Sergiu Hart
Hebrew University of Jerusalem

Yair Tauman
Tel Aviv University and SUNY at Stony Brook

Market Crashes without External Shocks*

It is shown that market crashes and bubbles can arise without external shocks. Sudden changes in behavior coming after a long period of stationarity may be the result of *endogenous* information processing. Except for the daily observation of the market, there is no new information, no communication, and no coordination among the participants.

I. Introduction

Market crash is usually considered an indication that the fundamentals of the economy have changed and recession is around the corner. This, however, need not be the case. For instance, in October 1987 Wall Street lost over 20% of its value in one day, but this was not followed by a recession. Moreover, in the days preceding the crash, there were no significant external events or "bad news" that could justify the dramatic price fall.

We argue here that market crashes (and, similarly, market bubbles) may well be the result of information processing by the participants—and nothing else. Moreover, in terms of market observables, it looks as if nothing is really changing. Still, underneath the surface, there is a gradual updating of information by the participants. Then, at a certain point in time, this causes a sudden change of behavior.

Specifically, the phenomenon we describe here has to do with the step-by-step advance in levels of "mu-

* We thank Kenneth J. Arrow, Robert J. Aumann, Avraham Beja, Yaacov Bergman, Marcel Brachfeld, Yoram Halevy, Ehud Kalai, Eugene Kandel, Andreu Mas-Colell, Abraham Neyman, and Dov Samet for interesting discussions on these topics, and the anonymous referee for useful suggestions. The research of S. Hart was partially supported by grants of the U.S.-Israel Binational Science Foundation and of the Israel Academy of Sciences and Humanities.
Contact the corresponding author, Sergiu Hart, at hart@huji.ac.il.

(*Journal of Business*, 2004, vol. 77, no. 1)

0021-9398/2004/7701-0001$10.00

tual knowledge" (what one knows about what the other knows, and so on). Each trading day increases this level through the daily market observables—prices, quantities traded, and so forth—which are common knowledge. However, the behavior does not necessarily change with the level of mutual knowledge. The behavior may be constant for all levels up to a certain critical level—where a jump occurs. Such a phenomenon has been exhibited by Geanakoplos and Polemarchakis (1982).[1] In their article, two agents communicate information to each other by repeatedly announcing and revising their posteriors. It is shown there that the agents can repeat exactly the same opinions, yet still manage to communicate relevant information. Here, we fit their example to our model and show that there is no need for any direct communication between the traders; observing the market suffices.

We emphasize, first, that the phenomenon exhibited here is completely endogenous. Except for the daily observation of the market, there is no new information—whether pertinent or irrelevant (like "sunspots"); also, there is no communication nor any coordination among the participants—whether expressed or tacit. And second, the stationary unchanging behavior of the market for arbitrarily long periods of time is no sign that nothing is happening. Underneath the surface, completely unobservable, information is being processed by the participants—which ultimately leads to a sudden change of behavior.

The behavior of the agents in the example will turn out to resemble rules actually used by traders in the market. For instance: "I am buying every day; but, if others keep selling every day, then at some point I will start selling too." Some of the so-called technical analysis is indeed of this kind. Intuitively, if others are willing to sell all the time, then the buyer will at some time have to take this fact into account: perhaps his assessment is incorrect after all. The framework of this article allows us to make such arguments precise.

Our analysis, in showing that there may be many periods of trade until the participants' information converges[2]—even when there is no new external information—may perhaps help to interpret the large volume of trade in some markets (like the global currency markets).

We present here a simple example. However, the phenomenon we highlight is robust and general. One may change almost every feature of the example (like the state space, the probabilities, the number of traders, the daily rules of behavior, the prices, and so on), without qualitatively affecting the result.

To summarize, no exogenous "shocks" are needed to explain sudden departures from stationary behavior; these may well be due to the participants updating their information, based on the market observables. This information processing is, however, not observable—there is no change in behavior—up to the point where it generates an abrupt switch of behavior: a "crash."

1. For other situations with similar behavior, see Geanakoplos and Sebenius (1983) and the survey of Geanakoplos (1994).

2. See 2 in Sec. III.

II. The Basic Example

There are two traders. Day after day they keep trading—one is selling and the other is buying. Then, at some time, the updating of information leads one of them to reverse his behavior—from buying to selling. This happens in the absence of any exogenous influence.

Let Ω be the set of all states of the world. Assume Ω contains nine states, say, $\Omega = \{1, 2, \ldots, 9\}$. The prior probability assessments of the two traders are denoted P_A and P_B, respectively; that is, Alice believes that the state ω in Ω has probability $P_A(\omega)$, and similarly for Bob. For simplicity, we take a uniform probability distribution for both: $P_A(\omega) = P_B(\omega) = 1/9$ for all ω in Ω. We emphasize that this "common prior" assumption plays no role in our analysis (see *a* and *b* in Sec. IV).

The private information of each trader is described as usual by a partition of the state space:[3] two states belong to the same cell of the partition if and only if the trader cannot distinguish between them. Call the two traders Alice and Bob. Alice's partition is

$$A_1 = \{1, 2, 3\}, \quad A_2 = \{4, 5, 6\}, \quad A_3 = \{7, 8, 9\},$$

and Bob's partition is

$$B_1 = \{1, 2, 3, 4\}, \quad B_2 = \{5, 6, 7, 8\}, \quad B_3 = \{9\}.$$

The interpretation is as follows: trader Bob, for example, cannot distinguish between states 1, 2, 3, and 4, nor between 5, 6, 7, and 8. If state 9 is the true state, then Bob will know that for sure. If, however, state 1 is the true state, then Bob will only know that it is either 1 or 2 or 3 or 4 (but not 5, 6, 7, 8, 9).

Consider the event $E = \{1, 5, 9\}$. For instance, assume that E is the event of a "bad" outcome (e.g., the company earnings will go down). Suppose that each one of the two traders behaves each day according to the following rule:

$$\begin{cases} \textit{Sell}, & \text{if the probability of } E \text{ is } 0.3 \text{ or more;} \\ \textit{Buy}, & \text{if the probability of } E \text{ is less than } 0.3. \end{cases}$$

Of course, the relevant probability is always computed given the current information.

The manner in which the cutoff point of 0.3 (or, for that matter, the whole policy) is determined is irrelevant to the analysis here; in particular, we abstract away from the stock price and the quantities.[4] Also, we note that it does not matter whether both traders use the same behavior strategy.

Assume that the true state of the world is $\omega_0 = 1$. Initially, Alice assesses

3. For a formal treatment, see Aumann (1999) and Aumann and Heifetz (2002).

4. For concreteness, one may assume that "sell" actually means "sell the quantity q for the price p," and "buy" means "buy the quantity q for the price p" (where q and p are the same for both decisions). Since, as we will see below, every day (up to the "crash") Alice will sell and Bob will buy, the price need not change. See 1 in Sec. III for some explicit demand functions.

TABLE 1 **The Information, the Probability Assessments, and the Actions**
(The States of the Event ***E*** Are Underlined)

		Alice				Bob			
Day	Common Information (before Day t)	A_1 $\underline{1}$ 2 3	A_2 4 $\underline{5}$ 6	A_3 7 8 $\underline{9}$	Action	B_1 $\underline{1}$ 2 3 4	B_2 $\underline{5}$ 6 7 8	B_3 $\underline{9}$	Action
$t = 1$	$\underline{1}$ 2 3 4 $\underline{5}$ 6 7 8 $\underline{9}$	1/3	1/3	1/3	*Sell*	1/4	1/4	1	*Buy*
$t = 2$	$\underline{1}$ 2 3 4 $\underline{5}$ 6 7 8	1/3	1/3	0	*Sell*	1/4	1/4		*Buy*
$t = 3$	$\underline{1}$ 2 3 4 $\underline{5}$ 6	1/3	1/3		*Sell*	1/4	1/2		*Buy*
$t = 4$	$\underline{1}$ 2 3 4	1/3	0		*Sell*	1/4			*Buy*
$t = 5$	$\underline{1}$ 2 3	1/3			*Sell*	1/3			*Sell*

the probability of E to be 1/3 (since at $\omega_0 = 1$ she knows that the true state is either 1 [which belongs to E], or 2 or 3 [which do not belong to E])—therefore, she gives an order to sell; Bob assesses this probability to be 1/4 (he knows the state is 1, 2, 3, or 4)—therefore, he gives an order to buy. So a transaction takes place.

We will show that this will happen not only on the first day, but also on each one of the first 4 days: the assessments of the two traders for the probability of E remain 1/3 and 1/4, respectively. On the fifth day, however, there is a sudden and major change: both assessments become 1/3, and both traders give orders to sell. So, a "crash" occurs after four seemingly "quiet and normal" days.

Let us see this in detail (see table 1).

Let $t = 1, 2, \ldots$ denote the "days," and Ω^t the common knowledge information at time t, before the traders choose their actions.

On day $t = 1$, we have $\Omega^1 = \Omega$. The assessments of Alice are

$$P_A(E \mid A_i \cap \Omega^1) = P_A(E \mid A_i) = 1/3, \quad \text{for } i = 1, 2, 3,$$

and those of Bob are

$$P_B(E \mid B_j \cap \Omega^1) = P_B(E \mid B_j) = \begin{cases} 1/4, & \text{if } j = 1, 2, \\ 1, & \text{if } j = 3. \end{cases}$$

Since $\omega_0 = 1$, the current information is A_1 for Alice and B_1 for Bob, so Alice sells and Bob buys. Note that Bob only computes $P_B(E \mid B_1)$ (he knows that the true state is in B_1). However, for the sequel, one also needs to know what Bob would have computed—and done—in the other states as well.

On day $t = 2$, it is common knowledge that Bob bought on the previous day; therefore, it is common knowledge that $B_3 = \{9\}$ did not occur.[5] There-

5. Note that, at $\omega_0 = 1$, both traders initially knew that the state is not 9 (so this was *mutually known*); however, it was *not commonly known*. To see why, let $F = \{1, 2, \ldots, 8\}$ and let $K_A F$ be the event "Alice knows F." Similarly, $K_B K_A F$ is the event "Bob knows that Alice knows F," etc. Then $K_A F = \{1, 2, \ldots, 6\}$, $K_B K_A F = \{1, 2, 3, 4\}$, $K_A K_B K_A F = \{1, 2, 3\}$, and $K_B K_A K_B K_A F = \phi$. Thus, it is never the case that Bob knows that Alice knows that Bob knows that Alice knows that the state is not 9. Hence, in no state of the world—in particular, at $\omega_0 = 1$—is F common knowledge.

fore, the new common knowledge information is $\Omega^2 = \{1, 2, \ldots, 8\}$. The new assessments are

$$P_A(E \mid A_i \cap \Omega^2) = \begin{cases} 1/3, & \text{if } i = 1, 2, \\ 0, & \text{if } i = 3, \end{cases}$$

for Alice, and

$$P_B(E \mid B_j \cap \Omega^2) = 1/4, \quad \text{for } j = 1, 2,$$

for Bob. Again, at $\omega_0 = 1$ a transaction takes place: Alice sells and Bob buys.

The new common knowledge information on day $t = 3$ is $\Omega^3 = \{1, 2, \ldots, 6\}$ (since Alice would not have sold at $t = 2$ if the state were in A_3), and we have on day 3:

$$P_A(E \mid A_i \cap \Omega^3) = 1/3, \quad \text{for } i = 1, 2,$$

and

$$P_B(E \mid B_j \cap \Omega^3) = \begin{cases} 1/4, & \text{if } j = 1, \\ 1/2, & \text{if } j = 2. \end{cases}$$

Thus, there is a transaction on day 3, and on day 4 it is common knowledge that the state is in $\Omega^4 = \{1, 2, 3, 4\}$ (since B_2 is commonly ruled out), and thus

$$P_A(E \mid A_i \cap \Omega^4) = \begin{cases} 1/3, & \text{if } i = 1, \\ 0, & \text{if } i = 2, \end{cases}$$

and

$$P_B(E \mid B_j \cap \Omega^4) = 1/4, \quad \text{for } j = 1.$$

Finally, on day 5 we get $\Omega^5 = \{1, 2, 3\}$ (since A_2 is commonly ruled out),

$$P_A(E \mid A_i \cap \Omega^5) = 1/3, \quad \text{for } i = 1,$$

and

$$P_B(E \mid B_j \cap \Omega^5) = 1/3, \quad \text{for } j = 1.$$

But now both Alice and Bob send orders to sell and a "crash" occurs.

What is happening in this example is the following. Initially, both Alice and Bob know that the state is[6] 1, 2, 3, or 4. However, this fact is *not* common knowledge between them. For example, from Bob's point of view, the state could well be 4, in which case Alice would have known that it is either 4, 5, or 6. So Bob does *not* know that Alice knows that it is 1, 2, 3, or 4. As time goes by, the trading increases the common knowledge (see the second column of table 1), until, on day 5, it reaches its conclusion: it is common knowledge that $\omega = 1, 2,$ or 3.

6. Alice knows even more: she knows that the state is not 4.

III. Extensions and Modifications

We discuss now a number of extensions and generalizations of the basic example.

1. For simplicity, we have abstracted away from prices and quantities. There is, however, no difficulty in introducing them explicitly. For a specific example, take the von Neumann–Morgenstern utility function of the traders to be $u(x) = \log(x + w)$, where x denotes monetary payoff and w initial "wealth." Assume that the price of the share will be a if E occurs, and b otherwise (with $a < b$). If the current price is p and the probability of E is π, then the net quantity demanded is easily obtained:

$$Q(p;\ \pi) = \arg\max_q \pi u[q(a - p)] + (1 - \pi)u[q(b - p)]$$

$$= \frac{w[\pi a + (1 - \pi)b - p]}{(b - p)(p - a)};$$

note that Q may well be negative—in which case the agent sells. For instance, take[7] $a = 1$, $b = 25$, and $w = 119$; then

$$Q(p;\ 1/4) = \frac{119(19 - p)}{(25 - p)(p - 1)},$$

and

$$Q(p;\ 1/3) = -\frac{119(p - 17)}{(25 - p)(p - 1)}.$$

Therefore, at all days t up to the last one, the equilibrium price[8] is $p^t = 18$; Bob buys one unit and Alice sells one unit ($q_B^t = Q(18;\ 1/4) = 1$ and $q_A^t = Q(18;\ 1/3) = -1$ for $t = 1, 2, 3, 4$). On the last day $t = 5$, the price drops to $p^5 = 17$ (at the bottom of the trading range [17, 19]), and there is no trade ($q_A^5 = q_B^5 = Q(17;\ 1/3) = 0$).

One may add to the model further economic elements, like varying prices and quantities, "limit" orders, and so on. This will make the analysis more complex, but—once the trading policies are appropriately defined (i.e., in such a manner that the information deduced remains the same as in the basic example)—it will not affect the phenomenon we exhibit.

2. For any positive integer n, a similar example consisting of $(n + 1)^2$ (instead of 9) states will yield $2n$ days where transactions occur, and a "crash" on day $2n + 1$ (see Aumann's example, at the bottom of p. 97 in Geanakoplos and Polemarchakis [1982]). So a "crash" can be preceded by arbitrarily many periods where trading occurs normally and nothing seems to change.

7. The only reason for these numbers is that all prices and quantities come out integers. Also, we assume for simplicity that the wealth of both traders is the same and does not change over time.

8. Determined by the market-clearing condition, i.e., $Q(p^t;\ 1/4) + Q(p^t;\ 1/3) = 0$.

3. We have assumed throughout that $\omega_0 = 1$; that is, a "bad" state (in E) is the true state. In the end (day 5), both traders indeed want to sell (the two traders, even by pooling their information, cannot distinguish between states 1, 2, and 3; they both sell since the probability assessment of E ends up being 1/3). However, exactly the same behavior would have resulted if the true state were $\omega_0 = 2$ or $\omega_0 = 3$—which are "good" states (not in E). Also, note that if $\omega_0 = 4$, then a "bubble" occurs (i.e., both buy) at $t = 4$.

4. If the state space is finite (or, more precisely, if the two information partitions are finite), then there can be only finitely many instances of information updating (after which, under the common prior assumption, the two assessments necessarily agree; this is the Agreement Theorem of Aumann [1976]).

5. We have made the example as simple as possible; in particular, there are only two traders. One may of course deal with more traders; one easy way is to have two types of traders, A (like Alice) and B (like Bob).

IV. Discussion

We conclude with a number of general remarks.

a) Robustness. It is important to emphasize that, even though the example seems to be very "special," it is not. One may change almost every feature without affecting the conclusion. For instance, it suffices for the prior probabilities to satisfy $|P_A(\omega) - 1/9| < 1/13$ and $|P_B(\omega) - 1/9| < 1/8$ for all ω in Ω. In contrast, this is not the case with, say, the Geanakoplos and Polemarchakis (1982) setup: any small perturbation invalidates it (since their posteriors need to be all equal, whereas ours need only satisfy an inequality: greater or less than 0.3).

b) No "common prior." A particular implication of the above robustness is that one may make the two priors P_A and P_B different. Thus our conclusion is independent of the "common prior" assumption, which underlies much of the literature (see, in particular, Aumann 1976; Geanakoplos and Polemarchakis 1982; and Geanakoplos and Sebenius 1983).

c) More sophisticated players. Our two traders update their probability assessments on the basis of the observed actions. Assume instead that they could make more sophisticated inferences. For instance, they would realize that the only way for an order to be executed in our example is for the two traders to give opposite orders. When this is taken into account, the decisions may be affected. For example, at $t = 4$, Bob reasons as follows: "The only case where my 'buy' order will be executed is when Alice sells; but that happens only when the state is 1, 2, or 3—and then I should sell, not buy!" This may "unravel" the whole example; for related analyses, see Geanakoplos and Sebenius (1983)[9] and the surveys of Morris and Shin (1997, 2003) as well as the references there.

9. Where it is shown that, under the common prior assumption, fully rational players will never trade.

Note that in our case the above reasoning applies when Alice and Bob are the only two traders in the market. If, however, there are other traders in addition to Alice and Bob (who are the two "major" informed traders), then transactions do not necessarily take place between the two. Still, each one observes the moves of the other (they are, after all, the major players in this market). Then the phenomenon in the example works as presented.

Another issue is that, in our example, Alice is initially better informed at $\omega_0 = 1$ than Bob: she knows that Bob's information is B_1. Therefore, she can construct ahead of time the whole process as in table 1. This can easily be remedied by taking $\omega_0 = 5$, or, if one wants the same number of periods before the crash, a bigger example as in 2 of Section III above.

References

Aumann, R. J. 1976. Agreeing to disagree. *Annals of Statistics* 4 (November): 1236–39.

———. 1999. Interactive epistemology, I and II. *International Journal of Game Theory* 28 (August): 263–314.

Aumann, R. J., and A. Heifetz. 2002. Incomplete information. In *Handbook of game theory*, vol. 3, ed. R. J. Aumann and S. Hart, pp. 1665–86. Amsterdam: Elsevier/North-Holland.

Geanakoplos, J. 1994. Common knowledge. In *Handbook of game theory*, vol. 2, ed. R. J. Aumann and S. Hart, pp. 1437–96. Amsterdam: Elsevier/North-Holland.

Geanakoplos, J., and H. M. Polemarchakis. 1982. We can't disagree forever. *Journal of Economic Theory* 28 (October): 192–200.

Geanakoplos J., and J. K. Sebenius. 1983. Don't bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association* 78 (June): 424–26.

Morris, S., and H. S. Shin. 1997. Approximate common knowledge and coordination: Recent lessons from game theory. *Journal of Logic, Language and Information* 6 (April): 171–90.

———. 2003. Global games: Theory and applications. In *Advances in economics and econometrics: Proceedings of the Eighth World Congress of the Econometric Society*, vol. 1, ed. M. Dewatripont, L. Hansen, and S. Turnovsky, pp. 56–114. Cambridge: Cambridge University Press.