) and this stronger version of (2) would imply (3).

Condition (4) requires that the probabilities that the bubble bursts at period t are bounded away from zero. It is easy to use (4) to show that

$$\lim_{T \rightarrow \infty} \text{Pr}[\tilde{p}_T > 0] = 0.$$

In other words, the bubble bursts in finite time with probability 1. In fact, one can show that (4) implies that the expected duration of the bubble is finite.

Finally, condition (5) simplifies the analysis substantially. To understand the meaning of this condition, suppose that the traders trade the one unit of the asset back and forth over time, trading in period t at price \tilde{p}_t . Then $\tilde{\omega}_t$ is precisely the trading profit through period t earned by the agent who sells the asset at period t . Condition (5), then, says that these profits necessarily are higher the longer the bubble continues. One implication is that profits increase along any possible path, so that prices always increase until the bubble bursts. Note that (5) is stronger than simply saying that profits increase along any possible path – it says that prior to knowing the price path that will be realized, each agent prefers that the bubble continue as long as possible (as long as he is the last seller!).

Notice also that when we use (1') instead of (1), we have $\tilde{\omega}_t = \underline{\omega}_t$ for all t since (1') implies that there is only one possible path for prices up to period t if the bubble does not burst before t . Furthermore, it is not hard to use (2) to show that this path of prices must be strictly increasing. Hence (1') and (2) imply (5).

One important and nontrivial implication of these conditions is the following. Given a stochastic process $\{\tilde{p}_t\}$, let \tilde{p} denote the random variable giving the last strictly positive price (where we define the realization to be zero if $p_1 = 0$). Then

Theorem 1. *For any stochastic process $\{\tilde{p}_t\}$ satisfying (1), (2), (3), and (4), $E(\tilde{p}) = \infty$.*

Since every nice stochastic process satisfies (1) through (4), the same holds for all nice and very nice stochastic processes.

We will show that every very nice stochastic process is generated as an equilibrium outcome given some F . To state this precisely, note that any strategies σ together with the priors on initial wealths F induces a stochastic process $\{\tilde{p}_t\}$ on prices. Let $\mathcal{P}(\sigma, F)$ denote this mapping. We say that a given stochastic process on prices, say $\{\tilde{p}_t\}$, is *rationalizable* if there exists F satisfying (A.1) and (A.2) and an equilibrium (σ, δ) such that $\mathcal{P}(\sigma, F) = \{\tilde{p}_t\}$.

Rather than characterizing rationalizable processes, we will focus on a stronger notion which requires the equilibrium to have a particular form. Given a very nice stochastic process $\{\tilde{p}_t\}$, we will say that strategy vector σ is a *canonical* strategy for $\{\tilde{p}_t\}$ if it has the property that trade occurs at period t at price p_t^* if this has always happened in the past and the buyer has enough wealth to purchase. More formally, for that history of length t where the sequence of prices has been $(p_1^*, \dots, p_{t-1}^*)$ and where the entire unit of the asset has been traded in each period, the trader holding the asset offers to trade the entire unit at price p_t^* and the other trader offers to buy it at this price if he has enough wealth – that is, if the offer is feasible in the sense of satisfying (A.3). Otherwise, each trader refuses to trade – i.e., each makes an offer which has his position in the asset unchanged. We will say that $\{\tilde{p}_t\}$ is *canonically rationalizable* if some canonical strategies for $\{\tilde{p}_t\}$ (together with some beliefs) form an equilibrium.

In other words, in the equilibria we construct, if trade occurs at period t , it occurs at price p_t^* . Every time trade occurs, the seller sells his entire holdings of the asset to the buyer. Trade occurs at period t iff trade has occurred in every past period at the equilibrium price and the trader who is supposed to buy the asset at period t has

accumulated enough wealth to pay p_t^* . Intuitively, then, there is a bubble until some trader cannot afford to continue trading at which point the bubble bursts.

Theorem 2. *Every very nice stochastic process is canonically rationalizable.*

We prove this theorem by fixing an arbitrary very nice stochastic process. We then construct an F and an equilibrium given F which together generate the given process for prices. As canonical rationalizability requires, the equilibrium we construct has the property that trade continues until one trader cannot afford to continue trading. In equilibrium, the traders both know how much money they will need at any given period in order to trade at that point, given that the bubble has not yet burst. Hence trader i 's initial wealth determines for him a finite "truncation date" – that is, a finite upper bound on the number of periods he will trade, given by the first period at which he will be unable to trade even if the bubble has not yet burst. Given this finite truncation, it is not difficult to use backward induction to show that each trader wishes to engage in the trades we construct for him.

To see this construction and the importance of the finite truncation device more concretely, consider the following example. Fix any $z > 1$ and define a stochastic process for prices as follows. Let

$$\tilde{p}_1 = \begin{cases} z, & \text{with probability } 1/z; \\ 0, & \text{otherwise.} \end{cases}$$

For $t = 2, 3, \dots$, let

$$\tilde{p}_t = \begin{cases} zp_{t-1}, & \text{with probability } 1/z; \\ 0, & \text{otherwise,} \end{cases}$$

where p_{t-1} denotes the realization of \tilde{p}_{t-1} . In other words, the period t price is z^t if the bubble does not burst before t . It is easy to see that this stochastic process is very nice. To show that this process is rationalizable for this economy, we construct wealth distributions and equilibrium strategies for the two players.

To construct these distributions, note that if the single unit of the asset is traded in every period up to and including $2n$, then trader 2 (who would be buying the asset in odd periods and selling it in even ones) will have wealth of

$$w_2 + \sum_{j=1}^{2n} (-1)^j p_j = w_2 + \sum_{j=1}^{2n} (-z)^j = w_2 - z \left[\frac{1 - z^{2n}}{1 + z} \right].$$

Trader 2 can purchase the asset in period $2n + 1$ iff

$$w_2 - z \left[\frac{1 - z^{2n}}{1 + z} \right] \geq z^{2n+1}$$

or

$$w_2 \geq z \left[\frac{1 + z^{2n+1}}{1 + z} \right].$$

The fact that trader 2 bought the asset in each period up to $2n - 1$ signals at least that

$$w_2 \geq z \left[\frac{1 + z^{2m+1}}{1 + z} \right]$$

for all $m < n$. Since the right-hand side is strictly increasing in m , we see that the probability that trader 2 has enough wealth to purchase the asset in period $2n + 1$ given that he has had enough wealth to purchase it in each past period is just

$$\Pr \left[w_2 \geq \frac{z + z^{2n+2}}{1 + z} \mid w_2 \geq \frac{z + z^{2n}}{1 + z} \right].$$

Canonical rationalizability requires that it is wealth constraints that cause trade to stop – that is, traders always trade if they can afford to purchase the asset. Hence this is exactly the probability that trade occurs at period $2n + 1$ given that it has occurred in every previous period. (Recall that w_1 and w_2 are independent random variables by assumption.) Since this is supposed to match the stochastic process above, this probability must be $1/z$. Hence we obtain equations:

$$\frac{1 - F_2 \left(\frac{z + z^{2n+2}}{1 + z} \right)}{1 - F_2 \left(\frac{z + z^{2n}}{1 + z} \right)} = \frac{1}{z}, \quad n = 1, 2, \dots \tag{6}$$

where F_2 is the marginal distribution function for w_2 . A similar procedure yields an set of equations for F_1 , the marginal distribution function for w_1 . It is not hard to show that the following distribution functions will solve these equations:

$$F_1(w) = 1 - \frac{1}{\sqrt{w(1+z)+z}}$$

and

$$F_2(w) = 1 - \frac{1}{\sqrt{w(1+z)-z}}.$$

To verify this, notice that the expression on the left-hand side of (6) is

$$\sqrt{\frac{z + z^{2n} - z}{z + z^{2n+2} - z}} = \frac{1}{z}.$$

The argument for F_1 is analogous. Notice that F_1 and F_2 are both well-behaved distribution functions and both converge to 1 as $w \rightarrow \infty$.

The strategies are as follows. Consider any odd period t and any history such that $x_{t-1}^1 = 1$ and p_k was z^k for every $k < t$. In this event, trader 1 makes the offer $(z^t, 0, 1)$. Since trader 1 has the asset at this point, this corresponds to trader 1 wishing to sell the asset at price z^t . If $x_{t-1}^1 \neq 1$ or some past price was zero or inconsistent with the stochastic process above, he “refuses to trade” – more formally, he offers $(0, x_{t-1}^1, 1 - x_{t-1}^1)$.⁸ If $p_k = z^k$ for every $k < t$, $x_{t-1}^2 = 0$, and trader 2 has at least z^t units

⁸ Since trader 1 is demanding that he retain the asset, the price he offers is irrelevant to him.

of money, then trader 2 matches trader 1's offer – that is, he offers $(z^t, 0, 1)$. Otherwise, trader 2 offers $(0, 1 - x_{t-1}^2, x_{t-1}^2)$. Even periods are handled analogously; the roles of the traders reverse with 1 buying and 2 selling, but the strategy of the buyer and of the seller have the same form as above.

Our construction of the initial distributions for wealth guarantees that if the traders follow these strategies, the induced stochastic process will be exactly the process we are seeking to rationalize. Furthermore, because the bubble bursts only when one trader does not have enough money to pay for the asset, the strategies have the appropriate structure for canonical rationalizability.

All that remains is to show that these strategies form an equilibrium. So suppose trader i expects trader j to follow his part of this proposed equilibrium. Consider any period $t \geq 1$ and suppose these strategies have been followed so far. If $p_k \neq z^k$ for some previous period k , then j will refuse to trade – that is, will make an offer whereby he retains his position in the asset. Hence i may as well do likewise.

Suppose instead that $p_k = z^k$ for every previous period k . It is easy to see that j 's strategy effectively leaves i only able to choose how long to trade. That is, in each period t , either j offers to trade at price z^t or he refuses to trade at all. Hence i 's choice is only whether to be willing to accept such a "proposal" or not. The strategy above calls for i to accept these proposals as long as it is feasible for him to do so. Is this optimal for i ?

It is easy to see that if i is supposed to sell at period t , he should always try to sell. If he does not, his continuation payoff is zero. If he does, then he can always refuse to purchase the asset again at period $t + 1$. This strategy yields a strictly positive amount in expectation at period t and zero thereafter and so is strictly preferred to not selling.

If i is supposed to buy at period t , then, at best, he earns zero continuation profits from period t onward. To see this, suppose i always offers to sell at the appropriate price and offers to buy only up to period $T \geq t$. Since his expected revenue in any period always exactly equals the price he paid the previous period, he earns zero expected profits for any finite T . It is easy to see that it is not feasible for him to adopt the strategy of always offering to buy. Recall that he can afford to buy at period t if his initial wealth plus his accumulated trading profits exceed z^t . In the case of $i = 2$ and $t = 2n + 1$, we showed above that this is possible iff

$$w_2 \geq z \left[\frac{1 + z^{2n+1}}{1 + z} \right].$$

Since the right-hand side goes to infinity as $n \rightarrow \infty$, there will be some finite date at which trader 2 will no longer be able to accept trader 1's proposal. Similarly, there is a finite limit on how long trader 1 can continue to buy the asset. Hence following his proposed equilibrium strategy is optimal.

At this point, we may well have tried the reader's patience sorely. How, one might ask, can both traders expect to earn positive gains? After all, since both are risk neutral, there are no gains from trade in this market. If there are no gains from trade possible, then any gain for one trader is a loss for the other. How have we overcome the impossibility theorems of Tirole (1982), Milgrom and Stokey (1982), and Sebenius and Geanakoplos (1983)?

To see what breaks down, let us attempt to imitate the proof of Tirole’s Proposition 1. Adapting his notation and terminology to the example above, his argument runs as follows. Let $h_i(w_1, w_2)$ denote the profits earned by trader i as a function of the two initial wealths in this equilibrium. As argued above, the martingale property implies that the expected continuation profits for a buyer are always zero, so that trader 1’s expected profits for the game are just the expected profits on the very first sale. Hence

$$E_{\tilde{w}_2}[h_1(w_1, \tilde{w}_2)|w_1] = \int_{\tilde{w}_2} h_1(w_1, \tilde{w}_2)dF_2(\tilde{w}_2) = E[\tilde{p}_1] = 1, \quad \forall w_1. \tag{7}$$

Similarly, since trader 2 is supposed to buy at period 1,

$$E_{\tilde{w}_1}[h_2(\tilde{w}_1, w_2)|w_2] = \int_{\tilde{w}_1} h_2(\tilde{w}_1, w_2)dF_1(\tilde{w}_1) = 0, \quad \forall w_2. \tag{8}$$

But for any pair of initial wealths, the profits earned by one trader is exactly equal to the loss earned by the other. That is, $h_2(w_1, w_2) = -h_1(w_1, w_2)$. Substituting into (8),

$$E_{\tilde{w}_1}[h_1(w_1, w_2)|w_2] = \int_{\tilde{w}_1} h_1(\tilde{w}_1, w_2)dF_1(\tilde{w}_1) = 0, \quad \forall w_2. \tag{9}$$

But equations (7) and (9) are inconsistent. If we multiply the integral in (7) by $dF_1(\tilde{w}_1)$ and integrate over \tilde{w}_1 , we see that

$$E_{\tilde{w}_1, \tilde{w}_2}[h_1(\tilde{w}_1, \tilde{w}_2)] = \int_{\tilde{w}_1} \int_{\tilde{w}_2} h_1(\tilde{w}_1, \tilde{w}_2)dF_1(\tilde{w}_1)dF_2(\tilde{w}_2) = 1.$$

However, if we multiply the integral in (9) by $dF_2(\tilde{w}_2)$ and integrate over \tilde{w}_2 , we obtain a contradiction:

$$E_{\tilde{w}_1, \tilde{w}_2}[h_1(\tilde{w}_1, \tilde{w}_2)] = \int_{\tilde{w}_2} \int_{\tilde{w}_1} h_1(\tilde{w}_1, \tilde{w}_2)dF_1(\tilde{w}_1)dF_2(\tilde{w}_2) = 0.$$

This is precisely how Tirole proves that all traders must have expected profits of zero. How is this paradox resolved?

The resolution is quite simple. Fubini’s Theorem (see, e.g., Royden (1968), pg. 269) states that when a double integral exists, its value is given by integrating first over one variable and then over the other in either order. An implication, then, is that if one gets a different answer depending on the order of integration (as we did above), there is no consistent way to define the double integral and, hence, it does not exist. That is, the ex ante expectation of trader 1’s equilibrium profits does not exist! Surprisingly, this is true even though his expected profits conditional on w_1 are well-defined for every w_1 and are even independent of w_1 . Thus conditional expected profits are not only well-defined, but are even common knowledge. It is also important to note that it is not true that ex ante profits are infinite (as they are in the St. Petersburg Paradox, for example) – they simply do not exist.

To see intuitively why the ex ante expectation does not exist, let us try to compute it directly. In the equilibrium constructed above, if $p_t > 0$, trader 1 receives z^t if t is odd and pays z^t if t is even – in short, he receives $(-1)^{t+1}z^t$. Hence his expected profits can be written as

$$\sum_{t=1}^{\infty} (-1)^{t+1}z^t \Pr[\tilde{p}_t > 0].$$

However, note that

$$\Pr[\tilde{p}_{t+1} > 0 | p_t > 0] = \frac{1}{z}.$$

Multiplying both sides by $z^{t+1}\Pr[\tilde{p}_t > 0]$ (and using the fact that $p_t > 0$ is necessary for $p_{t+1} > 0$), we obtain

$$z^{t+1}\Pr[\tilde{p}_{t+1} > 0] = z^t\Pr[\tilde{p}_t > 0].$$

Since this is true for all t , $z^t\Pr[\tilde{p}_t > 0] = z\Pr[\tilde{p}_1 > 0] = 1$ for all t . Thus trader 1's ex ante expected profits are given by

$$\sum_{t=1}^{\infty} (-1)^{t+1}.$$

But, of course, the infinite sum $1 - 1 + 1 - 1 + 1 - \dots$ does not converge. Hence the expectation is not defined.⁹

This also clarifies why the ex ante expectation does not exist while the conditional expectation always exists. When trader 1 conditions on his initial signal, he truncates the sum at a finite date since he knows that he will not trade past that date. For any finite T , of course, $\sum_{t=1}^T (-1)^{t+1}$ is perfectly well-defined and, as a result, so is trader 1's expected profit. Difficulties arise only when we ask for the ex ante expectation and hence cannot truncate.

It is important to note that the fact that the ex ante expectation does not exist is entirely due to this phenomenon. This nonexistence of the expectation has nothing to do with how traders evaluate infinite sequences of trades – i.e., what the utility of a trader is if he and the other trader swap \$1 back and forth forever has nothing to do with this calculation. *Every* sequence of trades with positive probability is finite in length. It is only when we try to compute the expectation over these infinitely many finite-length sequences that problems arise. Furthermore, the nonexistence has nothing to do with poorly defined probability distributions. In particular, the traders know and agree on the ex ante probability distribution over the time path of prices. In fact, each trader's conditional expected profits are even common knowledge, since trader i 's conditional expected profits are independent of w_i .

It is also important to point out that the nonexistence of this expectation is a necessary property of the kind of stochastic process we are considering, not something driven by our particular approach to constructing equilibria. More precisely,

Theorem 3. *If $\{\tilde{p}_t\}$ is any stochastic process which is a submartingale and has $E(\tilde{p}_1) > 0$, then*

$$E\left[\sum_{t=1}^{\infty} (-1)^t \tilde{p}_t\right]$$

does not exist.

⁹ Put differently, there is no number, including ∞ and $-\infty$, that we can consistently say is what this sum equals.

Any stochastic process that would seem plausible as a bubble we could rationalize would surely be a submartingale and have $E(\tilde{p}_1) > 0$. Furthermore, certainly the simplest description of trading would have at least some agent trading a constant¹⁰ quantity over time, buying and selling in alternating periods. But then this agent's trading profits would be equal or proportional to $\sum_{t=1}^{\infty} (-1)^t \tilde{p}_t$. Hence Theorem 3 implies that ex ante expected trading profits do not exist. Consequently, requiring this expectation to exist is tantamount to ruling out precisely the most intuitively plausible simple bubbles. Of course, when optimality is based on ex ante considerations alone, one has no choice but to require this expectation to exist. When interim optimality is well-defined, though, it is not obvious that this requirement is reasonable.

The real question, in short, is what, if anything, the nonexistence of the ex ante expectation means in economic terms.¹¹ Is existence a technical nicety or an important restriction which should be imposed?

There is no obvious answer to this question, but we think a strong case can be made for the former view. The traders in the model "live" in the interim world¹² – that is, they make decisions after they observe their initial wealth. The ex ante world is a fictitious construction used to ensure that the traders' beliefs about one another are consistent. Without constructing the ex ante world, there need not be any clear connection between trader i 's belief about trader j and trader j 's strategy. Like the common prior assumption, requiring that this expectation exists is a way to ensure that the perception of trader i is consistent with the perception of trader j . However, unlike the common prior assumption, the "inconsistency" caused by the nonexistence of this expectation does not imply that the probability distribution over price paths is poorly defined. If the traders have different priors, they may well disagree about the probability distribution over the price path. Hence there is no obvious sense in which a model without common priors "predicts" a particular probability distribution. As noted above, the nonexistence of the ex ante expectation does not imply that the predictions of the model are unclear.

Furthermore, while strong philosophical arguments have been made for the common prior assumption (see especially Aumann (1987)), these arguments do not extend to requiring the ex ante expectation to exist. The argument for common priors, basically, is that if two individuals have all the same information, then they should agree completely on the probability of any given event. In our equilibria, since the traders have common priors, they do agree completely on what the conditional expectations are. If they had the same information, they would agree on every possible probability. They even agree on ex ante probabilities. Thus this kind

¹⁰ Alternatively, we can interpret \tilde{p}_t as the amount of money received or spent by the agent in period t , rather than the price. Under this interpretation, the expectation shown not to exist is clearly just expected profits.

¹¹ One implication it certainly has is that we must define equilibrium as we did above. Nash equilibrium requires that strategies be ex ante optimal, but, of course, we cannot evaluate this. This is why our definition of equilibrium only requires interim optimality. Interestingly, the usual definition of sequential equilibrium requires only sequential – that is, interim – rationality, not ex ante optimality.

¹² In fact, attempts to derive equilibrium notions from decision-theoretic principles have all used interim optimality. See, for example, Aumann (1987) or Aumann and Brandenburger (1991).

of philosophical argument does not seem to imply that we should require the ex ante expectation to exist.

It is worth noting that the fact that truncating the horizon can lead to bubbles has been long known. This is precisely the reason why bubbles are possible with myopic traders. As Tirole (1982) notes, it is easy to construct bubbles where traders are optimal across any two periods. However, in these examples, expected lifetime profits are not defined so that one cannot say whether the traders are dynamically optimal. In our model, precisely the same difficulty arises, except that it is ex ante expected profits, not lifetime expected profits that are not well defined. As we argued above, we believe that a strong case can be made for the view that ex ante optimality does not add anything economically meaningful to the requirement of interim rationality. One could certainly not say the same about dynamic versus myopic optimality. Another way of truncating the horizon without violating the rationality of agents is by the use of overlapping generations economies. In these cases, agents only live two periods, so optimality across any two periods is equivalent to dynamic optimality. However, unlike our model, these models require infinitely many agents.

Nalebuff (1989) in his discussion of the “envelope switching problem” gives a result strongly related to the nonexistence of ex ante expected profits.¹³ The problem, essentially, is the following. A number is drawn at random from the set $\{1, 2, 4, 8, \dots\}$. Let z denote the number drawn. We flip a coin. If it comes up heads, we put z dollars in envelope 1 and $2z$ dollars in envelope 2. Otherwise, we put z in envelope 2 and $2z$ in envelope 1. Player i is then given envelope i . He looks inside and is asked if he would like to trade with the other player. There are well-defined probability distributions such that each player would answer this question affirmatively no matter how much money is in his envelope. This seems quite unintuitive since the symmetry of the situation suggests that neither should “envy” the other. Nalebuff resolves this paradox by showing that the problem only arises if expected utility of “playing this game” is infinite.

Earlier, we concluded that ex ante expected utility in our model is undefined, not infinite. The contrast between this conclusion and Nalebuff’s is misleading. Recall that we defined utility as trading profits – an appropriate definition given the risk neutrality of the agents. Nalebuff defines utility as the utility of total final wealth. It is easy to show that the fact that ex ante expected profits are undefined implies that the ex ante expectation of initial wealth and the ex ante expectation of the sum of initial wealth and trading profits – the appropriate analog here to expected utility as computed by Nalebuff – are both infinite.

Theorem 4: *If we have an equilibrium which canonically rationalizes a very nice stochastic process $\{\tilde{p}_i\}$ with priors (F_1, F_2) , then letting π_i denote the equilibrium trading profits of trader i , we have*

$$E[w_i] = \infty, \quad i = 1, 2$$

and

$$E[w_i + \pi_i] = \infty, \quad i = 1, 2.$$

¹³ See also Brams and Kilgour (1991) and Brams, Kilgour, and Davis (1991).

Proof: Both expectations must exist (where we include the possibility that the expectation is ∞) because the expectation of any random variable which is always nonnegative must exist. By assumption, w_i is nonnegative with probability 1. Furthermore, since our traders are not allowed to spend more money than they have, $w_i + \pi_i$ is also nonnegative with probability 1. Suppose that either $E[w_i]$ or $E[w_i + \pi_i]$ (or both) is not infinite. Then $E[w_i + \pi_i] - E[w_i]$ is well defined (again, including the possibility that this difference is either ∞ or $-\infty$). But then $E[\pi_i]$ is well-defined as it must equal this difference. But we already have seen that $E[\pi_i]$ does not exist. \square

In other words, the equilibria we construct do rely quite crucially on the fact that our assumptions make utility unbounded. What is the appropriate view of this assumption? Nalebuff, for example, concludes there is only a paradox in the envelope switching problem under the “monstrous hypothesis” of infinite expected utility. Furthermore, many axiomatic derivations of expected utility imply that utility is bounded, precisely because unbounded utility leads to the possibility of noncomparabilities or other problems due to infinite or undefined expectations. (See Fishburn (1988) for more details.) Finally, it is well-known that unbounded utility can lead to paradoxes such as the St. Petersburg paradox.

On the other hand, risk neutrality hardly seems like an unusual assumption, even though it allows unbounded utility. Putting the point differently, the assumptions which allow bubbles are not unusual assumptions in the literature, even though some of the implications of these assumptions have not been fully appreciated. More importantly, all relevant expectations are well-defined in the interim world, which is, arguably, what the real world corresponds to. No “monstrous hypothesis” seems to lurk there.

In fact, it is not hard to show that we can even bound the utility functions with probability 1 and achieve the same result.¹⁴ More specifically, suppose that neither trader knows the other’s initial wealth or utility function. Fix an equilibrium as constructed for the proof of Theorem 2. In those equilibria, a trader’s initial wealth determines an upper bound on how long he will trade. Hence it determines an upper bound on the profits he might earn and a lower bound on the losses he might incur. Suppose that we change our assumptions to requiring that given this initial wealth, the trader’s utility function is linear in wealth between these two points, but is flat outside this range. Then precisely the same strategies still form an equilibrium, rationalizing the bubble. But now utility is bounded both from above and from below with a probability 1. On the other hand, for larger initial wealths, the upper bound on utility will have to be higher and the lower bound lower. In other words, it is interim utility that is bounded, not *ex ante*. Again, if one views *ex ante* utility as a fictional construct, the unboundedness of *ex ante* utility should not be troubling.

Before turning to a generalization of Theorem 2, we discuss the role played by our assumptions on price determination. Our emphasis on interim versus *ex ante* optimality should not lead the reader to infer that this is the only important difference between our model and Tirole’s. With this change alone, Tirole’s results,

¹⁴ We are grateful to the referee for suggesting this possibility.

we believe, still hold. The other important change is our assumption that agents are not price takers. If agents were price takers, uncertainty about initial wealth could not generate bubbles – or, at least, could not generate bubbles by providing privately known finite truncation dates. To see why, suppose we continue with all our earlier assumptions, except that we now assume that the market is competitive. More precisely, each trader views prices as a stochastic process which he has information about but cannot affect by his trading choices.

Let p_t denote the period t price and \hat{p}_{t+1}^i trader i 's expectation of the period $t + 1$ price. The price-taking assumption implies that if $\hat{p}_{t+1}^i > p_t$, then trader i will buy as much as possible in the current period – an infinite quantity, if allowed. To deal with such unbounded demands, constraints on holdings are often imposed. Regardless of whether maximum or minimum position constraints are imposed, certainly it must be true that if trader i sells some of the asset at period t , then it must be true that $\hat{p}_{t+1}^i \leq p_t$. If i expected the period $t + 1$ price to be strictly larger, he would be better off waiting and selling at $t + 1$. Similarly, if trader j purchases some of the asset in period t , it must be true that $\hat{p}_{t+1}^j \geq p_t$. It is not hard to use arguments akin to Tirole's (1982) proof of his Proposition 2 to show that this implies that no trade can occur unless for both traders, $\hat{p}_{t+1}^i = p_t$.¹⁵

Our use of finite truncation dates violates this requirement. We have used uncertainty about initial wealth to create these truncation dates, but the point is more general. Suppose that trader i 's truncation date is $t + 1$ – that is, in the context of the model, he will not have sufficient cash to purchase the asset at date $t + 1$. Hence the price will fall to zero at that point. Then trader i must rationally forecast this, so that \hat{p}_{t+1}^i would have to be zero. Hence trade at date t would be impossible unless $p_t = 0$. In short, our finite truncation device plus price-taking would imply that all prices were zero if trade ever occurs.

To see how this argument breaks down in our model, consider again the equilibrium constructed above. Suppose that trader 1's initial wealth is such that he will purchase the asset at date 2 if he successfully sells it at date 1. Then he knows that if $p_1 = z$, then $p_2 = z^2 > p_1$. Hence he sells in period 1 knowing that the period 2 price will be larger. He is still willing to sell in period 1 because if he doesn't, the bubble bursts in period 1 and he will never sell the asset for a strictly positive price.

More generally, it is clear that a trader may be willing to sell even though he expects the price to be higher tomorrow if not selling affects the price adversely. Of course, the effect a trader has on the price is razor sharp in the bargaining game we analyze. However, the magnitude of the effect seems less important than its direction, leading us to conjecture that similar results would be obtained in models where individual trader's actions affect the price in a more continuous fashion.

On the other hand, the direction of the effect may also seem unintuitive. Why should a refusal to sell – or, more generally, a reduction in the amount one sells – cause prices to fall, rather than rise? Why should we expect this property to carry over to other models of price determination? In any model of bubbles, a buyer

¹⁵ Roughly, this is the same proof as Tirole gives to show that ex ante expected profits must be zero. However, unlike profits, prices should always be nonnegative. Hence expected prices must always be well defined.

will purchase the asset because he believes it likely that he will be able to resell it later at a higher price. In our model, he will only be able to resell it for more money later if the initial wealths of other traders were sufficiently high. His only way of learning about the initial wealth of another trader is to observe his purchases of the asset – the more the other trader spends, the higher his initial wealth must have been. Consequently, if he does not observe the other traders purchasing enough of the asset, he will assign a lower probability to the bubble continuing and hence will be less willing to purchase the asset. Hence if one trader sells less of the asset, this reduces the confidence of all traders in the bubble continuing and so may well reduce prices. More generally, if the private information of each trader determines a finite truncation date, then when trade in the asset slows, traders may suspect that others are dropping out of the market. This will increase each trader's subjective probability that the bubble will burst soon, thus hastening the crash.¹⁶ In short, this effect is intuitive and seems generalizable to other models of price determination so long as private information creates privately known truncation dates.

It is very easy to show that the results generalize to certain alternative models of price determination. First, a large variety of bargaining rules would lead to the same outcome. For example, it is easy to show that the same results would be obtained if we assumed that only an agent possessing some of the asset could make an offer, with a potential buyer only able to accept or reject. Similarly, if the offers do not match exactly but are “compatible,” we could assume that a compromise proposal is adopted, as in Chatterjee and Samuelson's (1983) bargaining game or in the literature on *k*-double auctions (see, for example, Satterthwaite and Williams (1989)). Second, it is straightforward to show that the same results hold for a demand submission game as studied by Kyle (1989), Jackson (1991), and Peck, Shell, and Spear (1989). This game differs in that demand functions are submitted instead of offers. Agreement is determined by market clearing. To see why this approach would not be much different than what we have here, note that a trader could always submit a demand function where he only offers to trade a certain quantity at a certain price. Hence the kind of offers we consider could still be made, so that the demand functions essentially include the offers we consider. It should not be surprising that enriching the set of offers which can be made does not reduce the set of rationalizable processes. One advantage of the demand submission game is that it provides a relatively simple and plausible way to extend our framework to more than two traders.

4. Equilibria with sunspots

If we add sunspots – i.e., intrinsically irrelevant uncertainty – we can greatly expand the class of canonically rationalizable processes. Formally, we change the model described in Section 2 in the following way. At the beginning of each period, before offers are made, both traders observe the realization of a random variable which is independent of the initial wealths and has no direct effect on payoffs at all. Let $\{\tilde{s}_t\}$

¹⁶ This effect is consistent with the experimental findings of Smith, Suchanek, and Williams (1988) who find that volume tends to decrease just before the bursting of a bubble.

denote the stochastic process for the “signals.” Now a history for trader i at period t also includes the realizations of \tilde{s}_r for $r \leq t$. Hence i 's offer may depend on these past realizations.

Just as before, a stochastic process for signals, a pair of strategies, and a probability distribution on initial wealths determines a stochastic process for prices. We will say that $\{\tilde{p}_i\}$ is *s-rationalizable* if there exists $\{\tilde{s}_i\}$, $F = (F_1, F_2)$, and an equilibrium σ which induces the stochastic process $\{\tilde{p}_i\}$. We also define an analogous notion to canonical rationalizability, which we call canonical s-rationalizability. Loosely, in the equilibria which support such a process, the bubble can burst either because of the way traders react to sunspots or, as before, because one trader runs out of money. More formally, $\{\tilde{p}_i\}$ is *canonically s-rationalizable* if it is s-rationalizable with strategies satisfying the condition that given any realization of wealths and sunspots such that the bubble bursts at date t , *either* both traders refuse to trade at t , regardless of their wealths, *or* for one of the agents, say i , the offer made by j is not feasible for i .

The main result of this section is:

Theorem 5. *Every nice stochastic process $\{\tilde{p}_i\}$ is canonically s-rationalizable.*

As with Theorem 2, this result is proved by fixing a nice stochastic process and then constructing wealth distributions and sunspot distributions which support an equilibrium generating this process. The construction is very simple. The distribution of sunspots in a given period given a particular history is chosen to mimic the distribution of prices which the stochastic process specifies. The strategies are chosen so that the seller of the asset in any period always offers to sell at a price equal to that period's sunspot. The buyer for the period offers to buy at a price equal to the sunspot *if* he has sufficient wealth. Exactly as in the case without sunspots, the fact that each agent has finite wealth guarantees that there is a finite date after which he will no longer be able to afford to buy. Hence our finite truncation device works in essentially the same fashion as it did before.

It is also easy to see that because Theorem 3 covers a class of stochastic processes even broader than the nice processes, it and Theorem 4 also apply to this model. Hence for any priors which canonically s-rationalize a nice stochastic process, we must have $E[w_i] = \infty$ for $i = 1, 2$ and we must have $E[w_i + \pi_i] = \infty$ for $i = 1, 2$, where π_i is agent i 's equilibrium trading profits.

A natural question to ask is whether this model can rationalize stochastic processes outside the class of nice processes. The assumption of finite supports is almost certainly unnecessary, though the analysis is technically much more complex without it. The rest of condition (1), that prices are nonnegative, is also not necessary. Intuitively, ownership of units of the asset simply determines whose turn it is to receive transfers of money from the other traders. A negative price, then, just means that we are changing the order of turns.

The requirement that the price be a submartingale is used in our construction to guarantee that traders are willing to buy when we need them to. However, the assumption of a nonnegative price is similarly used to get traders to sell when we need them to. Hence the same intuition which suggests that negative prices are possible suggests that we can also relax this assumption to some extent.

Condition (4), that the probability of the bubble bursting at period t is bounded from below is also certainly not necessary. In fact, it is straightforward to replace this condition with weaker but less intuitive assumptions. On the other hand, one implication of this assumption – that prices collapse to fundamentals in finite time with probability 1 – may, perhaps in some weakened form, be necessary. It certainly is necessary given the way we use finite truncation dates for each trader to construct equilibrium bubbles.

Condition (5) is quite useful, but unnecessary. It enables us to construct wealth distributions so that the wealth constraints enter in a very simple way. We construct these distributions so that the support of F_1 is $\{0, \bar{\omega}_2, \bar{\omega}_4, \dots\}$ and the support of F_2 is $\{0, \bar{\omega}_1, \bar{\omega}_3, \dots\}$. Note that $\bar{\omega}_t$ is the maximum revenue accumulated by trading up to and including period t by the agent who sells at date t . Equivalently, it is the maximum amount paid in total up to and including t by the agent who buys at date t . Hence if the buyer at that date is endowed with wealth of at least $\bar{\omega}_t$, he can certainly afford to buy at that date. By condition (5), if he is endowed with wealth of $\bar{\omega}_{t-2}$ or less, he certainly cannot afford to purchase at date t , no matter what price path is realized. Hence (5) together with this construction of wealth distributions guarantees that an agent can afford to purchase at date t for one realized price path if and only if he can afford to do so for all possible price paths.

Finally, condition (3), that a price of zero is an absorbing state, may be necessary in a weakened form given our approach. The whole idea of canonical rationalizability is that the bubble bursts when one trader has insufficient wealth. Hence the date at which the bubble bursts reveals an upper bound on the initial wealth of one of the traders. As seen earlier, expected initial wealth must be infinite in our model. Hence we cannot construct a bubble using uncertainty about initial wealth if there is a finite upper bound on initial wealth. Therefore, we cannot construct an equilibrium like this in which a new bubble starts once the old one has burst. On the other hand, one can construct examples with more than two traders with more than one bubble. Also, if the uncertainty is about some variable other than initial wealth, equilibria with multiple bubbles may be more straightforward.

5. Conclusions

Many simple alterations of the model are possible. For example, since we can rationalize any strict submartingale satisfying certain conditions, it can be true that all traders expect strictly positive gains from trade. Interestingly, this suggests a role for an outside party who sets up these trades as a “broker.” If the broker’s fees per transaction are small enough, all traders still wish to trade and the broker makes strictly positive profits. Also, though it is more complex, the results can be qualitatively extended to the case of risk averse traders. Finally, there are certainly many alternative price-setting institutions which would generate similar equilibria.

It is also worth noting that while private information plays an important role, the information does not have to be about initial wealth. A simple, though perhaps less economically interesting, alternative is to suppose that each trader has a fixed length of “life” which is private information. Clearly, this could be used to generate

the same finite truncation device we exploit here, without having to assume the existence of cash constraints in trading.

To conclude, we have shown that bubbles are possible with rational traders who have common priors. We require two departures from Tirole's (1982) framework. First, our construction requires that (at least some) traders are not price takers. Second, the ability to construct rational equilibrium bubbles hinges, surprisingly enough, on the distinction between ex ante and interim optimality. Since the two criteria typically coincide, ex ante optimality has generally been required rather than what we see as the more reasonable requirement of interim optimality. As we show, the choice of optimality criteria has surprisingly important consequences.

Appendix

Proof of Theorem 1:

For a stochastic process $\{\tilde{p}_t\}$ satisfying (1) through (4),

$$\begin{aligned} E(\tilde{p}) &= \sum_{t=1}^{\infty} \sum_{p' \in \text{supp}(\tilde{p}_t), p' \neq 0} p' \Pr[\tilde{p}_t = p', \tilde{p}_{t+1} = 0] \\ &= \sum_{t=1}^{\infty} \sum_{p' \in \text{supp}(\tilde{p}_t), p' \neq 0} p' \Pr[\tilde{p}_t = p'] \Pr[\tilde{p}_{t+1} = 0 | p_t = p']. \end{aligned}$$

By (4), then,

$$E(\tilde{p}) \geq \sum_{t=1}^{\infty} \sum_{p' \in \text{supp}(\tilde{p}_t), p' \neq 0} p' \Pr[\tilde{p}_t = p'] \varepsilon = \sum_{t=1}^{\infty} \varepsilon E(\tilde{p}_t).$$

By the submartingale property, $E(\tilde{p}_t) \geq E(\tilde{p}_1)$ for all t . Also, (1) implies that $E(\tilde{p}_1) > 0$. Hence

$$E(\tilde{p}) \geq \sum_{t=1}^{\infty} \varepsilon E(\tilde{p}_1) = \infty. \quad \square$$

Proof of Theorem 2:

Fix a very nice stochastic process $\{\tilde{p}_t\}$. Recall that under the canonical strategies, if trade has occurred in every period $k \leq t$, then trade occurs at $t + 1$ iff the accumulated wealth of the buyer in period $t + 1$ is at least p_{t+1}^* . Let

$$w_t^* = \sum_{j=1}^t (-1)^{j+1} p_j^*, \quad t = 1, 2, \dots$$

Notice that w_t^* is simply the total amount of money trader 1 has received from trader 2 up to and including period t if the canonical strategies have been followed and if the bubble has not burst before t . Hence under the canonical strategies, the probability of trade at period t conditional on trade in all past periods is

$$\begin{cases} \Pr[w_2 \geq w_t^* | w_2 \geq w_{t-2}^*], & \text{if } t \text{ is odd;} \\ \Pr[w_1 \geq -w_t^* | w_1 \geq -w_{t-2}^*], & \text{otherwise.} \end{cases}$$

where w_{-1}^* and w_0^* are defined to be 0.

Any very nice stochastic process must have $p_{t+1}^* > p_t^*$. Hence $w_{2n+1}^* > 0$ and $-w_{2n}^* > 0$. Furthermore, note that (2) and (4) imply

$$w_{2n+1}^* - w_{2n-1}^* = p_{2n+1}^* - p_{2n}^* \geq p_{2n+1}^* - (1 - \varepsilon)p_{2n+1}^* = \varepsilon p_{2n+1}^* > \varepsilon p_{2n}^*.$$

Hence w_{2n+1}^* is strictly increasing in n and goes to infinity as $n \rightarrow \infty$. A similar argument applies to $-w_{2n}^*$.

Let

$$\varepsilon_t = \Pr[\tilde{p}_t = 0 | p_1 = p_1^*, \dots, p_{t-1} = p_{t-1}^*].$$

If the process is canonically rationalizable, there are distribution functions for wealths such that the probabilities generated by the canonical strategies match the sequence of ε_t 's. In other words, we must find priors satisfying

$$1 - \varepsilon_{2n-1} = \frac{\Pr[w_2 \geq w_{2n-1}^*]}{\Pr[w_2 \geq w_{2n-3}^*]}, \quad n = 1, 2, \dots$$

and

$$1 - \varepsilon_{2n} = \frac{\Pr[w_1 \geq -w_{2n}^*]}{\Pr[w_1 \geq -w_{2n-2}^*]}, \quad n = 1, 2, \dots$$

For ease of notation, redefine variables as follows. For $n = 1, 2, \dots$, let $H_{2n-1} = \Pr[w_2 \geq w_{2n-1}^*]$ and let $G_{2n} = \Pr[w_1 \geq -w_{2n}^*]$. Then we can write these equations as

$$H_{2n-1} = (1 - \varepsilon_{2n-1})H_{2n-3}, \quad n = 1, 2, \dots$$

and

$$G_{2n} = (1 - \varepsilon_{2n})G_{2n-2}, \quad n = 1, 2, \dots$$

This gives us a pair of difference equations. The initial value for the first difference equation is $H_{-1} = \Pr[w_2 \geq 0] = 1$. Similarly, the initial value for the second equation, G_0 , must be 1. Solving the difference equations with these initial conditions yields:

$$H_{2n-1} = \prod_{k=1}^n (1 - \varepsilon_{2k-1}), \quad n = 1, 2, \dots \tag{20a}$$

and

$$G_{2n} = \prod_{k=1}^n (1 - \varepsilon_{2k}), \quad n = 1, 2, \dots \tag{20b}$$

Hence if there are priors such that this process is canonically rationalized, they must satisfy (20a) and (20b). Consider (20a) first. Since $\varepsilon_t \in (0, 1)$ for all t ,¹⁷ $H_{2n-1} \in (0, 1)$ for all n and is strictly decreasing in n . Furthermore, condition (4) implies $1 - \varepsilon_t \leq 1 - \varepsilon$ for all t , so that

$$\prod_{k=1}^n (1 - \varepsilon_{2k-1}) \leq (1 - \varepsilon)^n.$$

Hence $H_{2n-1} \rightarrow 0$ as $n \rightarrow \infty$.

Define F_2 by $F_2(w) = 1 - H_{2n-1}$ for all $w \in [w_{2n-1}^*, w_{2n+1}^*)$, $n = 0, 1, \dots$. As noted above, $H_{2n-1} \in (0, 1)$, so that $F_2(w) \in [0, 1]$ for all w . Since H_{2n-1} is strictly decreasing and w_{2n-1}^* strictly increasing in n , $F_2(w)$ is weakly increasing in w , as a distribution function must be. Obviously, F_2 is continuous from the right. Finally, the fact that $H_{2n-1} \rightarrow 0$ and $w_{2n-1}^* \rightarrow \infty$ as $n \rightarrow \infty$ implies that $F_2(w) \rightarrow 1$ as $w \rightarrow \infty$. Hence this is a proper distribution function satisfying (10). The analogous arguments show that a step function F_1 constructed using the G_{2n} series is a distribution function satisfying (10).

By construction, then, if the canonical strategies for $\{\tilde{p}_t\}$ are followed when F_1 and F_2 are the priors constructed above, the stochastic process generated will be $\{\tilde{p}_t\}$. Hence we only need to show that the canonical strategies, together with some beliefs, form an equilibrium given these priors.

Fix any beliefs consistent with Bayes' Rule and these strategies. We claim that the canonical strategies plus these beliefs are an equilibrium. First, consider a history for which the sequence of past prices has zero probability under the strategies above or for which the traders' holdings are inconsistent with these strategies. According to the canonical strategies, each trader refuses to trade – i.e., submits an offer with his position in the asset unchanged – in all subsequent periods. Given this behavior by trader j , trader i cannot trade regardless of what he does. Hence each trader is choosing a best reply.

¹⁷ We must have $\varepsilon_t > 0$ by (4) and $\varepsilon_t < 1$ by (2).

Now consider any other history. If the bubble has already burst, an analogous argument to the above implies that following the canonical strategies is optimal. So suppose the bubble has not burst yet. Recall that each trader knows his initial wealth and knows that, given that the opponent follows the canonical strategies, he will be unable to trade past the period in which w_{2n+1}^* or $-w_{2n}^*$ (depending on the trader) exceeds his initial wealth. Since initial wealth is finite with probability 1 and each of these sequences goes to infinity, there is necessarily a finite date past which the trader knows he will not trade. Hence the conditional expected payoff to following the canonical strategy or any deviation from it is well-defined.

It is not hard to see that if i is supposed to sell in the current period, it is optimal for him to offer to sell at p_t^* . Otherwise, his continuation profits are certainly zero. If he does so, his expected continuation profit is at least $p_t^*(1 - \varepsilon_i) > 0$ as he can always refuse to trade in every future date.

If i is supposed to buy in period t and his accumulated wealth as of period t is less than p_t^* , he cannot purchase the asset at this price. Hence given that the other trader follows his equilibrium strategy, i may as well follow the canonical strategy. Suppose i 's accumulated wealth is at least p_t^* then. Then if he offers to purchase the asset at p_t^* , offers to sell at p_{t+1}^* the next period, and refuses to trade thereafter, his expected continuation profits are $(1 - \varepsilon_{t+1})p_{t+1}^* - p_t^* \geq 0$. Hence certainly offering to buy at price p_t^* is weakly better than not doing so. Therefore, following the canonical strategy is optimal. \square

Proof of Theorem 3:

Suppose not. Let

$$\theta = E \left[\sum_{t=1}^{\infty} (-1)^t \tilde{p}_t \right].$$

Clearly, since the expectation does exist, it must be true that

$$\sum_{t=1}^{\infty} (-1)^t \tilde{p}_t$$

exists with probability one, so that

$$\theta = \sum_{t=1}^{\infty} (-1)^t E(\tilde{p}_t) = \sum_{n=1}^{\infty} [E(\tilde{p}_{2n}) - E(\tilde{p}_{2n-1})].$$

By the submartingale assumption, $E(\tilde{p}_{2n}) \geq E(\tilde{p}_{2n-1})$. Hence $\theta \geq 0$.

Note also, though, that

$$\theta = -E(\tilde{p}_1) + \sum_{n=1}^{\infty} [E(\tilde{p}_{2n}) - E(\tilde{p}_{2n+1})].$$

Using the submartingale assumption again, $E(\tilde{p}_{2n}) \leq E(\tilde{p}_{2n+1})$, so $\theta \leq -E(\tilde{p}_1)$. But $E(\tilde{p}_1) > 0$, so $\theta < 0$, a contradiction. \square

Proof of Theorem 5:

Fix any nice stochastic process $\{\tilde{p}_t\}$. We first construct wealth and sunspot distributions. These two processes are independent of one another, so we define them one at a time. First, let

$$\Pr[w_1 = 0] = \varepsilon$$

and

$$\Pr[w_1 = \bar{\omega}_{2n}] = \varepsilon(1 - \varepsilon)^n, \quad n = 1, 2, \dots$$

Note that

$$\Pr[w_1 = 0] + \sum_{n=1}^{\infty} \Pr[w_1 = \bar{\omega}_{2n}] = \varepsilon + \varepsilon(1 - \varepsilon) \sum_{n=0}^{\infty} (1 - \varepsilon)^n = 1.$$

Similarly, let

$$\Pr[w_2 = 0] = \varepsilon$$

and

$$\Pr[w_2 = \bar{\omega}_{2n-1}] = \varepsilon(1 - \varepsilon)^n, \quad n = 1, 2, \dots$$

Again, the probabilities sum to 1, so (A.1) and (A.2) are satisfied.

Construct distributions for the sunspots as follows. The support of \tilde{s}_t is the same as the support of \tilde{p}_t . For any $s \neq 0$,

$$\Pr[\tilde{s}_t = s | s_1, \dots, s_{t-1}] = \frac{1}{1 - \varepsilon} \Pr[\tilde{p}_t = s | p_1 = s_1, \dots, p_{t-1} = s_{t-1}].$$

Finally,

$$\Pr[\tilde{s}_t = 0 | s_1, \dots, s_{t-1}] = 1 - \sum_{s \in \text{supp}(\tilde{s}_t), s \neq 0} \Pr[\tilde{s}_t = s | s_1, \dots, s_{t-1}].$$

This expression is nonnegative as long as

$$\sum_{s \in \text{supp}(\tilde{s}_t), s \neq 0} \Pr[\tilde{s}_t = s | s_1, \dots, s_{t-1}] \leq 1$$

or

$$\frac{1}{1 - \varepsilon} \sum_{s \in \text{supp}(\tilde{p}_t), s \neq 0} \Pr[\tilde{p}_t = s | p_1 = s_1, \dots, p_{t-1} = s_{t-1}] \leq 1.$$

Rearranging,

$$\sum_{s \in \text{supp}(\tilde{p}_t), s \neq 0} \Pr[\tilde{p}_t = s | p_1 = s_1, \dots, p_{t-1} = s_{t-1}] \leq 1 - \varepsilon$$

or

$$\Pr[\tilde{p}_t = 0 | p_1 = s_1, \dots, p_{t-1} = s_{t-1}] \geq \varepsilon$$

which holds by condition (4). Hence this generates a well-defined stochastic process for the sunspots.

Construct the strategies as follows. Let $h_s = (s_1, \dots, s_t)$ denote a sequence of sunspot realizations with positive probability under the process defined above. Given this sequence, at period t , if one unit of the asset has been exchanged in each past period r at a price equal to s_r , then the agent who owns the asset offers to sell at price s_t . The agent with none of the asset offers to buy at this price if his accumulated wealth to this point is at least s_t . That is, if $w_j + \omega_{t-1}(h_{t-1}^s) \geq s_t$ where

$$\omega_t(h_t^s) = \sum_{j=0}^{t-1} (-1)^j s_{t-j}.$$

Otherwise, he refuses to buy. If there has ever been a period of no trade in the past or a period in which the price differed from that period's sunspot realization, then both traders refuse to trade.

It is easy to see that these strategies have the structure required for canonical s-rationalizability. Trade always occurs until either $s_t = 0$ – in which case neither trader is willing to trade, no matter what his wealth is – or one trader cannot afford to purchase the asset at a price the other is offering to sell for. All that remains, then, is to show that (1) if these strategies are followed, the stochastic process of prices generated will be the process we began with and (2) these strategies constitute an equilibrium.

To show the former, notice that under these strategies, given a history of sunspot realizations $h_{t-1}^s = (s_1, \dots, s_{t-1})$, the probability that the period t price is p for $p \neq 0$ is

$$\Pr[\tilde{s}_t = p | s_1, \dots, s_{t-1}] \Pr[w_j \geq p - \omega_{t-1}(h_{t-1}^s) | w_j \geq \omega_{t-2}(h_{t-2}^s)]$$

where $h_{t-2}^s = (s_1, \dots, s_{t-2})$ and j is the trader who buys at period t – that is, $j = 1$ if t is even and $j = 2$ if t is odd. Note that

$$p - \omega_{t-1}(h_{t-1}^s) = \omega_t(s_1, \dots, s_{t-1}, p).$$

Let $h_t^s(p) = (s_1, \dots, s_{t-1}, p)$. Given that the strategies above are followed, $h_t^s(p)$ must be a possible sequence of prices according to the stochastic process $\{\tilde{p}_t\}$. It is easy to see that $\omega_t(h_t^s(p))$ is precisely the realization of $\tilde{\omega}_t$ associated with this sequence of prices. Hence

$$\omega_t(h_t^s(p)) \in [\omega_r, \bar{\omega}_t]$$

and analogously for $\omega_{t-2}(h_{t-2}^s)$. By condition (5) of the definition of a nice stochastic process, $\omega_r > \bar{\omega}_{r-2}$ for all r . Hence $\bar{\omega}_r > \bar{\omega}_{r-2}$ for all r . Given the assumed distribution on w_j , then, $w_j \geq \omega_t(h_t^s(p))$ if and only if

$w_j \geq \bar{\omega}_r$. Similarly, $w_j \geq \omega_{t-2}(h_{t-2}^s)$ if and only if $w_j \geq \bar{\omega}_{t-2}$. Hence

$$\begin{aligned} \Pr[w_j \geq p - \omega_{t-1}(h_{t-1}^s) | w_j \geq \omega_{t-2}(h_{t-2}^s)] &= \Pr[w_j \geq \bar{\omega}_t | w_j \geq \bar{\omega}_{t-2}] \\ &= \frac{\Pr[w_j \geq \bar{\omega}_t]}{\Pr[w_j \geq \bar{\omega}_{t-2}]} = 1 - \varepsilon. \end{aligned}$$

So the probability of price p is $\Pr[\bar{p}_t = p | p_1 = s_1, \dots, p_{t-1} = s_{t-1}]$, just as it should be. Clearly, since the probability of any $p \neq 0$ is matched appropriately, the probability of $p_t = 0$ is matched.

The proof that these strategies form an equilibrium almost exactly parallels the analogous proof for Theorem 2 and so is omitted. \square

References

- Allen, F., Gorton, G.: Churning bubbles. *Rev. Econ. Stud.* **60**, 813–836 (1993)
- Allen, F., Morris, S., Postlewaite, A.: Finite bubbles with short sale constraints and asymmetric information. *J. Econ. Theory* **61**, 206–229 (1993)
- Aumann, R.: Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* **55**, 1–18 (1987)
- Aumann, R., Brandenburger, A.: Epistemic conditions for Nash equilibrium. Harvard Business School working paper, 1991
- Bertocchi, G.: Bubbles and inefficiencies. Brown University working paper, 1988
- Blanchard, O.: Speculative bubbles, crashes, and rational expectations. *Econ. Lett.* **3**, 387–389 (1979)
- Brams, S., Kilgour, D.: The box problem: To switch or not to switch. New York University, Department of Politics, working paper, 1991
- Brams, S., Kilgour, D., Davis, M.: Unraveling in games of sharing and exchange. New York University, Department of Politics, working paper, 1991
- Chatterjee, K., Samuelson, W.: Bargaining under incomplete information. *Operat. Res.* **31**, 835–851 (1983)
- Dow, J., Madrigal, V., Werlang, S.: Preferences, common knowledge, and speculative trade. Working paper, 1990
- Fishburn, P.: Nonlinear preference and utility theory. Baltimore: Johns Hopkins University Press 1988
- Glosten, L., Milgrom, P.: Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *J. Financ. Econ.* **14**, 71–100 (1985)
- Jackson, M.: Equilibrium, price formation, and the value of private information. *Rev. Financ. Stud.* **4**, 1–16 (1991)
- Jackson, M., Peck, J.: Speculation and price fluctuations with private, extrinsic signals. *J. Econ. Theory* **55**, 274–295 (1991)
- Kreps, D., Wilson, R.: Sequential equilibria. *Econometrica* **50**, 863–894 (1982)
- Kyle, A.: Continuous auctions and insider trading. *Econometrica* **53**, 1313–1335 (1985)
- Kyle, A.: Informed speculation with imperfect competition. *Rev. Econ. Stud.* **56**, 317–355 (1989)
- Milgrom, P., Stokey, N.: Information, trade, and common knowledge. *J. Econ. Theory* **26**, 17–27 (1982)
- Nalebuff, B.: Puzzles: The other person's envelope is always greener. *J. Econ. Perspect.* **3**, 171–181 (1989)
- Nash, J.: Two-person cooperative games. *Econometrica* **21**, 128–140 (1953)
- Neeman, Z.: Common beliefs and the existence of speculative trade. Northwestern University working paper, 1993
- Peck, J., Shell, K., Spear, S.: The market game: Existence and structure of equilibria. Cornell University working paper, 1989
- Royden, H.: Real Analysis. New York: Macmillan Publishing Co. 1968
- Rubinstein, A.: Comments on the interpretation of game theory. *Econometrica* **59**, 909–924 (1991)
- Satterthwaite, M., Williams, S.: Bilateral trade with the sealed bid k -double auction: existence and efficiency. *J. Econ. Theory* **48**, 107–133 (1989)
- Sebenius, J., Geanakoplos, J.: Don't bet on it: Contingent agreements and asymmetric information. *J. Am. Stat. Assoc.* **78**, 424–426 (1983)

- Smith, V., Suchanek, G., Williams, A.: Bubbles, crashes, and endogeneous expectations in experimental spot asset markets. *Econometrica* **56**, 1119–1151 (1988)
- Spear, S.: Are sunspots necessary? *J. Polit. Econ.* **97**, 965–973 (1989)
- Tirole, J.: On the possibility of speculation under rational expectations. *Econometrica* **50**, 1163–1181 (1982)
- Tirole, J.: Asset bubbles and overlapping generations. *Econometrica* **53**, 1071–1100 (1985) (corrected version: **53**, 1499–1528 (1985))
- Weil, P.: Confidence and the real value of money in an overlapping generations economy. *Q. J. Econ.* **102**, 1–22 (1987)