Queen's Economics Department Working Paper No. 851

Ex Ante versus Interim rationality and the existence of bubbles

Sugato Bhattacharyya Barton L. Lipman

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

4-1992

Ex Ante versus Interim Rationality and the Existence of Bubbles

Sugato Bhattacharyya
Carnegie Mellon University

Barton L. Lipman
Queen's University

Queen's University Discussion Paper #851

Current Draft
April 1992

We wish to thank Abhijit Banerjee, James Dow, Itzhak Gilboa, Fallaw Sowell, Sanjay Srivastava, Stan Zin and 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 for helpful comments.

ABSTRACT

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 even though it is common knowledge that the asset has no fundamental value at all. In the equilibria we construct, agents purchase the asset at successively higher prices (in expectation) until the bubble "bursts" and no subsequent trade occurs. In equilibrium, each trader has a finite "truncation date" and the date at which the bubble bursts is a function of these. Since no trader knows everyone's truncation date, none knows when the bubble will burst. As we show, these random truncations can arise from extrinsic uncertainty (*i.e.*, sunspots) or intrinsic uncertainty (such as uncertainty regarding the initial wealth of other traders). There are two key differences between our model and Tirole's which enable us to use this device to construct equilibrium bubbles. First, Tirole requires *ex ante* optimality, while we only require every trader's strategy to be optimal conditional on his information (specifically, on his truncation date) — *i.e.*, interim optimal. Since each trader knows his information before he actually trades, this would seem to be the relevant definition of optimality. Second, Tirole considers rational expectations equilibria, while we analyze a demand submission game.

I. 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 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], or Jackson and Peck [1991]. Allen and Gorton [1988] showed that bubbles can occur because of the preferences that agency problems induce for portfolio managers. Also, Dow, Madrigal, and Werlang [1990] have

---

[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.' "

shown that Milgrom and Stokey's no–trade results do not hold in general without expected utility preferences. Finally, Allen and Postlewaite [1991] 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.

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 (in expectation) until the bubble "bursts" and no subsequent trade occurs. In equilibrium, each trader has a finite "truncation date" and the date at which the bubble bursts is a function of these. Since no trader knows everyone's truncation date, none knows when the bubble will burst. As we show, these random truncations can arise from extrinsic uncertainty (*i.e.*, sunspots) or intrinsic uncertainty (such as uncertainty regarding the initial wealth of other traders). Either way, the agents are willing to trade because no one knows in advance exactly when the bubble will burst.

There are four differences between our model and Tirole's [1982] model, two important and two 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 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.

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 information (which includes his equilibrium truncation date) — that is, each is interim optimal. (Each is also sequentially rational in the sense of Kreps and Wilson [1982].) Since each trader knows his initial information 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.

The second important difference is the modeling of price determination: Tirole uses rational expectations equilibria, while we analyze a demand submission game. As we explain in Section III, our approach to constructing equilibrium bubbles does not work with rational expectations equilibrium, not because of rationality considerations, but because of price–taking. 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. In Section II, we explain the model with extrinsic uncertainty. In Section III, 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. In Section IV, we show that similar results can be generated with uncertainty about intrinsically relevant variables. Since these equilibria are

---

2 As we explain in Section IV, 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.

more complex, we do this for the special case of two traders. We offer some concluding remarks in Section V. All proofs are contained in the Appendix.

II. The Model.

There is a finite set of traders, $i = 1, \ldots, I$, all risk neutral. The two goods in the economy are shares of an asset and "money" or "wealth." Trader i's initial endowment is $e_i \geq 0$ shares of the asset. Until Section IV, we assume that traders can borrow and lend costlessly and so do not need to specify i's initial endowment of wealth. Let $\bar{x} = \sum_i e_i$ denote the aggregate endowment of the asset, where we assume $\bar{x} > 0$. It is common knowledge that the asset pays no dividends ever. Trade may occur at dates, $t = 1, 2, \ldots$. One can interpret the model as taking place in continuous but finite time (as in Allen and Gorton [1988]) 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 lim inf if the limit does not exist[3]) 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

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

where $x_0^i \equiv e_i$. 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.

In each period t, each trader i privately observes a signal. The signal

[3] Using the lim inf guarantees that the traders are "pessimistic." It is easy to construct infinite bubbles if both traders use the lim sup since each would view exchanging a dollar back and forth forever favorably. With the lim inf, this problem does not arise.

1982, pp. 863–894.

Kyle, A., "Continuous Auctions and Insider Trading," ***Econometrica***, 53, November 1985, pp. 1313–1335.

Kyle, A., "Informed Speculation with Imperfect Competition," ***Review of Economic Studies***, 56, July 1989, pp. 317–355.

Milgrom, P., and N. Stokey, "Information, Trade, and Common Knowledge," ***Journal of Economic Theory***, 26, 1982, pp. 17–27.

Nalebuff, B., "Puzzles: The Other Person's Envelope is Always Greener," ***Journal of Economic Perspectives***, 3, Winter 1989, pp. 171–181.

Peck, J., K. Shell, and S. Spear, "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, July 1991, pp. 909–924.

Sebenius, J., and J. Geanakoplos, "Don't Bet on It: Contingent Agreements and Asymmetric Information," ***Journal of the American Statistical Association***, 78, 1983, pp. 424–426.

Spear, S., "Are Sunspots Necessary?," ***Journal of Political Economy***, August 1989, pp. 965–973.

Tirole, J., "On the Possibility of Speculation under Rational Expectations," ***Econometrica***, September 1982, pp. 1163–1181.

Tirole, J., "Asset Bubbles and Overlapping Generations," ***Econometrica***, September 1985, pp. 1071–1100 (corrections: November 1985, pp. 1499–1528).

Weil, P., "Confidence and the Real Value of Money in an Overlapping Generations Economy," ***Quarterly Journal of Economics***, February 1987, pp. 1–22.

BIBLIOGRAPHY

Allen, F., and G. Gorton, "Rational Finite Bubbles," Wharton School, University of Pennsylvania working paper, 1988.

Allen, F., and A. Postlewaite, "Rational Expectations and Stock Market Bubbles," University of Pennsylvania working paper, 1991.

Aumann, R., "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, January 1987, pp. 1–18.

Aumann, R., and A. Brandenburger, "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," *Economics Letters*, 1979. pp. 387–389.

Brams, S. and D. Kilgour, "The Box Problem: To Switch or Not to Switch," New York University, Department of Politics, working paper, 1991.

Brams, S., D. Kilgour, and M. Davis, "Unraveling in Games of Sharing and Exchange," New York University, Department of Politics, working paper, 1991.

Dow, J., V. Madrigal, and S. Werlang, "Preferences, Common Knowledge, and Speculative Trade," working paper, 1990.

Fishburn, P., *Nonlinear Preference and Utility Theory*, Baltimore: Johns Hopkins University Press, 1988.

Glosten, L., and P. Milgrom, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 1985, pp. 71–100.

Jackson, M., "Equilibrium, Price Formation, and the Value of Private Information," *Review of Financial Studies*, 4, Number 1, 1991, pp. 1–16.

Jackson, M., and J. Peck, "Speculation and Price Fluctuations with Private, Extrinsic Signals," *Journal of Economic Theory*, 55, December 1991, pp. 274–295.

Kreps, D., and R. Wilson, "Sequential Equilibria," *Econometrica*, 50, July

is the realization of a random variable $\tilde{s}_t^i$. These signals are observed *before* any trading decisions at date t. We use s_t^i to denote a generic realization of $\tilde{s}_t^i$. (This notational device of removing a tilde from a random variable to denote a generic realization will be used throughout.) For simplicity, we also assume that each trader observes a signal $\tilde{s}_0^i$ before date 1. (The signals $\tilde{s}_0^i$ and $\tilde{s}_1^i$ can be joined together at the cost of more complex notation.) Let $\tilde{s}_t = (\tilde{s}_t^1, \ldots, \tilde{s}_t^I)$. The stochastic process $\{\tilde{s}_t\}_{t=0}^{\infty}$ is defined on a probability space $(\Theta, \mathcal{T}, \mu)$ which is known to all agents. In particular, all agents use the same prior μ in forming beliefs about the signals received by others. For notational brevity, we let $\mathcal{S}$ denote the stochastic process $\{\tilde{s}_t\}$ and the underlying probability space $(\Theta, \mathcal{T}, \mu)$. For each i and t, let S_t^i denote the support of the random variable $(\tilde{s}_0^i, \ldots, \tilde{s}_t^i)$. Since there is no "intrinsically relevant" uncertainty (*i.e.*, uncertainty about preferences, endowments, asset returns, etc.), these signals are sunspots.

The sequence of trades and prices is determined by a demand submission game similar to that studied by Kyle [1989], Jackson [1991], and Peck, Shell, and Spear [1989]. At each date t, all traders submit demand functions $y_t^i : \mathbf{R}_+ \to \mathbf{R}$, where $y_t^i(p)$ is interpreted as the quantity of the asset trader i demands to hold at price p. Note that $y_t^i(p)$ may be negative, so short sales are allowed. If there is a unique p such that $\sum_i y_t^i(p) = \bar{x}$, then trade takes place at that price. In this case, $p_t = p$ and trader i's new position is $x_t^i = y_t^i(p)$. Our results do not depend on what happens if there is more than one such p. Because it economizes on notation, we assume that if there are multiple such p but only one which involves exchange (*i.e.*, changes in traders' holdings), this p is chosen. If there are multiple p which involve exchange, we assume that no trade takes place, so that $x_t^i = x_{t-1}^i$. Since there is no trade in this event, no trader's utility is affected by the price so we adopt the convention that $p_t = 0$ in this case. Finally, if no market clearing price exists, no trade occurs, so that $x_t^i = x_{t-1}^i$ and $p_t = 0$. Let Y denote the set of functions $y : \mathbf{R}_+ \to \mathbf{R}$ — that is, the set of demand functions a trader could submit.

A *history for trader i* summarizes everything i knows at a particular point in the game. Hence it lists all his past signals, the demand functions he has submitted, his holdings of the asset at each past date, and the sequence of past prices. More formally, for $t \geq 1$,

$$h_t^i = (s_0^i, \ldots, s_t^i, y_1^i, \ldots, y_{t-1}^i,$$
$$x_1^i, \ldots, x_{t-1}^i, p_1, \ldots, p_{t-1}) \in S_t^i \times Y^{t-1} \times \mathbf{R}^{t-1} \times \mathbf{R}_+^{t-1}$$

is a t length history for i if for all $k < t$, $y_k^i(p_k) = x_k^i$. In other words, such a sequence is a history if it is 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 demand functions he submits or his position in the asset over time, while prices are observed by all traders. Our results would not be changed if all trades and demand submissions were commonly observed or if some summary statistics of trades (such as volume) were commonly observed.

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

Each trader has beliefs about the histories of the other traders. Formally, let $H_t^{\sim i} = \prod_{j \neq i} H_t^j$ and let $\Delta(H_t^{\sim i})$ denote the set of probability distributions defined on the Borel sets of $H_t^{\sim i}$. A *belief for trader i* is a sequence of functions $\delta^i = (\delta_1^i, \delta_2^i, \ldots)$ satisfying the following conditions. First, $\delta_t^i : H_t^i \to \Delta(H_t^{\sim i})$. That is, $\delta_t^i(h_t^i)$ is a probability distribution over $H_t^{\sim i}$ as a function of i's history. Second, if $h_t^{\sim i} = (h_t^1, \ldots, h_t^{i-1}, h_t^{i+1}, \ldots, h_t^I)$ is in the support of $\delta_t^i(h_t^i)$, then the sequence of prices in each h_t^j exactly matches the sequence of prices in h_t^i. Finally, if $h_t^{\sim i}$ is in the support of $\delta_t^i(h_t^i)$, then for every $k \leq t-1$, $\bar{x} - \sum_{j \neq i} x_k^j = x_k^i$. In other words, $\delta_t^i(h_t^i)$ can only give positive probability to histories for the other traders which are consistent with the prices i has observed, the trades i has made, and market clearing. Let $\delta = (\delta^1, \ldots, \delta^I)$.

erate a very nice stochastic process $\{\tilde{p}_t\}$. Clearly,

$$E(w_2) \geq \sum_{n=0}^{\infty} x_{2n-1}\{\Pr[w_2 \geq x_{2n-1}] - \Pr[w_2 \geq x_{2n+1}]\}$$

since the right–hand side computes the expectation as if all probability in the interval $[x_{2n-1}, x_{2n+1})$ were concentrated at the minimum of the interval. As shown in the proof of Theorem 3, $\Pr[w_2 \geq x_{2n-1}]$ must solve equation (20a). Hence

$$\begin{aligned} x_{2n-1}\{\Pr[w_2 \geq x_{2n-1}] - \Pr[w_2 \geq x_{2n+1}]\} &= x_{2n-1} \prod_{k=0}^{n-1}(1-\epsilon_{2k+1})\epsilon_{2n+1} \\ &\geq \epsilon x_{2n-1} \prod_{k=0}^{n-1}(1-\epsilon_{2k+1}). \end{aligned}$$

It is easy to show that $x_{2n-1} \geq \epsilon p_{2n-1}^*$. Also,

$$E[\tilde{p}_{2n-1}] = p_{2n-1}^* \prod_{t=1}^{2n-1}(1-\epsilon_t) \geq E[\tilde{p}_1].$$

where the inequality is from the submartingale property. Hence

$$\begin{aligned} x_{2n-1}\{\Pr[w_2 \geq x_{2n-1}] - \Pr[w_2 \geq x_{2n+1}]\} &\geq \epsilon^2 E[\tilde{p}_1] \frac{\prod_{k=0}^{n-1}(1-\epsilon_{2k+1})}{\prod_{t=1}^{2n-1}(1-\epsilon_t)} \\ &= \frac{\epsilon^2 E[\tilde{p}_1]}{\prod_{k=1}^{n-1}(1-\epsilon_{2k})}. \end{aligned}$$

As shown in the proof of Theorem 3, $\prod_{k=1}^{n-1}(1-\epsilon_{2k}) \to 0$ as $n \to \infty$. Hence

$$\lim_{n \to \infty} x_{2n-1}\{\Pr[w_2 \geq x_{2n-1}] - \Pr[w_2 \geq x_{2n+1}]\} = \infty,$$

implying $E(w_2) = \infty$. An analogous argument shows $E(w_1) = \infty$ as well. ∎

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 x_{2n+1} or $-x_{2n}$ (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, just as in the proof of Theorem 2, 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 submit a demand function with $y_t^i(p_t^*) = 0$. Otherwise, his continuation profits are certainly zero. If he does submit such a demand function, his expected continuation profit is at least $p_t^*(1-\epsilon_t) > 0$ as he can always submit a demand function of 0 at every p 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 demands the asset at p_t^* and submits a demand function with $y_{t+1}^i(p_{t+1}^*) = 0$ in the next period and zero demands thereafter, his expected continuation profits are $(1-\epsilon_{t+1})p_{t+1}^* - p_t^* \geq 0$. Hence certainly submitting a demand with $y_t^i(p_t^*) = 1$ is weakly better than not doing so. Therefore, following the canonical strategy is optimal. ▌

Proof of Theorem 4.

Suppose F_1 and F_2 are priors such that the canonical strategies gen-

Any strategies σ together with the stochastic process for signals $\mathcal{S}$ generates a stochastic process on signals, prices, and trades. This, in turn, generates a stochastic process on private histories for each trader. Let $\psi_t^i(y_1^i, \ldots, y_{t-1}^i, \sigma^{\sim i}, \mathcal{S})$ denote the induced distribution on $\prod_i H_t^i$ when the sequence of demand submissions by i is $y_1^i, \ldots, y_{t-1}^i$ and the strategies of the other traders are $\sigma^{\sim i}$. Let $\mathcal{H}_t^i(y_1^i, \ldots, y_{t-1}^i, \sigma^{\sim i}, \mathcal{S})$ denote the support of ψ_t^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 $\bar{\sigma}^{\sim i}$ given $\bar{\delta}_t^i(h_t^i)$. Second, for every $h_t^i \in \mathcal{H}_t^i(y_1^i, \ldots, y_{t-1}^i, \bar{\sigma}^{\sim i}, \mathcal{S})$ such that $y_1^i, \ldots, y_{t-1}^i$ is the sequence of demand submissions in h_t^i, $\bar{\delta}_t^i(h_t^i)$ is generated from $\psi_t^i(y_1^i, \ldots, y_{t-1}^i, \bar{\sigma}^{\sim i}, \mathcal{S})$ 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.

III. Equilibria with Sunspots.

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 note that any strategies σ together with the stochastic processes for signals $\mathcal{S}$ induces a stochastic process $\{\tilde{p}_t\}$ on prices. Let $\mathcal{P}(\sigma, \mathcal{S})$ denote this mapping. We say that a given stochastic process on prices, say $\{\tilde{p}_t\}$, is ***rationalizable*** if there exists $\mathcal{S}$ and an equilibrium (σ, δ) such that $\mathcal{P}(\sigma, \mathcal{S}) = \{\tilde{p}_t\}$.

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 $\mathbf{R}_+$ where $\text{supp}(\tilde{p}_1) \neq \{0\}$.

(2) $\{\tilde{p}_t\}$ is a submartingale. In other words, for every $(p_1, \ldots, p_{t-1}) \in \text{supp}(\tilde{p}_1, \ldots, \tilde{p}_{t-1})$,

$$\text{E}[\tilde{p}_t \mid p_1, \ldots, p_{t-1}] \geq p_{t-1}.$$

(3) $\Pr[\tilde{p}_{t+1} = 0 \mid p_t = 0] = 1$.

(4) For all $(p_1, \ldots, p_{t-1}) \in \text{supp}(\tilde{p}_1, \ldots, \tilde{p}_{t-1})$ such that $p_{t-1} > 0$,

$$\Pr[\tilde{p}_t = 0 \mid p_1, \ldots, p_{t-1}] = \epsilon_t \in (0, 1).$$

(5) There exists $\epsilon > 0$ such that $\epsilon_t \geq \epsilon$ for all t.

Condition (1) is for analytical simplicity. 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 rationalizable) 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) greatly simplifies the analysis. It says that the stochastic process satisfies a weak form of history independence in that if the bubble has not burst prior to period t, then the probability it bursts at t is independent of the exact sequence of past prices. It also requires this probability to be strictly between zero and one. This condition does allow for many other forms of history dependence; even ignoring histories where the bubble bursts before t, the probability distribution over nonzero period t prices can vary widely with the history of prices up to period t. Condition (5) further simplifies matters by requiring these probabilities to be bounded away from zero. It is easy to see (5) implies

$$\text{(6)} \qquad \lim_{T \to \infty} \prod_{t=1}^{T} (1 - \epsilon_t) = 0.$$

Hence if there are priors such that this process is canonically rationalized, they must satisfy (20a) and (20b). Consider (20a) first. Since $\epsilon_t \in (0, 1)$ for all t, $H_{2n-1} \in (0, 1)$ for all n and is strictly decreasing in n. Furthermore, condition (5) implies $1 - \epsilon_t \leq 1 - \epsilon$ for all t, so that

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

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

Define F_2 by $F_2(w) = 1 - H_{2n-1}$ for all $w \in [x_{2n-1}, x_{2n+1})$, $n = 0, 1, \ldots$. 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 x_{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 $H_{2n-1} \to 0$ and $x_{2n-1} \to \infty$ as $n \to \infty$ implies that $F_2(w) \to 1$ as $w \to \infty$. 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 submits a constant demand equal to his holdings 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.

where x_{-1} and x_0 are defined to be 0.

Any very nice stochastic process must have $p^*_{t+1} > p^*_t$. Hence $x_{2n+1} > 0$ and $-x_{2n} > 0$. Furthermore, note that

$$x_{2n+1} - x_{2n-1} = p^*_{2n+1} - p^*_{2n} \geq p^*_{2n+1} - (1-\epsilon_{2n+1})p^*_{2n+1} = \epsilon_{2n+1}p^*_{2n+1} \geq \epsilon p^*_1.$$

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

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

$$1 - \epsilon_{2n-1} = \frac{\Pr[w_2 \geq x_{2n-1}]}{\Pr[w_2 \geq x_{2n-3}]}, \quad n = 1, 2, \ldots$$

and

$$1 - \epsilon_{2n} = \frac{\Pr[w_1 \geq -x_{2n}]}{\Pr[w_1 \geq -x_{2n-2}]}, \quad n = 1, 2, \ldots .$$

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

$$H_{2n-1} = (1 - \epsilon_{2n-1})H_{2n-3}, \quad n = 1, 2, \ldots$$

and

$$G_{2n} = (1 - \epsilon_{2n})G_{2n-2}, \quad n = 1, 2, \ldots .$$

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:

(20a) $$H_{2n-1} = \prod_{k=1}^{n}(1 - \epsilon_{2k-1}), \quad n = 1, 2, \ldots$$

and

(20b) $$G_{2n} = \prod_{k=1}^{n}(1 - \epsilon_{2k}), \quad n = 1, 2, \ldots .$$

In other words, the bubble bursts in finite time with probability 1. In fact, one can show that (5) implies that the expected duration of the bubble is finite.

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) through (4) and (6),* $\mathrm{E}(\tilde{p}) = \infty$.

Since (6) is weaker than (5), the same holds for all nice stochastic processes.

Our main result is:

Theorem 2. *Every nice stochastic process $\{\tilde{p}_t\}$ is rationalizable.*

We prove this theorem by fixing an arbitrary nice stochastic process. We then construct stochastic processes for signals and an equilibrium given these processes which generate the given process for prices. The way the construction is carried out is that trader i's initial signal, $\tilde{s}^i_0$, gives him a finite "truncation date" — that is, a finite upper bound on the number of periods he will trade. 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. Let $I = 2$, $e_1 = 1$, and $e_2 = 0$. 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 $t = 2, 3, \ldots$, 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 satisfies conditions (1) through (5) above and so is nice. To show that this process is rationalizable for this economy, we construct signal distributions and equilibrium strategies for the two players. The two–point support allows us to do this with very simple signal distributions.

To construct these distributions, we let $\tilde{s}_t^i = 0$ with probability 1 for $t = 1, 2, \ldots$, and $i = 1, 2$. The initial signals, $\tilde{s}_0^1$ ($\tilde{s}_0^2$), are distributed geometrically on the even (odd) integers. More specifically,

$$\Pr[\tilde{s}_0^1 = 2n] = \Pr[\tilde{s}_0^2 = 2n-1] = (z-1)\left(\frac{1}{z}\right)^n, \quad n = 1, 2, \ldots$$

where $\tilde{s}_0^1$ and $\tilde{s}_0^2$ are independent random variables. In the equilibrium we construct, the agents trade at the equilibrium prices until one of them "quits." Trader i plans to quit — *i.e.*, refuse to purchase the asset — at period s_0^i. Hence the period at which the bubble actually bursts will be the smaller of s_0^1 and s_0^2. It is not hard to show that these strategies do generate the stochastic process described above.

More formally, the strategies are as follows. For any odd period t, if $x_{t-1}^1 = 1$ and p_k was z^k for every $k < t$, then trader 1 submits the demand schedule

$$y_t^1(p) = \begin{cases} 0, & \text{for } p = z^t; \\ 1, & \text{otherwise.} \end{cases}$$

Since trader 1 has the asset at this point, this corresponds to trader 1 wishing to sell the asset at price z^t and to retain it otherwise. (Recall that the demand functions are demands for final positions, not trades. Hence a trader 1's demand for 1 unit corresponds to not wanting to trade.) If $x_{t-1}^1 \neq 1$ or some past price was zero or inconsistent with the stochastic process above, he submits the demand $y_t^1(p) = x_{t-1}^1$ for all p. If $p_k = z^k$

the quantity he is supposed to buy (either his initial endowment plus that of his one trading partner or his initial endowment plus his share of that of his trading partner). Given the signal s_t^i, i knows that

$$\sum_{j \neq i} y_t^j(p) = \begin{cases} \bar{x} - x_*^i, & \text{for } p = s_t^i; \\ \bar{x}; & \text{otherwise.} \end{cases}$$

If i demands x_*^i at price s_t^i, the market clears at this price and he purchases this quantity. Otherwise, just as before, the market breaks down and he earns zero continuation profits.

To show that following $\sigma_t^i(h_t^i)$ is optimal, we must show that his continuation profits from purchasing x_*^i are nonnegative. Clearly, a sufficient condition for this is that the expectation of $\tilde{p}_{t+1}$ given h_t^i and given that i follows $\sigma_t^i(h_t^i)$ is at least s_t^i. It is not hard to see that this is true iff

$$\mathrm{E}[\tilde{p}_{t+1} \mid s_0^i, p_1, \ldots, p_t] \geq p_t.$$

By the Lemma, however, $\mathrm{E}[\tilde{p}_{t+1} \mid s_0^i, p_1, \ldots, p_t] = \mathrm{E}[\tilde{p}_{t+1} \mid p_1, \ldots, p_t]$. Hence this follows from the assumption that $\{\tilde{p}_t\}$ is a submartingale. ▌

Proof of Theorem 3.

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

$$x_t = \sum_{j=1}^{t} (-1)^{j+1} p_j^*, \quad t = 1, 2, \ldots .$$

Notice that x_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

$$\begin{cases} \Pr[w_2 \geq x_t \mid w_2 \geq x_{t-2}], & \text{if } t \text{ is odd;} \\ \Pr[w_1 \geq -x_t \mid w_1 \geq -x_{t-2}], & \text{otherwise.} \end{cases}$$

Essentially regular histories.

Terminal histories. For any terminal $h_t^i \in H_{**}^i(t)$, every trader $j \neq i$ submits a constant demand of x_{t-1}^j at period t and in all future periods. Hence i's continuation profits are zero whatever he does, so he may as well follow $\sigma_t^i(h_t^i)$.

Nonterminal histories. Since h_t^i is essentially regular, i believes with probability one that all traders $j \neq i$ are using the strategies σ^j. Hence if he follows σ^i, the rest of the stochastic process $\{\tilde{p}_t\}$ will be generated. Recall, though, that trader i's initial signal, $\tilde{s}_0^i$, is constructed so that i knows that the bubble will burst — *i.e.*, prices will converge to zero permanently — no later than period s_0^i. Hence the expected payoff to continuing to follow σ^i is necessarily well–defined since it is the sum of finitely many random variables, each with a finite support. Furthermore, as we explain in more detail below, the expected payoff to any form of deviation is also well–defined. As a result, i's best reply is also well–defined.

First, suppose that $i \in A$ and t is odd or $i \in B$ and t is even, so that i is supposed to sell at period t. Given the signal s_t^i, i knows that

$$\sum_{j \neq i} y_t^j(p) = \begin{cases} \bar{x}, & \text{for } p = s_t^i; \\ \bar{x} - x_{t-1}^i, & \text{otherwise.} \end{cases}$$

If i demands 0 at price s_t^i, then the price will be s_t^i and he will sell his holdings. Otherwise, the market will not clear (or will clear at a price other than s_t^i with all traders maintaining their current positions). Since this creates a history which is obviously not regular, there will be no subsequent trade in this case. Hence if i does not demand 0 at price s_t^i, his continuation profits are zero. If he does demand 0, he can always refuse to engage in any trade from period $t+1$ on and will thus earn continuation profits of at least $s_t^i x_{t-1}^i \geq 0$. Hence it is optimal for him to follow $\sigma_t^i(h_t^i)$.

Finally, consider the case where $i \in A$ and t is even or $i \in B$ and t is odd. In this case, $x_{t-1}^i = 0$ and i is supposed to buy at t. Let x_*^i denote

for every $k < t$, $x_{t-1}^2 = 0$, and $t < s_0^2$, then trader 2 submits the following demand function:

$$y_t^2(p) = \begin{cases} 1, & \text{for } p = z^t; \\ 0, & \text{otherwise.} \end{cases}$$

That is, he wishes to purchase the one unit of the asset at price z^t and none at any other price. Otherwise, trader 2 submits the demand function $y_t^2(p) = x_{t-1}^2$ for all p. In any even period t, if $x_{t-1}^2 = 1$ and $p_k = z^k$ for all $k < t$, trader 2 submits

$$y_t^2(p) = \begin{cases} 0, & \text{for } p = z^t; \\ 1, & \text{otherwise.} \end{cases}$$

Otherwise, he submits $y_t^2(p) = x_{t-1}^2$ for all p. If $p_k = z^k$ for all $k < t$, $x_{t-1}^1 = 0$, and $t < s_0^1$, trader 1 submits

$$y_t^1(p) = \begin{cases} 1, & \text{for } p = z^t; \\ 0, & \text{otherwise.} \end{cases}$$

If any of these three conditions is violated, he submits $y_t^1(p) = x_{t-1}^1$ for all p.

Consider the stochastic process induced by these strategies. In period 1, trader 1 owns the one unit of the asset. Hence he will offer to sell it for z. If $s_0^2 = 1$, trader 2 refuses to buy the asset at any price. In this event, by convention, $p_1 = 0$ and no trade occurs. If $s_0^2 \geq 3$, he demands one unit of the asset at a price of z, so trade occurs at $p_1 = z$. The probability that $\tilde{p}_1 = z$, then, is just $1 - \Pr[\tilde{s}_0^2 = 1] = (1/z)$. If trade does not occur at this period, it will not occur ever again and so the price will remain zero. If trade occurs, we move on to period 2 with $p_1 = z$, $x_1^1 = 0$, and $x_1^2 = 1$. At this point, trade will occur at price z^2 iff $s_0^1 \geq 4$, which occurs with probability $1/z$. At the next period, trade only occurs if $s_0^2 \geq 5$. Conditional on $s_0^2 \geq 3$, the probability of this event is again $1/z$, etc. This is precisely the stochastic process we sought to rationalize.

Do these strategies form an equilibrium? 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 submit a demand function in which he refuses to trade at any price. Hence i may as well do so.

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 do so at any price. Hence i's choice is only whether to be willing to accept such a "proposal" or not. The strategy above calls for i to use his initial signal to determine how long to accept these proposals. Should he follow his signal in this fashion?

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. Alternatively, suppose he adopts the strategy of always offering to buy — *i.e.*, he chooses $T = \infty$. Then he loses money for sure since the probability j eventually refuses to buy is 1. 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

believes that all his trading partners are submitting constant demand functions of x_t^j for all p. Finally, i believes that the other traders, unaware that a deviation has occurred, are still following $\hat{\sigma}_t^j$. Hence i believes that

$$\sum_{j \neq i} y_t^j(p) = \bar{x} - x_t^i$$

for every p and that this will be true at every future period regardless of what he does. Hence he can never trade and so may as well follow $\sigma_t^i(h_t^i)$.

have probability zero according to $\psi_k^i(y_1^i, \ldots, y_{k-1}^i, \sigma^{\sim i}, S)$. Therefore, we can assign any (feasible) beliefs for i on this history.

We will assume that for any such history, trader i infers that the other deviators are his "trading partners." If $i \in A$, he infers that trader $j = g(i)$ deviated as well as the other traders in $g^{-1}(j)$. If $i \in B$, he infers that the traders in $g^{-1}(i)$ deviated. In both cases, i infers that no other trader deviated. It is easy to see that any trader who did not deviate will be unaware that a deviation has occurred. For any such history for i, say h_t^i, $\sigma_t^i(h_t^i)$ is the constant function x_{t-1}^i for all p. (Hence, in particular, i believes his trading partners will also submit constant demands equal to their current holdings.)

Let δ be any beliefs satisfying these conditions and which are consistent with σ and Bayes' Rule. As discussed above, it is easy to see that if σ is played, the induced stochastic process for prices is $\{\tilde{p}_t\}$.

Completing the Proof.

We now show that (σ, δ) is an equilibrium. Fix any i and suppose that all traders $j \neq i$ follow strategy σ^j. We show that it is optimal for i to follow σ^i at each possible history.

Obviously not regular histories. Fix any h_t^i that is obviously not a regular history for i. For any such history, $\delta_t^i(h_t^i)$ must put probability one on h_t^j being obviously not regular for j for every j. Hence i knows that every trader j will submit the constant function x_{t-1}^j at t and at all future periods. Hence regardless of the demand function i submits, he cannot make any trades at t or any future period. Hence he may as well follow $\sigma_t^i(h_t^i)$.

Other not regular histories. Consider any other $h_t^i \notin H_{**}^i(t)$. By assumption, i believes that the deviators are he and his trading partners and that no other traders are aware that a deviation has occurred. Furthermore, i

[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 $h_i(s_0^1, s_0^2)$ denote the profits earned by trader i as a function of the two initial signals in this equilibrium and let G_i denote the distribution function for i's initial signal. 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

$$(7) \qquad \mathrm{E}_{\tilde{s}_0^2}[h_1(s_0^1, \tilde{s}_0^2) \mid s_0^1] = \int_{\tilde{s}_0^2} h_1(s_0^1, \tilde{s}_0^2)\, dG_2(\tilde{s}_0^2) = \mathrm{E}[\tilde{p}_1] = 1, \quad \forall s_0^1.$$

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

$$(8) \qquad \mathrm{E}_{\tilde{s}_0^1}[h_2(\tilde{s}_0^1, s_0^2) \mid s_0^2] = \int_{\tilde{s}_0^1} h_2(\tilde{s}_0^1, s_0^2)\, dG_1(\tilde{s}_0^1) = 0, \quad \forall s_0^2.$$

But for any pair of initial signals, the profits earned by one trader is exactly equal to the loss earned by the other. That is, $h_2(s_0^1, s_0^2) = -h_1(s_0^1, s_0^2)$. Substituting into (8),

$$(9) \qquad \mathrm{E}_{\tilde{s}_0^1}[h_1(\tilde{s}_0^1, s_0^2) \mid s_0^2] = \int_{\tilde{s}_0^1} h_1(\tilde{s}_0^1, s_0^2)\, dG_1(s_0^1) = 0, \quad \forall s_0^2.$$

But equations (7) and (9) are inconsistent. If we multiply the integral in (7) by $dG_1(\tilde{s}_0^1)$ and integrate over $\tilde{s}_0^1$, we see that

$$\mathrm{E}_{\tilde{s}_0^1, \tilde{s}_0^2}[h_1(\tilde{s}_0^1, \tilde{s}_0^2)] = \int_{\tilde{s}_0^1} \int_{\tilde{s}_0^2} h_1(\tilde{s}_0^1, \tilde{s}_0^2)\, dG_1(\tilde{s}_0^1)\, dG_2(\tilde{s}_0^2) = 1.$$

However, if we multiply the integral in (9) by $dG_2(\tilde{s}_0^2)$ and integrate over $\tilde{s}_0^2$, we obtain a contradiction:

$$\mathrm{E}_{\tilde{s}_0^1, \tilde{s}_0^2}[h_1(\tilde{s}_0^1, \tilde{s}_0^2)] = \int_{\tilde{s}_0^2} \int_{\tilde{s}_0^1} h_1(\tilde{s}_0^1, \tilde{s}_0^2)\, dG_1(\tilde{s}_0^1)\, dG_2(\tilde{s}_0^2) = 0.$$

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 s_0^1 are well–defined for every s_0^1 and are even independent of s_0^1. 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 — 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

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

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 $1 - 1 + 1 - 1 + 1 - \ldots$ does not converge. Hence the expectation is not defined.

We will say that $h_t^i \in H_t^i \setminus H_*^i(t)$ is ***an essentially regular history for i*** if there is a regular t length history for i, say $\hat{h}_t^i(h_t^i)$ which differs only in some of its y_k^i components. In other words, it is a history on which i submitted a different demand function than that called for by $\hat{\sigma}^i$, but the deviation had no effect in the sense that trades and prices followed the equilibrium path. For such a history, let $\sigma_t^i(h_t^i)$ equal $\hat{\sigma}_t^i(\hat{h}_t^i(h_t^i))$. Let $H_{**}^i(t)$ denote the union of $H_*^i(t)$ with the set of essentially regular t length histories for i.

We will say that $h_t^i \in H_t^i \setminus H_*^i(t)$ is ***obviously not a regular history for i*** if there is an $k \leq t-1$ such that $p_k \neq s_k^i$. In this case, every trader knows that some trader has deviated. For such a history, let $\sigma_t^i(h_t^i)$ be the constant function x_{t-1}^i for all p.

Let $\bar{H}_t^i$ denote the remaining histories in $H_t^i \setminus H_*^i(t)$ and fix any $h_t^i \in \bar{H}_t^i$. By definition, $p_k^i = s_k^i$ for all $k \leq t-1$. Furthermore, h_t^i differs from histories in $H_*^i(t)$ in some component other than the demands i has submitted in the past. Hence it must be true that i's past holdings of the asset differ from his equilibrium holdings even though the sequence of prices does not differ from the equilibrium sequence. For each $k \leq t-1$, let h_k^i denote the projection of h_t^i onto H_k^i. Let $r+1$ denote the first period such that $h_{r+1}^i \notin H_*^i(r+1)$. Hence $h_r^i \in H_*^i(r)$. Could it be true that $y_r^i = \hat{\sigma}_r^i(h_r^i)$ — *i.e.*, that trader i submitted his equilibrium demand function at period r? If so, the fact that $p_r = s_r^i$ implies that trader i's position must x_r^i must equal $\hat{\sigma}_r^i(h_r^i)(s_r^i)$. But then the history h_{r+1}^i must be a regular history for i, a contradiction. Therefore, i deviated at period r.

Could it be true that only trader i deviated at period r? If so, then all other traders' holdings must be those given in (19) above. But since the holdings must add to $\bar{x}$, this implies that i's holdings cannot differ from those given in (19), a contradiction. Therefore, some other trader must also deviate from $\hat{\sigma}$ in period r for the history h_{r+1}^i to be realized. More precisely, for every k between $r+1$ and t, any feasible $\delta_k^i(h_k^i)$ must assign probability zero to the event that for all $j \neq i$, $h_k^j \in H_{**}^j(s)$. Hence h_k^i must

s_t^i (which is equal to the realization of $\tilde{p}_t$ for all i) and refuse to buy or sell any quantity at any other price. The traders in set B each offer to buy the endowment of their "trading partner" in set A plus retain their own initial endowment, refusing to buy or sell at any other price. For every even period, the same occurs with the roles of A and B reversed. When $\#A > \#B$, the only change is that a given $i \in B$ may have more than one "partner" in A. These "partners" then split i's endowment when they buy. If all traders follow these strategies, then the price sequence generated will match the sequence of signals. Since these were chosen to match the stochastic process for prices we wanted to rationalize, the generated stochastic process will be the desired one. However, the $\hat{\sigma}$ strategies may not be an equilibrium. Our construction works by having the traders follow these strategies unless there has been some deviation from them in the past and letting the bubble collapse if deviation occurs. To define this precisely, we must first define equilibrium histories given these strategies.

The set of t length histories for trader i which will be his equilibrium path histories will be denoted $H_*^i(t)$ and referred to as the *regular t length histories for i*. H_*^i will be the collection of these histories for i. We define H_*^i recursively. Let $H_*^i(1) = H_1^i$. (Recall that $H_1^i = S_1^i$ — that is, the support of $(\tilde{s}_0^i, \tilde{s}_1^i)$. Hence all these histories may occur in equilibrium.) Given $H_*^i(t-1)$, let $H_*^i(t)$ denote the t length histories

$$h_t^i = (s_0^i, s_1^i, \ldots, s_t^i, y_1^i, \ldots, y_{t-1}^i, x_1^i, \ldots, x_{t-1}^i, p_1, \ldots, p_{t-1}) \in H_t^i$$

such that the following conditions hold. First, the projection of h_t^i onto H_{t-1}^i, say h_{t-1}^i, is an element of $H_*^i(t-1)$. Second, $p_{t-1} = s_{t-1}^i$ and $y_{t-1}^i = \hat{\sigma}_{t-1}^i(h_{t-1}^i)$. Finally, if h_{t-1}^i is nonterminal,

$$(19) \qquad x_{t-1}^i = \begin{cases} e_i + e_{g(i)}, & \text{if } t \text{ is odd and } i \in A; \\ e_i + (1/\#g^{-1}(i)) \sum_{j \in g^{-1}(i)} e_j, & \text{if } t \text{ is even and } i \in B; \\ 0, & \text{otherwise.} \end{cases}$$

Finally, let $H_*^i = \cup_{t=1}^{\infty} H_*^i(t)$. For every i and every $h_t^i \in H_*^i(t)$, let $\sigma_t^i(h_t^i) = \hat{\sigma}_t^i(h_t^i)$.

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 s_0^i.

The real question, then, is what, if anything, the nonexistence of the *ex ante* expectation means in economic terms.[4] Is existence a technical nicety or a meaningful and important restriction?

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

[4] One implication it certainly has is that we must define equilibrium as we did above. It is standard to require that strategies be *ex ante* optimal, but, of course, we cannot evaluate this. This is why our definition of equilibrium only requires interim optimality.

interim world[5] — that is, they make decisions after they observe their initial signals. 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 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

[5] 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]

Construction of Equilibrium Strategies and Histories.

Choose any function $g : A \to B$ which is onto. Since $\#A \geq \#B$, such functions exist. We construct equilibrium strategies by first specifying strategies which are used along the equilibrium path. We then use these strategies to identify which histories are equilibrium histories — *i.e.*, exactly which histories these strategies are used for. Finally, we complete the strategies by defining them for other histories. This procedure is complicated by the fact that no trader observes the demand functions submitted by other traders nor the trades made by other traders.

Call a t length history h_t^i *terminal* if $s_t^i = 0$ where s_t^i is the realization of $\tilde{s}_t^i$ in h_t^i. (In what follows, we will often omit this kind of explanation. Hence if we refer to a history and realizations of certain signals (or other possible components of a history), these realizations (or other components) should be interpreted as being those in the history.) Otherwise, it is *nonterminal*. Define the strategy $\hat{\sigma}^i$ as follows. For every nonterminal history $h_t^i \in H_t^i$, if $i \in A$ and t is odd or if $i \in B$ and t is even,

$$\sigma_t^i(h_t^i)(p) = \begin{cases} 0, & \text{for } p = s_t^i; \\ x_{t-1}^i, & \text{otherwise.} \end{cases}$$

If h_t^i is nonterminal, $i \in A$, and t is even,

$$\sigma_t^i(h_t^i)(p) = \begin{cases} e_i + e_{g(i)}, & \text{for } p = s_t^i; \\ x_{t-1}^i, & \text{otherwise.} \end{cases}$$

If h_t^i is nonterminal, $i \in B$ and t is odd,

$$\sigma_t^i(h_t^i)(p) = \begin{cases} e_i + (1/\#g^{-1}(i)) \sum_{j \in g^{-1}(i)} e_j, & \text{for } p = s_t^i; \\ x_{t-1}^i, & \text{otherwise.} \end{cases}$$

Finally, if h_t^i is terminal, $\sigma_t^i(h_t^i)$ is the constant function x_{t-1}^i for every p.

To see the intuition, suppose $\#A = \#B$ so that g is one–to–one. Then we can think of $g(i)$ as i's "trading partner" in equilibrium. Intuitively, in period t, if t is odd, the traders in set A offer to sell their holdings at price

Finally, we construct the joint distribution. To do so, we assume that n_A and n_B are conditionally independent in the sense that for any $W \in \mathcal{W}$, any $N_A \subseteq \mathrm{N}_A$, and any $N_B \subseteq \mathrm{N}_B$,

$$(17) \qquad \frac{\mu(W, N_A, N_B)}{q(W)} = \left[\sum_{n \in N_A} \mu_A(n \mid W)\right]\left[\sum_{n \in N_B} \mu_B(n \mid W)\right].$$

This assumption, together with those above, completely specifies the joint distribution. Given any measurable set $T \in \mathcal{T}$, (17) defines the measure of T as a function of the conditional distributions. Equations (15) and (16) completely define the conditional distributions as functions of the marginals given in (13) and (14).

One of the most important aspects of this construction is summarized by the following lemma. The proof of this lemma is available on request.

Lemma. *For any even (odd) t, any $i \in A$ ($i \in B$), any $s_0^i \in \mathrm{supp}(\tilde{s}_0^i)$, and $(p_1, \ldots, p_t) \in \mathrm{supp}(\tilde{p}_1, \ldots, \tilde{p}_t)$,*

$$(18) \qquad \mathrm{E}[\tilde{p}_{t+1} \mid \tilde{s}_0^i = s_0^i, \underline{p}] = \mathrm{E}[\tilde{p}_{t+1} \mid \underline{p}]$$

where $\underline{p} = (p_1, \ldots, p_t)$.

To understand the meaning of this lemma, consider $t = 2$, $i \in A$, and $s_0^i = 4$. Recall that the initial signals give an upper bound on the date at which the bubble bursts. Hence given this initial signal, trader i knows that the price will be 0 from period 4 on. Suppose, though, that it is period 2, he has seen the period 2 signal, and the price is still strictly positive. Should he buy at price p_2? Clearly, this depends on the expectation of $\tilde{p}_3$ conditional on his information. His information at this point consists of his initial signal plus the observation of p_1 and p_2. (Recall that his period 1 and 2 signals, in equilibrium will just equal these prices and so will convey no additional information.) The lemma simply says that his initial signal does not affect this expectation. As we will see, this fact will guarantee that trader i purchases the asset in this situation.

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.

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, we believe, still hold. The other important change is our assumption that agents are not price takers. If agents were price takers, our finite truncation device could not generate bubbles. To see why, suppose that there are two traders, where trader 1 is endowed with one unit of the asset and trader 2 with none as in the example above, but now suppose 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.

The price–taking assumption implies that any trader for whom the expected value of the next period's price is strictly higher than the current price 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 is true no trader who expects a higher price in the next period will sell in the current period. Similarly, no trader who expects a lower price in the next period will buy in the current period. It is easy to use arguments akin to those of Milgrom and Stokey or Geanakoplos and Sebenius to show that this implies that no trade can occur unless it is common knowledge that for *both* traders, the expected value of tomorrow's price is equal to today's price. But if trader i has a finite date at which he can truncate his horizon, this condition cannot hold at the next-to-last date for i. If i plans to stop trading after the current date, he must expect the next price to be zero.

To see how this argument breaks down in our model, consider again the equilibrium constructed above. Suppose that $s_0^1 \geq 4$, so that that trader 1 intends to 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. While the effect a trader has on the price is quite sharp in the demand submission game we analyze, we expect similar results to hold with prices which depend continuously on a trader's actions.

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, must hold in our model simply because the demand functions agents submit are only defined for such prices. If we change this part of the game, equilibria in which prices are sometimes negative are possible. 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.

and that

$$\sum_{n=1}^{\infty} \mu_B(n) = 1.$$

We now use the marginals to define various conditional distributions. We will let $\mu_A(n \mid W)$ denote the conditional probability that $n_A = n$ given $\omega \in W$ and analogously for $\mu_B(n \mid W)$. For any $W \in \mathcal{W}$ such that $W \subseteq W_{2k-1}$,

(15a) $$\mu_A(n \mid W) = \Pr[n_A = n \mid n_A \geq k]$$

and

(15b) $$\mu_B(k \mid W) = 1.$$

For any $W \in \mathcal{W}$ such that $W \subseteq W_{2k}$,

(15c) $$\mu_A(k \mid W) = 1$$

and

(15d) $$\mu_B(n \mid W) = \Pr[n_B = n \mid n_B \geq k + 1].$$

Finally, we can extend these definitions to construct the conditional distributions for any $W \in \mathcal{W}$ by

(16) $$\mu_A(n \mid W) = \sum_{t=1}^{\infty} \frac{q(W \cap W_t)}{q(W)} \mu_A(n \mid W \cap W_t)$$

and analogously for n_B.

It is easy to see that these conditional distributions are well-defined. Also, notice that these distributions are constructed so that

$$\tilde{T} = \min\{2n_A, 2n_B - 1\}$$

with probability 1.

One can show by induction on N that

$$(13) \qquad \sum_{n=1}^{N} \mu_A(n) = 1 - Z_{N+1}^A.$$

(The proof of this claim is available on request from the authors.) In light of this,

$$\lim_{N\to\infty} \sum_{n=1}^{N} \mu_A(n) = 1 \text{ iff } \lim_{N\to\infty} Z_N = 0$$

or

$$\lim_{N\to\infty} \prod_{j=1}^{N} \frac{\Pr[\tilde{T} \geq 2j+1]}{\Pr[\tilde{T} \geq 2j]} = 0.$$

Note that

$$\frac{\Pr[\tilde{T} \geq 2j+1]}{\Pr[\tilde{T} \geq 2j]} = \Pr[\tilde{T} \geq 2j+1 \mid \tilde{T} \geq 2j] = 1 - \epsilon_{2j}.$$

By condition (5), $1 - \epsilon_{2j} \leq 1 - \epsilon$. Hence

$$\lim_{N\to\infty} \prod_{j=1}^{N} \frac{\Pr[\tilde{T} \geq 2j+1]}{\Pr[\tilde{T} \geq 2j]} \leq \lim_{N\to\infty} (1-\epsilon)^N = 0.$$

Therefore, the marginal on N_A is well–defined and puts probability 1 on $n_A < \infty$.

The marginal on N_B is defined analogously. More specifically,

$$(14) \qquad \mu_B(n) = \frac{\Pr[\tilde{T} = 2n-1]}{\Pr[\tilde{T} \geq 2n-1]} \prod_{j=1}^{n-1} \frac{\Pr[\tilde{T} \geq 2j]}{\Pr[\tilde{T} \geq 2j-1]}.$$

Analogously to the above, one can show that

$$\sum_{n=1}^{N} \mu_B(n) = 1 - Z_{N+1}^B$$

where

$$Z_n^B = \prod_{j=1}^{n-1} \frac{\Pr[\tilde{T} \geq 2j]}{\Pr[\tilde{T} \geq 2j-1]}$$

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.

The condition that zero prices is an absorbing state is also not necessary. It is straightforward to construct examples in which a new bubble may start once the old one has burst. Condition (4), that the probability of the bubble bursting at period t is independent of the sequence of past prices, is also certainly not necessary. A general proof without this assumption is difficult, but examples without it are easy to construct.

Similarly, condition (5), that the probability the bubble bursts at t given that it hasn't already burst is uniformly bounded away from zero, is also 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.

IV. Wealth Uncertainty.

Theorem 2 shows that uncertainty about intrinsically irrelevant phenomena (*i.e.*, sunspots) can generate a huge class of bubbles. The advantage of focusing on this kind of uncertainty is that we can isolate the role of strategic uncertainty (that is, uncertainty about what others will do) from uncertainty about intrinsically relevant variables like dividend streams. On the other hand, it does seem to us that uncertainty regarding variables of "real" interest is more plausible as a description of real markets.[6] In this section, we show that uncertainty about intrinsically relevant variables can also produce bubbles.

To illustrate the point most straightforwardly, we consider a very simple economy, similar to the example discussed in Section III. There are two traders, 1 and 2. As before, the two goods are the asset and "money" or "wealth." Trader 1's initial endowment is one unit of the asset and w_1 units of money. Trader 2 is not endowed with any of the asset, only with w_2 units of money. Each trader knows his own wealth, but neither knows the other's endowment of money. They have a common prior over the pair of wealths. For simplicity, we assume that the wealths are independently distributed, so that the common prior can be written as $F_1(w_1)F_2(w_2)$, where F_i is the (cumulative) distribution function for trader i's initial wealth. We assume that

$$\text{(10)} \qquad \text{For } i = 1, 2, \quad F_i(w) = 0 \;\; \forall w < 0 \;\text{ and }\; \lim_{w \to \infty} F_i(w) = 1.$$

Hence the wealth of each trader is finite with probability 1, though the supports of the priors may not have finite upper bounds. Later, we show that the equilibria we construct require the supports to not be bounded.

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

6 Bertocchi [1988] makes a similar argument. See also Spear [1989].

For every $i \in A$ and every $\omega \in O$,

$$\tilde{s}_0^i(\omega, n_A, n_B) = 2n_A.$$

Finally, for every $i \in B$ and $\omega \in E$,

$$\tilde{s}_0^i(\omega, n_A, n_B) = 2n_B - 1.$$

Less formally, the traders in A receive an even signal which is equal to the realization of $\tilde{T}$ if this date is even. The traders in B receive an odd signal equal to the realization of $\tilde{T}$ if this is odd. We will define the probabilities on n_A and n_B so that traders in A receive a signal larger than $\tilde{T}$ if the realization is odd and analogously for the traders in B. Hence the realization of $\tilde{T}$ will exactly equal the smallest realization of the initial signals.

We define the probability measure μ on Θ in three steps. First, we define the marginal distributions on Ω, $\mathbf{N}_A$, and $\mathbf{N}_B$. Second, we define conditional distributions for n_A and n_B given ω. Finally, we show that there is a consistent joint distribution yielding the appropriate marginals and conditionals. For the first step, the marginal on Ω is taken to be q — that is, the same as the probability measure for Ω which defines the stochastic process for prices.

The marginal on $\mathbf{N}_A$ will be denoted μ_A and is given by:

$$\text{(12)} \qquad \mu_A(n) = \frac{\Pr[\tilde{T} = 2n]}{\Pr[\tilde{T} \geq 2n]} \prod_{j=1}^{n-1} \frac{\Pr[\tilde{T} \geq 2j + 1]}{\Pr[\tilde{T} \geq 2j]} .$$

(Throughout, we maintain the convention that $\prod_{j=1}^{0} y_j = 1$ for any $\{y_j\}$ sequence.) It is easy to see that this specification has $\mu_A(n) \in [0, 1]$ for all n.

To show that $\mu_A(n)$ sums to one, let

$$Z_n^A = \prod_{j=1}^{n-1} \frac{\Pr[\tilde{T} \geq 2j + 1]}{\Pr[\tilde{T} \geq 2j]} .$$

$\tilde{T}$, then, is the date at which the bubble bursts. By condition (5), $\tilde{T}$ is finite with probability 1. For each T, let

$$W_T = \{\omega \in \Omega \mid \tilde{T}(\omega) = T\}.$$

(Note that each $W_T \in \mathcal{W}$.) Also, let

$$O = \bigcup_{n=1}^{\infty} W_{2n-1}$$

and

$$E = \bigcup_{n=1}^{\infty} W_{2n}.$$

These are just the sets of states in which the bubble bursts at an odd or even date respectively. Note that $\Pr[\omega \in O] > 0$ and $\Pr[\omega \in E] > 0$.

The state space for the signals is taken to be $\Theta = \Omega \times \mathrm{N} \times \mathrm{N}$ where $\mathrm{N} = \{1, 2, \ldots\}$. A generic element of this set will be denoted (ω, n_A, n_B). For clarity, we will refer to the first N as N_A and the second as N_B. The σ-algebra for Θ will be denoted $\mathcal{T}$ and is the direct product of $\mathcal{W}$, the power set of N_A, and the power set of N_B. The stochastic process for the signals is defined as follows. For every $t \geq 1$ and for all i,

$$\tilde{s}_t^i(\omega, n_A, n_B) = \tilde{p}_t(\omega).$$

In words, prior to the submission of demands at period t, each trader observes a common signal revealing what the period t price is "supposed" to be. The initial signals will give upper bounds on $\tilde{T}$, the date at which the bubble bursts. To define the initial signals, we partition the set of traders into two nonempty sets A and B. This can be done in any fashion which satisfies $\sum_{i \in A} e_i > 0$ and $\#A \geq \#B$. Then for all $i \in A$ and every $\omega \in E$,

$$\tilde{s}_0^i(\omega, n_A, n_B) = \tilde{T}(\omega).$$

Analogously, for every $i \in B$ and every $\omega \in O$,

$$\tilde{s}_0^i(\omega, n_A, n_B) = \tilde{T}(\omega).$$

returns as unbounded.[7]

The preference assumptions are exactly as in the previous section. Again, it is common knowledge that the asset never pays any dividends, traders are risk neutral, and do not discount the future. Trade is modeled exactly as before with one change. Now we assume that the exchange agreed upon occurs at the date of agreement. Hence when a trade takes place, the traders' wealths adjust accordingly. We now assume, in addition, 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.

Strategies are defined analogously to the definitions given in Section II with two changes. First, there are no signals observed, so that strategies depend only on initial wealth and the history of trades and prices. Second, as noted above, a strategy which calls for spending more than one's current wealth is not feasible and hence is excluded from the strategy set. More precisely, if the history of prices and i's positions is given by $p_1, \ldots, p_{t-1}$ and $x_1^i, \ldots, x_{t-1}^i$, then the demand function i submits in period t, $y_t^i(p)$, must satisfy:

$$p[y_t^i(p) - x_{t-1}^i] \leq w_i - \sum_{k=1}^{t-1} p_k[x_k^i - x_{k-1}^i], \quad \forall p$$

where $x_0^i = e_i$.

As in the previous section, we show that a large class of stochastic

7 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 levels of the other agent, then the assumption of unbounded supports is appropriate.

processes for prices can be rationalized. For simplicity, though, we consider a smaller class than the nice stochastic processes defined in Section III. We say that a stochastic process for prices $\{\tilde{p}_t\}$ is *very nice* if it is nice and satisfies:

(1′) For all t, $\text{supp}(\tilde{p}_t) = \{0, p_t^*\}$ for some $p_t^* > 0$.

Of course, condition (1′) is much stronger than condition (1). As we will see, the two–point support for $\tilde{p}_t$ leads to a very intuitive construction of equilibria.

In the equilibria we construct, if trade occurs at period t, it occurs at price p_t^*. Trade occurs at period t iff trade has occurred in every past period at the equilibrium price and the trader who is to buy the asset at period t has accumulated enough wealth to pay p_t^* to the other trader. Intuitively, then, there is a bubble until one trader cannot afford to continue trading at which point the bubble bursts.

To support this outcome, the strategies of the traders are that the trader holding the asset at period t, say trader i, submits demand schedule

$$y_t^i(p) = \begin{cases} 0, & \text{for } p = p_t^*; \\ 1, & \text{otherwise.} \end{cases}$$

If for all $k < t$ trade occurred in period s at price p_k^* and if j (the trader who is not holding the asset) has current wealth of at least p_t^*, he submits demand schedule

$$y_t^j(p) = \begin{cases} 1, & \text{for } p = p_t^*; \\ 0; & \text{otherwise.} \end{cases}$$

If either of these conditions is violated, he submits $y_t^j(p) = 0$ for all p. Given a very nice stochastic process $\{\tilde{p}_t\}$, we refer to these strategies as the *canonical strategies for* $\{\tilde{p}_t\}$. If there is a pair of distributions for initial wealth, F_1 and F_2, satisfying assumption (10) such that the canonical strategies for $\{\tilde{p}_t\}$ form an equilibrium which generates $\{\tilde{p}_t\}$, then we say that $\{\tilde{p}_t\}$ is *canonically rationalizable.*

Hence

$$\sum_{t=1}^{\infty} \frac{\epsilon_t}{1-\epsilon_t}(1-\epsilon_t^K) < \infty.$$

Since $\sum_t \epsilon_t < \infty$, it must be true that $\epsilon_t \to 0$ as $t \to \infty$. That is, fixing any $\underline{\epsilon} \in (0,1)$, there exists T such that $\epsilon_t \leq \underline{\epsilon}$ for all $t \geq T$. Therefore,

$$\sum_{t=T}^{\infty} \frac{\epsilon_t}{1-\epsilon_t}(1-\epsilon_t^K) \geq (1-\underline{\epsilon}^K) \sum_{t=T}^{\infty} \frac{\epsilon_t}{1-\epsilon_t},$$

implying

$$\text{(11)} \qquad \sum_{t=1}^{\infty} \frac{\epsilon_t}{1-\epsilon_t} < \infty.$$

Note that for any $x \geq 1$, $\log(x) \leq x - 1$. (To see this, note that it obviously holds at $x = 1$ and that the derivative of the left–hand side is less than the derivative of the right–hand side for all $x \geq 1$.) Hence (11) implies

$$\sum_{t=1}^{\infty} \log\left(\frac{1}{1-\epsilon_t}\right) < \infty$$

or

$$\log\left(\prod_{t=1}^{\infty} \frac{1}{1-\epsilon_t}\right) < \infty.$$

Therefore,

$$\prod_{t=1}^{\infty} (1-\epsilon_t) > 0,$$

contradicting (6). Hence $\sum_t \epsilon_t = \infty$ and $\mathrm{E}[\tilde{p}] = \infty$. ▌

Proof of Theorem 2.

Suppose $\{\tilde{p}_t\}$ is a nice stochastic process defined on a probability space $(\Omega, \mathcal{W}, q)$.

Construction of Signal Distributions.

Define

$$\tilde{T}(\omega) = \begin{cases} \min\{t \mid \tilde{p}_k(\omega) = 0 \ \ \forall k \geq t\}, & \text{if such a } t \text{ exists;} \\ \infty, & \text{otherwise.} \end{cases}$$

Appendix

Proof of Theorem 1.

For a stochastic process $\{\tilde{p}_t\}$ satisfying (1) through (4) and (6),

$$E(\tilde{p}) = \sum_{t=1}^{\infty} E[\tilde{p}_t \mid p_t \neq 0, p_{t+1} = 0] \Pr[\tilde{p}_t \neq 0, \tilde{p}_{t+1} = 0].$$

For any $p' \in \text{supp}(\tilde{p}_t)$ with $p' \neq 0$,

$$\Pr[\tilde{p}_t = p' \mid p_t \neq 0, p_{t+1} = 0] = \frac{\Pr[\tilde{p}_{t+1} = 0 \mid p_t = p'] \Pr[\tilde{p}_t = p']}{\Pr[\tilde{p}_{t+1} = 0 \mid p_t \neq 0] \Pr[\tilde{p}_t \neq 0]}.$$

By condition (4), $\Pr[\tilde{p}_{t+1} = 0 \mid p_t = p'] = \epsilon_{t+1}$ for all $p' \neq 0$ and hence $\Pr[\tilde{p}_{t+1} = 0 \mid p_t \neq 0] = \epsilon_{t+1}$ as well. Therefore, the right–hand side is $\Pr[\tilde{p}_t = p'] / \Pr[\tilde{p}_t \neq 0]$. Hence

$$E[\tilde{p}_t \mid p_t \neq 0, p_{t+1} = 0] = \frac{E[\tilde{p}_t]}{\Pr[\tilde{p}_t \neq 0]}.$$

Substituting this into the first equation and using the submartingale property,

$$E[\tilde{p}] \geq E[\tilde{p}_1] \sum_{t=1}^{\infty} \frac{\Pr[\tilde{p}_t \neq 0, \tilde{p}_{t+1} = 0]}{\Pr[\tilde{p}_t \neq 0]}.$$

By property (4) again,

$$E[\tilde{p}] \geq E[\tilde{p}_1] \sum_{t=1}^{\infty} \epsilon_{t+1}.$$

Since (1) implies $E[\tilde{p}_1] > 0$, we see that $E[\tilde{p}] = \infty$ if $\sum_t \epsilon_{t+1} = \infty$.

Suppose that this is not true — that $\sum_{t=1}^{\infty} \epsilon_t < \infty$. Since $\epsilon_t < 1$ for all t, this implies

$$\sum_{t=1}^{\infty} \epsilon_t^k < \infty$$

for all $k \geq 1$. In particular, for any finite K,

$$\sum_{k=1}^{K} \sum_{t=1}^{\infty} \epsilon_t^k < \infty.$$

Theorem 3. *Every very nice stochastic process is canonically rationalizable.*

Theorem 3 says that there are distribution functions for initial wealth which rationalize the process, but does not say much about the properties of these distributions. The definition of canonical rationalizability requires these distributions to put probability one on initial wealth being finite. However, the next result says that we cannot strengthen this requirement very much. In particular, the supports cannot be bounded.

Theorem 4. *Suppose $\{\tilde{p}_t\}$ is a very nice stochastic process. For any independent priors on w_1 and w_2 such that the canonical strategies together with these priors generate $\{\tilde{p}_t\}$, $E(w_i) = \infty$ for $i = 1, 2$.*

In other words, even ignoring the issue of the optimality of the canonical strategies, if some priors can generate a very nice stochastic process, expected initial wealths must be infinite. This property is clearly related to Theorem 1 which showed that for a class of stochastic processes including all nice and hence all very nice processes, $E(\tilde{p}) = \infty$.

Theorem 4 provides the key linking our work with the "envelope switching problem" discussed by Nalebuff [1989]. (See also Brams and Kilgour [1991] and Brams, Kilgour, and Davis [1991].) The problem, essentially, is the following. A number is drawn at random from the set $\{1, 2, 4, 8, \ldots\}$. 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.

In Section III, 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. In Section III, initial wealth is irrelevant to the model and hence it is not straightforward to construct an appropriate comparison. However, in the model with wealth uncertainty, as Theorem 4 showed, expected initial wealth is infinite. In fact, it is easy to show that this implies that the expectation of the sum of initial wealth and trading profits — the appropriate analog here to expected utility as computed by Nalebuff — is infinite. To see this, note that the expectation of initial wealth plus profits must exist (including the possibility that the expectation is ∞). This is true because the expectation of any random variable which is always nonnegative must exist. Since traders are not allowed to spend more money than they have, wealth plus profits are nonnegative with probability 1. So suppose that this expectation is not infinite. Then, letting π_i denote equilibrium trading profits,

$$E[w_i + \pi_i] - E[w_i] = -\infty$$

implying $E[\pi_i] = -\infty$ — that is, the expectation is well-defined and is $-\infty$. But we already have seen that $E[\pi_i]$ does not exist.

Nalebuff concludes there is only a paradox under the "monstrous hypothesis" of infinite expected utility. It is worth noting that 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.) 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. Furthermore, 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.

V. 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 as long as the utility functions are unbounded from above. Finally, there are certainly many alternative price–setting institutions which would generate similar equilibria.

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.