

American Economic Association

Business as Usual, Market Crashes, and Wisdom After the Fact

Author(s): Andrew Caplin and John Leahy

Source: *The American Economic Review*, Vol. 84, No. 3 (Jun., 1994), pp. 548-565

Published by: American Economic Association

Stable URL: <http://www.jstor.org/stable/2118067>

Accessed: 09/07/2010 11:55

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=aea>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *The American Economic Review*.

<http://www.jstor.org>

Business as Usual, Market Crashes, and Wisdom After the Fact

By ANDREW CAPLIN AND JOHN LEAHY*

We present a three-stage model of market dynamics. In the first stage, routine behavior tends to keep information of common interest trapped in private hands. In the second stage, private information reaches a threshold that triggers some agents to alter their behavior; these actions release information to the market. The final stage involves the market's response to this news as other participants react to the initial departure from routine behavior. We present an application to industry investment. We also outline applications to the international debt crisis, to bank runs, and to political upheavals. (JEL D80, E32)

In this paper we consider how markets aggregate private information among agents who alter their behavior infrequently. Individuals often gather or receive private information, which they convey to others in the market through changes in their behavior. Much individual behavior, however, may be characterized as business as usual. It is dominated by commitments and routines that individuals can change only at some cost. These commitments and routines tend to impede the dissemination of private information, since they limit an individual's response to news. Information must change significantly before an individual is willing to bear the cost of adjustment. During the time in which individual behavior remains constant, it may be very difficult for others to ascertain the information that the individual possesses.

In such situations, the evolution of public knowledge will tend to be discontinuous. Each time an individual alters a routine, the market learns a great deal about the individual's private information, and this news

may prompt other agents to alter their behavior as well. If sufficiently many agents decide to act and their cumulative actions reveal information that is sufficiently novel, this process of action and reaction may lead to a radical change in market sentiment.

In the model that we present, this change in sentiment takes the form of a crash in which all agents choose to abandon the market. A crash of this type may tempt observers to second-guess market participants on several counts. Some observers may be troubled by the timing of the market's collapse. They may believe that such a significant change in market sentiment ought to have an equally significant cause. In this case, however, the absence of an immediate cause is not a reflection of irrational behavior on the part of market participants. The news that triggers such a crash is analogous to the straw that broke the camel's back. It may take only a small piece of additional information to cause an agent to break with business as usual. The market's reaction to this change in behavior then magnifies and propagates the initial disturbance.

Other observers may claim that the participants should have been able to spot the trouble coming. As evidence of the market's shortsightedness, they may point to the fact that much of the information revealed in the crash was in the hands of individual agents long before the crash took place. Again such criticism does not imply market irrationality. Rather, the criticism is itself a

*Caplin: Department of Economics, Columbia University, New York, NY 10027; Leahy: Department of Economics, Harvard University, Cambridge, MA 02138. We thank Ricardo Caballero, Prajit Dutta, Philip Jefferson, Andrew Postlewaite, David Romer, and Steven Strongin for helpful discussions, two anonymous referees for their comments, and the National Science Foundation for financial support.

manifestation of wisdom after the fact. Although much pertinent information is indeed known to individuals prior to the crash, routine behavior prevents agents from putting the pieces of the puzzle together. The crash is precisely the process by which the market aggregates the dispersed information.

These three stages—business as usual, market crashes, and wisdom after the fact—define our framework for analyzing the interaction between private information and inertia. We conclude this introduction by discussing three concrete economic settings in which this framework applies, motivating our modeling strategy, and relating our approach to the literature.

A. *Three Examples*

The first example involves bank runs of the type that were relatively common in the United States prior to the implementation of federal deposit insurance. Here the informational problem was that depositors did not know the investment portfolio of their bank. During the business-as-usual phase in which banks were able to satisfy their obligations, depositors could only make indirect inferences concerning their particular bank's solvency based on an analysis of general business conditions, as well as information on the solvency of other banks. In such a situation, depositors would rationally interpret the failure of another bank as a signal that their own bank was in trouble, especially if general business conditions were unfavorable. In this way a combination of weak business conditions together with a few initial failures could trigger a general run on banks.

This interpretation of bank runs is supported by the studies of Gary Gorton (1988), Frederic Mishkin (1991), and Charles Calomiris and Gorton (1991). Gorton shows that general bank runs were typically preceded by declines in economic indicators such as increased liabilities of failed businesses and declines in stock prices. These declines in general business conditions raised questions concerning the solvency of the banking system as a whole. Mishkin

shows that the actual panic was typically triggered by the failure of a major financial institution. Calomiris and Gorton provide further support for this informational interpretation of crashes. They argue that panics ended only when clearing-house associations provided the public with additional information on the solvency of individual banks.

The second setting involves bank loans to LDC's. In the late 1970's and early 1980's, many banks made sizeable loans to a variety of developing nations. They faced continuing demands for extensions of credit and uncertainty as to the ability of the borrowers to repay the loans. Lending continued until August 1982, when Mexico experienced difficulty servicing its loans. Thereafter, more than 40 nations defaulted, and voluntary lending practically ceased (see Jeffrey Sachs, 1988).

Our framework provides one possible perspective with which to view these events. The LDC situation involved banks that had incomplete information concerning debtors' ability to repay, and debtors who had little incentive to reveal potential future difficulties. This business-as-usual phase continued as long as countries continued to service their loans. While there were signals of trouble, such as high interest rates, declining commodity prices, and the recession in the developed nations, there remained the possibility that borrowers could ride out these troubles with a little additional financing.¹ When Mexico found itself unable service its debts in August 1982, creditors treated this as a signal that the troubles of many debtors were severe, and the crash in lending ensued. In the aftermath, many commentators criticized the earlier loans, saying that the banks should have foreseen the impending crisis.

¹In June 1981, The Chairman of Citibank, Walter Wriston argued in reference to LDC lending, "In my view, this fear that banks have reached a limit will turn out to be wrong tomorrow, as it always has in the past" (William Greider, 1987 p. 433).

The third setting involves mass political actions, such as the demonstrations against Communist regimes in Eastern Europe studied by Susanne Lohmann (1992a,b). Here the problem was that citizens did not know how widespread dissatisfaction with the government was. Given the risks involved in revealing opposition to the government, most people adopted the business-as-usual attitude of acquiescence. At some stage, private dissatisfaction reached a level high enough to trigger some individuals to express their views publicly in the form of demonstrations. These initial demonstrations and possibly the encouragement of bystanders helped to reveal the extent of opposition. This encouraged others who were discontent to join the demonstrations and led to the eventual collapse of the regime.

While these examples differ in important ways, they also share important qualitative features of our framework. In each case there are individuals who would like to know what other agents know but are prevented from learning by the other agents' inaction. The private information of these other agents gets released to the wider market of interested individuals only when a threshold is crossed. The release of the information then leads to feedback effects as the interested individuals incorporate it into their own decisions.

B. *Modeling Strategy*

Our strategy is to abstract from the detailed settings described above and to focus on a simple environment that captures the essential common features of the examples. These features are (i) the presence of private information concerning a state variable of public interest, (ii) a process by which the private information evolves over time, and (iii) inertia that traps this private information until a critical threshold has been crossed.

To address these issues as simply as possible, we consider a setting in which a large number of firms in an industry undertake investment projects. Firms receive private signals concerning the profitability of invest-

ment in the industry. This private information evolves as the project progresses. In addition to their private information, each firm can observe whether the other firms in the industry are continuing to invest or have instead stopped investing and canceled their projects. We assume that investment is irreversible, so that firms are reluctant to reveal their pessimism by canceling their project until that pessimism has reached a critical threshold.

During the period before any projects are canceled, there is almost no information-sharing: this is the business-as-usual period in which each firm must decide whether or not to continue with its project largely on the basis of the information it receives locally. At some point in the evolution of the industry, there may be some firms that are so pessimistic that they are prepared to suspend operations rather than throw more good money after bad. These first cancellations transmit information to those that have stayed in the market and become a catalyst for further learning; in our model, this takes the form of a market crash.

C. *Literature Review*

There are several antecedents to the framework that we develop. The most direct is the (S,s) literature on aggregation with fixed adjustment costs, including papers by Alan Blinder (1981), Caplin (1985), Caplin and Daniel Spulber (1987), Ricardo Caballero and Eduardo Engel (1991), and Caplin and Leahy (1991). At the microeconomic level, fixed costs of adjustment lead to inaction. The aggregation literature analyzes how individual inaction influences aggregate dynamics. In the current model, firms are reluctant to scrap capital because investment is irreversible. As in the (S,s) literature, firms choose the optimal time to overcome their inertia and act. The key difference is that we add private information so that action releases information to the rest of the market. We then study how these actions and the information that they reveal interact in the aggregate.

A separate literature studies the connection between irreversible investment and

learning. Joseph Zeira (1987, 1990), Rafael Rob (1991), and Caplin and Leahy (1993) describe situations in which agents are ignorant of the profitability of the market into which they are investing. In such situations there is a tendency for investment to take place gradually as agents learn from the success or failure of earlier investments. In these models, all investors have the same information. Our focus in this paper is on the aggregation of information. We therefore require a heterogeneity of beliefs that is not needed in the study of gradualism.

A third related literature involves recent work on informational cascades by Abhijit Banerjee (1992), Sushil Bikhchandani et al. (1992), In Ho Lee (1992), and Ivo Welch (1992). These authors share our concern with the fact that information is dispersed and that agents learn from one another. They show how rational inference from the behavior of others can lead to follow-the-leader behavior, and how deviations from expected behavior can lead to drastic changes in market sentiment.

The main difference between our model and these other models arises from our concern with the endogeneity of decision times. In our view, it is impossible to explain the onset of a crash without considering why agents act when they do. Models of informational cascades skirt this issue by fixing the order of moves among market participants: agents with extreme beliefs must wait their turn before conveying their information to others. In contrast, we endogenize decision times so that those with the most extreme news are the first to reveal their information. With endogenous decision times, one must confront the issue of why all information is not revealed immediately. This concern explains our focus on the adjustment costs that rationalize inertia.

Finally there is a collection of papers that explore how markets may aggregate private information that is of public interest. These papers do not allow private information to evolve. They take differences of opinion as given and do not explain how such differences develop without being noticed at an earlier stage. Jeremy Bulow and Paul Klemperer (1991) examine this issue in the con-

text of an auction in which each agent has a fixed private valuation. David Romer (1990) considers how stock-market trading patterns may reveal private valuations. Christophe Chamley and Douglas Gale (1992) show how investment decisions may reveal private information in a strategic setting. Again, we believe that it is difficult to explain the timing of a crash without also explaining how private information evolves to the point at which a crash is possible.

We present the specific model of industry investment in Section I. In Section II we characterize equilibrium. Section III provides conditions under which an equilibrium exists and takes up issues of comparative dynamics and welfare. Section IV discusses the relationship between our specific model and the more general economic principles that underlie the relationship between information and adjustment costs. Our conclusions are presented in Section V.

I. The Model

We build a model of irreversible investment with time to build. Much of the model is standard; the novel feature is the information structure. We consider a continuum of firms, each of which has the opportunity to undertake a project that will produce a single unit of a consumption good. Firms must spread production of the good over T periods. During this time, each firm gathers information on the state of final demand and on this basis decides whether or not the additional investments needed to complete the project are worthwhile.

Figure 1 illustrates the production technology. At any time s , a project can be in one of three states: active, suspended, or canceled. In order to become active, a firm must pay an entry cost equal to $\kappa > 1$ in period zero. In period 1, a firm has two options. It can choose either to continue in the active state or to enter the suspended state. The decision to stay active involves the payment of an installment cost that we normalize to unity, while suspension is costless. In period 2, an active firm faces the same choice between continuation and suspension, whereas a firm that chose to sus-




FIGURE 1. THE PRODUCTION TECHNOLOGY

pend its project in period 1 must decide whether to cancel the project outright or to reactivate at a cost $\mu\kappa + 2$. Here, $\mu\kappa \in [1, \kappa]$ represents the costs of reinitiation in excess of the two installments that the firm would have paid had it never suspended its project.

In general, a firm that is active in period $s \leq T - 1$ chooses whether to continue or to suspend, whereas a firm that is in the suspended state chooses whether to cancel or to reactivate. The only exception to this rule occurs in the final period T when a firm in the active state chooses between finishing the project and cancellation. After this final decision has been made, the state of demand is realized, and firms with active projects sell their finished good at the market-clearing price.

The choices that a firm makes depend on its perception of the market. This perception is based on three different sources of information. The first is the firm's prior knowledge concerning the state of demand. We assume that *ex ante* there are two equiprobable demand states with zero demand in the low-demand state (L) and inverse demand function $P(Q)$ in the high-demand state (H). We assume that $P(0)$ is sufficiently high that initial entry is worthwhile.

The second source of information is the private information that a firm receives by virtue of its market position. We assume that each firm receives a private signal con-

cerning the state of demand immediately upon paying the initiation cost and each subsequent payment of the continuation cost. A firm that has continued through all prior periods and is about to make a period- s decision therefore had access to s private signals. We denote this s -vector of signals by σ^s . For analytic simplicity, we assume that a firm that suspends and reinitiates receives two signals upon reinitiation, so that it has the same number of signals as a firm that remains active throughout. We also assume that there is no new signal received before the final continuation decision in period T , so that the maximum number of signals is $T - 1$.

The signals themselves take a very simple form. A firm can receive one of two possible signals: a good signal or a bad signal. In each period, the probability that a firm receives the good signal when demand is high is equal to $p \in (0.5, 1)$, and the probability of receiving the good signal when demand is low is $q = 1 - p$. We assume that these signals are conditionally independent (given the true state of demand) over time and across firms. The parameter p may be thought of as a measure of the information content of the signal. If p is close to 1 then each signal reveals a great deal of information on the true state of demand, whereas if p is close to 0.5, each signal is relatively uninformative.

The final source of information is the observation of the past decisions of all other

firms. Before making its first continuation decision in period 1, a firm sees the total mass of firms that entered at time zero. Before making its period-2 decision, the firm sees the proportion of the initial entrants that decided to continue in period 1. Generally, in making its period- s decision, a firm knows the entire history of continuation, suspension, and cancellation for all periods up to and including period $s - 1$.

The information that a firm can derive from these observations depends on the firm's beliefs concerning the strategies pursued by other firms. A strategy π specifies a decision in each possible contingency that the firm may face. For any period $1 \leq s \leq T$, let \mathcal{I}_s denote the set of all of the possible period- s information sets with generic element $I_s \in \mathcal{I}_s$.² The typical period- s information set contains s signals, and the history of other firms' decisions through period $s - 1$. For each element $I_s \in \mathcal{I}_s$, the strategy π specifies a probability $c(I_s) \in [0, 1]$ of continuation if the firm is in the active state or a probability $z(I_s) \in [0, 1]$ of reinitiation if the firm is in the suspended state. This definition therefore allows for mixed strategies. We assume that any mixing is carried out independently across firms.

To complete the description of the model, we need to describe the way in which firms interact in the market. We assume that each individual project is infinitesimal relative to the market as a whole, so that each firm takes prices as given. We also assume that these firms are risk-neutral so that they maximize the expected value of profits. Finally, we assume that there is free entry, so that *ex ante* expected profits are zero. In this market setting, we look for a symmetric Nash equilibrium. In such an equilibrium, a strategy must be optimal given the information revealed when other firms pursue the same strategy.

² \mathcal{I}_s is the filtration generated by the private signals and the continuation and suspension decisions of others.

Definition: An *equilibrium* is a quantity of initial entry E and a strategy π such that:

- (a) The strategy π is optimal given full use of all information revealed when all others apply strategy π and the quantity of initial entry is E .
- (b) The strategy π and the initial entry quantity E give rise to zero expected profits as viewed from time zero.

In sum, the exogenous data of the model are the demand function $P(Q)$, the length of the project T , the cost of entry κ , the reinitiation-cost parameter μ , and the information content of each signal p . In the next two sections, we provide conditions on the data that ensure existence of an equilibrium strategy π . We show that all equilibrium strategies have a simple and intuitive form and that the amount of information that is revealed in equilibrium is simple to describe and varies in reasonable ways with the parameters of the model.

II. Characterization of Equilibria

To solve for the set of equilibria, we first examine the information that a firm can deduce from the actions of other firms when all firms follow a particular strategy π . Note that nothing can be learned from the observation of others until some firms choose to suspend. As long as all firms remain active, there is no way to distinguish between the behavior of a firm that has received positive information and a firm that has received bad news. Once the first suspensions have taken place, however, the information structure changes dramatically. Given that there is a continuum of firms receiving independent signals and following identical strategies, the proportion of firms that suspend is a deterministic function of the demand state. Provided the mass of firms that suspend is different in the high- and the low-demand state, these first suspensions reveal the true state of demand.

The only strategies for which the first suspensions do not reveal the state of demand are those in which the mass of suspensions is the same in the high- and low-

demand states. Given that it is always the most pessimistic firms that suspend and that there are more pessimistic firms in the low-demand state, the only way to generate an equal mass of suspensions is for all firms to suspend simultaneously. Such mass suspensions, however, are inconsistent with rationality. On the one hand, it cannot be rational for all of these suspensions to be followed by cancellation, since then the strategy makes it impossible to receive revenue and therefore guarantees negative profits. On the other hand, any firm that chooses to reenter does so without any new information and therefore would have been better off had it continued throughout and avoided the reentry cost $\mu\kappa$. The first suspension must therefore reveal the true state of demand in any equilibrium of our model.

Given that information is revealed in this manner, there is a simple process for identifying the set of equilibria. We consider some fixed time $t \in [1, T - 1]$ and explore the conditions for t to be a first-suspension time in some equilibrium.³

To develop the necessary conditions, we divide the investment period into three phases. We first analyze the optimal strategy for all periods $s > t$, the *fully informed phase*. We then use backward induction to determine the optimal strategy in period t itself, the *first-suspension time*. We then consider the optimal strategy in all periods $s < t$, the *private-information phase*. We complete the conditions for t to be a first suspension time with a free-entry condition. Finally, we consider the possibility that there is no suspension in equilibrium.

³In formal terms, we define the first-suspension time corresponding to π and E by

$$t = \min_s \{s \in [1, T] : c(I_s) < 1 \text{ for some } I_s \in \mathcal{I}_s(E)\}$$

where $\mathcal{I}_s(E)$ is the subset of \mathcal{I}_s corresponding to entry-level E . The fact that there is no new signal in period T and that final suspension is irrevocable rules out equilibria with first suspension in T , since suspending in period $T - 1$ and then quitting dominates suspending in period T .

A. The Fully Informed Phase

The optimal decision for a firm at times $s > t$ is trivial since all firms are fully informed about the true state of demand. If the state of demand is high, the firm will pay whatever costs are necessary to complete the project. If demand is low, it will not pay any additional continuation costs. Note that the willingness of a suspended firm to reenter and complete the project in the case of high demand follows from the assumption that the reinitiation cost $\mu\kappa$ is less than the entry cost κ . If such a firm were initially willing to pay the entry cost given a 50-percent chance that demand would be high, it is surely willing to pay the reentry cost at a time when high demand is certain and many of the continuation costs are sunk costs.

B. The First-Suspension Time

A firm in period t knows that the true state of demand will be revealed by the mass of firms that choose to suspend in that period. It knows that if demand is high, then it will pay all subsequent continuation costs, whereas if demand turns out to be low, it will incur no further costs. Given these expectations, the decision to suspend rather than to continue leads to an additional expense of $\mu\kappa$ if demand is in fact high, and a saving of 1 if demand is low. This implies that suspension is the superior strategy if and only if

$$(1) \quad \mu\kappa h(\sigma^t) \leq 1 - h(\sigma^t)$$

where $h(\sigma^t)$ is the probability that demand is high given the private signal vector σ^t . According to (1), a firm suspends its project only if it is sufficiently pessimistic. The cut-off probability for suspension is equal to $\bar{h} \equiv 1/(1 + \mu\kappa)$.⁴ Since we have assumed that $\mu\kappa \geq 1$, we know that some firms will always be optimistic enough to continue. The num-

⁴Unless (1) is satisfied with equality, mixed strategies will not be optimal.

ber of firms that choose to suspend in period t will therefore vary with the state of demand, and these suspensions will in fact reveal the true state of demand to all firms.

The first condition on the parameters necessary for t to constitute an equilibrium first-suspension time is that the set of firms that wish to suspend in period t is non-empty. Since $h(\sigma^t)$ is strictly increasing in the number of good signals, this requirement is satisfied if a firm with no good signals chooses to suspend. Applying (1) to a firm with t bad signals, we arrive at the following condition:

$$\left(\frac{q^t}{p^t + q^t} \right) \mu \kappa \leq 1 - \frac{q^t}{p^t + q^t}$$

or, defining $\lambda = p/q$,

$$(2) \quad \mu \kappa \leq \lambda^t.$$

It is clear that this inequality places a lower bound T_L on the set of possible equilibrium first-suspension times. The first-suspension time must occur late enough that pessimistic firms are willing to suspend and reveal information.

C. Free Entry

We can now use our knowledge of the form of the equilibrium strategy to calculate a firm's expected profits given that t is the first-suspension time. In equilibrium, the level of entry will ensure that these profits are zero. If t is the first-suspension time, each firm will pay initiation and continuation costs totaling $\kappa + t - 1$ with certainty. A firm will pay the remaining $(T + 1 - t)$ continuation costs and receive the market-clearing price P_H if the true demand state is high. The additional expenses that a firm can anticipate are the reinitiation costs that the firm pays if it suspends when demand is high and the additional continuation cost that the firm pays if it continues in period t even though the true demand state is low.

The free-entry condition is therefore

$$(3) \quad \frac{1}{2}P_H - \kappa - (t - 1) - \frac{1}{2}(T + 1 - t) \\ - \mu \kappa \cdot \Pr(\{h(\sigma^t) < \bar{h}\} \cap \{H\}) \\ - \Pr(\{h(\sigma^t) \geq \bar{h}\} \cap \{L\}) = 0$$

where all probabilities are *ex ante* probabilities. Equation (3) determines the equilibrium price in the high-demand state p_H . Given the inverse demand function $P(Q)$, it is now a trivial matter to solve for the initial entry quantity E :

$$P(E) = P_H.$$

D. The Private-Information Phase

We now consider the optimal decisions in the privately informed stage, $1 \leq s \leq t - 1$. In order for t to be an equilibrium first-suspension time, it is necessary that no firm prefer to suspend in any earlier period. We show that this condition provides an upper bound on the set of first-suspension times that is conceptually simple, but algebraically complex.

Given that all other firms continue through period $t - 1$, a firm that chooses to suspend in any period $s \leq t - 1$ can gain no new information before deciding whether or not to reenter in period $s + 1$. This means that suspension followed by reentry is dominated by paying the two continuation costs. In effect, the only alternatives at such a time s are continuation or outright exit. Therefore a necessary condition for t to be a first-suspension time is that for any time $s \leq t - 1$ and any sequence of signals σ^s , the expected value of continuing until time t is nonnegative.

There is one useful simplification to this condition. It suffices to check that the most pessimistic firm prefers to continue rather

than quit.⁵ Period t is therefore a first-suspension time only if

$$(4) \quad U(s, t) \geq 0 \quad \forall s \in [1, t-1]$$

where $U(s, t)$ denotes the expected value in period s of continuing through period $t-1$ and then behaving optimally for a firm with no good signals.

Our knowledge of the optimal strategy from period t on allows us to express $U(s, t)$ as follows:

$$(5) \quad U(s, t) = \frac{q^s}{p^s + q^s} [P_H - (T+1-t)] \\ - (t-s) \\ - \mu\kappa \cdot \Pr_s(\{h(\sigma^t) \leq \bar{h}\} \cap \{H\}) \\ - \Pr_s(\{h(\sigma^t) > \bar{h}\} \cap \{L\}).$$

The first term on the right-hand side in (5) is expected revenue less the continuation costs that all firms pay only if demand is high. Here p_H is determined from equation

(3), the free-entry condition. The second term is the expenditure through period $t-1$. The third and fourth terms capture the two types of errors that a firm can make. The third term is the probability that a firm with s bad signals in period s will suspend in period t in spite of the fact that demand is high. The fourth term is the probability that the firm will pay the continuation cost in period t when demand is in fact low.

An important qualitative feature of $U(s, t)$ is that it is monotonically decreasing in t , so that a firm with s bad signals does better continuing until the first-suspension time if this time comes sooner.⁶ The intuition underlying this monotonicity is simple: a firm that believes demand to be low would prefer find out the truth immediately, rather than pay further costs that it believes will be wasted. The monotonicity of $U(s, t)$ implies that the set of inequalities in (4) imposes an upper bound T_U on the set of possible equilibrium first-suspension times. The first-suspension time must be early enough that pessimistic firms are willing to wait for it.

E. The Case of No Suspension

To complete the analysis, we consider the possibility that there may be an equilibrium strategy π that involves no suspension. This case is particularly simple, since the equilibrium involves each firm choosing to con-

⁵This follows from the observation that the value of an optimal strategy is increasing in the number of good signals. Consider two firms, one with $i+1$ good signals and the other with i good signals. In calculating expected profits, each firm must consider four possibilities: that it will continue in period t when demand is high, that it will continue when demand is low, that it will suspend when demand is high, and that it will suspend when demand is low. Now one possible strategy for the firm with $i+1$ good signals is to mimic the decision of the firm with i good signals. In this case, both firms continue and suspend in period t with the same probability. The only remaining difference between the expected profits of these two firms is that the firm with $i+1$ good signals believes that the high state of demand is more likely. Since the high state is also more profitable (to see this, note that since the firm with i good signals can always quit the value of its optimal strategy must be positive, and since it receives no revenue its value in the low state is negative), the firm with $i+1$ good signals expects to earn higher profits following the strategy than does the firm with i good signals. Since the firm with $i+1$ good signals can only do better by following an optimal policy, its expected profits must be higher than the firm with i good signals. This argument indicates that it is sufficient that a firm with all bad signals choose to continue through period $t-1$.

⁶We thank one of our referees for a simple proof of this fact. Let $U(s, t, i)$ denote the value in period s of continuing until period t and then behaving optimally for a firm with i good signals. We show that $U(s, t, i) \geq U(s, t+1, i)$ by induction. Note that $U(t, t, i) \geq U(t, t+1, i)$ since $U(t, t, i)$ is the value of an unrestricted optimal policy. Hence the inequality holds for $s=t$. Note further that, for $t > S$,

$$U(S, t, i) = \Pr(g|i, S)U(S+1, t, i+1) \\ + \Pr(b|i, S)U(S+1, t, i)$$

where g and b indicate good and bad signals respectively. Therefore if the inequality holds for $s=S+1$ it holds for $s=S$. This completes the proof. The statement in the text corresponds to $i=0$.

tinue in each period solely on the basis of its own private information. The necessary and sufficient condition for this to be an equilibrium is that a firm with no good signals prefers continuation to exit in all periods. This is the case if the expected revenue is greater than the remaining continuation costs:

$$(6) \quad \left(\frac{q^s}{p^s + q^s} \right) P_H \geq (T + 1 - s)$$

$$\forall s \in [1, T - 1].$$

The free-entry condition allows us to substitute $p_H = 2(\kappa + T)$ and solve implicitly for the initial quantity of entry:

$$P(E) = 2(\kappa + T).$$

III. Analysis of Equilibrium

In this section we consider conditions on the data that ensure that an equilibrium exists. We first discuss the reasons why existence may fail and then show how our model overcomes these issues for a broad range of parameter values. We close the section with a discussion of comparative statics and welfare.

At the heart of our model is the possibility that individual agents will take expensive acts and thereby reveal useful information to others. Given the presence of a continuum of firms, there are always many firms with the same sequence of signals at any given first-suspension time. Having invested so much already, there is a natural tendency for a firm to want to continue one more period and learn the true state of the market from the actions of others without incurring the suspension cost itself. This possibility of free-riding places us very close to the original setting of Sanford Grossman and Joseph Stiglitz (1980) in which there is no equilibrium. Example 1 shows that the free-rider problem can indeed result in nonexistence of equilibrium in our model.

Example 1: Set $T = 17$, $\lambda = 1.4$, $\mu = 1$, and $\kappa = \lambda^{16} + \varepsilon$ for some small $\varepsilon > 0$. With κ

greater than λ^{16} , equation (2) implies that the only possible equilibrium is one in which no firms suspend. But in this case a firm that has had 15 bad signals prefers immediate exit to payment of the final three continuation costs:

$$\frac{P_H}{\lambda^{15} + 1} \approx \frac{2(\lambda^{16} + 17)}{\lambda^{15} + 1} \approx 2.9993 < 3.$$

Hence no equilibrium exists.

In terms of the equations of Section II, nonexistence occurs when the lower bound T_L is higher than the upper bound T_U . In such cases, it takes so long for firms to be willing to suspend and reveal information that some would rather quit the market outright at an earlier stage.

Fortunately, there are several forces at play that mitigate the effects of the free-rider problem in our model. One important factor is that, since firms that suspend have the opportunity to reenter, these firms are able to make use of the information that their actions help to release in the unlikely event that their suspension was a mistake. For this reason reductions in μ lessen the incentive to free-ride by reducing the costs of reentry. An examination of equations (2) and (4) shows that reductions in μ reduce T_L and increase T_U , thereby expanding the set of equilibrium first-suspension times. In the limit, with $\mu\kappa = 1$, there is always an equilibrium with suspension in period 1.

Another important factor is the way in which the private information sorts agents. Since a firm that suspends pays the reentry cost only if demand is high, the cost of revealing information is contingent on the state of demand. Since different firms have different views concerning the likelihood of the high-demand state, the cost of revealing information also differs across firms. Those who act to reveal the information are precisely those that least expect to pay the reentry cost. This provides a natural basis for a separating equilibrium in which optimistic firms continue and pessimistic firms suspend.

Finally, increases in the initiation cost κ increase the market-clearing price, which in turn reduces the cutoff probability for suspension. These effects tend to raise T_U and help to ensure the existence of equilibrium.

In Theorem 1 we provide a set of sufficient conditions on the data that guarantee existence of equilibrium. We show that existence is assured for any time horizon T if the individual signals are sufficiently informative, and it is ensured for any value of λ if the horizon is sufficiently short. The proof of Theorem 1 is in the Appendix.

THEOREM 1: Consider any set of data κ , T , p , μ , and $P(Q)$. Let t be the unique integer such that $\lambda^{t-1} \leq \mu\kappa < \lambda^t$. If $t > T-1$, then there is an equilibrium with no suspension. If $t \leq T-1$ and if either

(a) $T \leq 16$

or

(b) $\lambda \geq 1 + \mu/2$

then there is an equilibrium with first-suspension time t .

Conditions (a) and (b) have straightforward interpretations. Condition (a) places a limit on how far into the future first suspension can take place. It eliminates the possibility that pessimistic firms would be unwilling to wait for information to be revealed. Condition (b) places a lower bound on the informativeness of the signal. It circumvents the free-rider problem by ensuring that the most pessimistic firms are sufficiently pessimistic that they will want to suspend.

Conditions (a) and (b) are far from necessary for existence. For example, in the case of $\mu = 1$, we can manipulate equations (2), (3), and (4) to restrict further the region of nonexistence to the shaded region of the (T, λ) -plane in Figure 2. Outside the shaded area there exists an equilibrium for all values of κ and μ . Existence problems can only occur in the shaded region, and even then only for particular choices of μ and κ .

At a deeper level, we regard the possibility of nonexistence in our model as a computational rather than a conceptual issue. It




FIGURE 2. THE REGION OF NONEXISTENCE FOR $\mu = 1$

is an artifact of our assumption that there is a continuum of agents. At the heart of the existence problem is the fact that with a continuum of firms, the true state of demand is revealed in the first period in which there is a positive probability of suspension. If there were only a finite number of firms pursuing mixed strategies, the probability that other firms would suspend and reveal information could be kept arbitrarily low. The desire to free-ride would influence the particular strategies that firms choose, but would not endanger existence. While the finite case removes the existence problem, it radically increases the computational complexity of the analysis.

To understand better the quantitative implications of Theorem 1, we compute the entire set of equilibrium suspension times for a simple parameterized example.

Example 2: Let $\mu = 1$ and consider a three-year project with a two-month shutdown, so that $T = 17$. Assume that $\kappa = 8$ so that the initiation cost is 32 percent of total costs and the remaining 17 periods each involve 4 percent of total expenses. Finally, assume that at the end of a year in which the news is unremittingly bad, a firm is 80-percent sure that the project is going to fail, so that

$$\frac{p^6}{p^6 + q^6} = 0.8$$

or

$$\lambda = \sqrt[3]{2} \approx 1.25.$$

Given this data $\kappa = \lambda^9$. This rules out all equilibria in which suspension occurs before period 9. A mechanical check of the inequality in equation (4) reveals that the upper bound on equilibrium suspension times is 12. In the period-9 and period-10 equilibria only the firms with no good signals suspend, whereas in the period-12 equilibrium firms with one good signal or no good signals suspend. There are two possible equilibria with suspension in period 11 since firms with one good signal are indifferent between suspension and continuation. We conclude that the truth is revealed between 18 months and two years after the firms initiate their projects, even though some firms are 80-percent sure of their project's failure after only one year, and 96-percent sure after two years.

The comparative statics of the model are quite straightforward. All else being equal, the higher the initiation cost, the higher is the price in the high-demand state and the more bad news it takes to convince a firm to suspend or quit. An increase in the initiation cost therefore raises both the lower bound T_L and the upper bound T_U on the set of equilibrium suspension times and increases the range of parameters for which there is an equilibrium with no suspensions. In Example 2, a reduction in κ from 8 to 6 lowers T_L from 9 to 4 and lowers T_U from 12 to 9. Increasing κ to 12, on the other hand, raises the lower bound to 11 and raises the upper bound to 13, while a κ of 27 makes possible an equilibrium with no suspension.

Changes in p have similar effects. The more convincing is each individual signal, the fewer signals a firm needs to justify suspension. An increase in p therefore reduces T_L and T_U . In the foregoing example, p is equal to 0.55. Increasing p to 0.66 reduces T_L from 9 to 3 and reduces T_U from 12 to 4. Even though suspension takes




FIGURE 3. EFFECTS OF CHANGES IN μ ON T_U AND T_L IN EXAMPLE 2

place considerably earlier under this parameterization, firms still wait until they are nearly 90-percent certain that demand is low before choosing to suspend.⁷

Finally, it is trivial to note that for any given κ , reductions in μ unambiguously expand the set of equilibrium suspension times. This effect is illustrated in Figure 3, which shows how changes in μ affect T_U and T_L in Example 2. For example, substituting $\mu = 0.5$ into Example 2 reduces T_L to 6 and raises T_U to 13. The vertical line at 0.125 indicates that μ must be greater than $1/\kappa$.

We turn briefly to welfare issues. The natural comparison is between the equilibria in our model and an alternative equilibrium with full information-sharing. The main difference between these equilibria lies in the timing of the cessation of investment in the low-demand state. The failure to share information implies that resources are wasted on continuation when demand is low, so that expected costs are higher than necessary. This raises the equilibrium price in the high-demand state and reduces entry. It is therefore clear that an equilibrium with

⁷Note that for sufficiently high p revelation will occur in the first period. Generally, a necessary condition for delay is that pessimism not grow too quickly.

first suspensions in any period $t \geq 2$ is far from socially optimal. The alternative strategy in which all those that receive a first bad signal suspend with arbitrarily small probability $\varepsilon > 0$ reveals the information in time for the period-2 decision at arbitrarily low social cost and therefore dominates from the social viewpoint. It is also straightforward to confirm that the government can in principle ensure such sharing at arbitrarily low social cost by taxing first-period continuation for a small proportion of the population and dividing the revenues among all market participants.

IV. Overview and Extensions

In this section we go beyond the simple model of the paper and take a broader view of the mechanism connecting business as usual and information transmission. We have made radical simplifying assumptions on the nature of the signals, the channels for the communication of information between firms, and the market structure. We now discuss some of the consequences of relaxing these simplifying assumptions.

A. The Signal Process

Our assumptions on the signal process ensure that, if private signals were publicly observable, the true state of demand would be known in period 1. Moreover, if the true state of demand were revealed to be low in period 1, all firms would immediately quit. The main difference between an equilibrium with private information and a world with full information-sharing is therefore the *amount of time* that it takes for investment and the market to crash. The *fact* of the crash is exogenous to the model.

It is conceptually straightforward to amend the signal process so that even fully informed agents would only slowly learn the true state of demand. In such cases, complete information would frequently result in a smooth learning process, while business as usual and private information would give rise to discontinuous learning. Example 3 describes one method of introducing aggregate uncertainty without disturbing the

equations or the results of the paper in any essential fashion.

Example 3: Suppose that \hat{t} is an equilibrium first-suspension time in the model of Section I given the data κ , μ , p , T , and $P(Q)$. Now amend the model of Section I, so that in each period $s \in [1, \hat{t} - 1]$ the true state of demand switches with probability $1/(2\hat{t} - 2)$. If these switches are publicly observable, a firm's decisions in periods $s \geq \hat{t}$ are exactly as in Section II, and since the possibility that the state of demand will change makes a firm with bad signals less pessimistic, the decision to remain in the market until \hat{t} will be unchanged as well. Only with complete information is the evolution of the market different, since even if the news is initially bad some firms will be willing to continue in the hope that the market will improve.

Example 3 remains special in that the degree of aggregate uncertainty fades completely by the first suspension time. With more general signal processes, uncertainty concerning the true state of demand may only be fully resolved in period T . In such cases, we may expect firms in our model to suspend operations and reveal information in several distinct periods. After the first suspensions, there would be a new consensus on the probability that demand is high, and an appropriate number of quits. In the following periods, the firms that remain would gather new private information until the next suspensions occurred, whereupon the market would again aggregate information. Typically, there would be a number of discrete changes in market sentiment in our framework, whereas with complete information suspensions would take place gradually as knowledge of the true state evolved.

B. Channels of Communication

A critical assumption underlying the model is that there are limited channels for indirect information transmission among firms, so that "actions speak louder than words." We make the simplifying assumption that the only way that firms learn from

one another is through observation of market choices. In addition, we assume that the market choices take a particularly simple form: continuation or suspension.

With a larger choice set, the potential for information transmission increases. For example, if a firm were able to alter its intensity of production, its level of effort might convey additional information to others in the market. It is easy to see how effort could fully convey a single private signal. The presence of additional uncertainty, however, would again provide a role for suspension to reveal information, although the inference problem facing firms would become considerably more complex.

In practice, observation of a firm's actions is not the only way to learn about its views. Stock markets, futures markets, and direct markets for information all play a role in the aggregation of dispersed private information. Private contractual relations or collusion may also encourage the sharing of knowledge. We believe, however, that such channels rarely convey information perfectly, so that changes in behavior are still likely to be informative. In assessing the extent to which routine behavior suppresses private information in a given market, the central issue is to understand the limitations, if any, on the efficiency of these other channels. In part, this is an empirical issue: how is information actually aggregated? In part, it involves the theoretical task of explaining why private information remains hidden in any given situation. These are topics of current research.

Another important issue for future research to address is the cost of gathering and processing information. We have assumed that the signals arrive costlessly, and that the signals are easy to interpret. In practice, it is frequently difficult and costly to gather private information. Again, it is conceptually straightforward to alter the model to allow for such costs of information processing. In such cases, changes in the behavior of competitors may stimulate the gathering or processing of important new information, rather than aggregating information that is already known. On one level, this amendment would make little differ-

ence, since the behavior of the market would be very similar. A deviation from business as usual would trigger a discrete change in market sentiment. What would change is that it would be far more difficult to see how private markets could stimulate the optimal level of information-sharing.⁸

C. Market Structure

For analytic simplicity, we have assumed that there is a continuum of competitive firms. Moving to cases with a finite number of firms that operate strategically enriches the model in several important ways. One advantage of the strategic framework is that it removes the existence problem, as discussed in Section III. Another advantage is that strategic considerations may help to rationalize the lack of information-sharing that is assumed in our model. When the number of firms is small, all firms have an incentive to convince their competitors that demand is low, so that other firms will exit, and the market will become more profitable for those that remain. This incentive to lie reduces the effectiveness of direct channels of communication.

There is also a related strategic incentive for a firm to continue when its signals are bad: the hope that other firms will leave the market first. Such a war of attrition may further delay the aggregation of private information, and while this incentive does not rule out possible information-sharing, it suggests that any mechanism inducing truthful revelation would have to be more sophisticated than a simple opinion poll.

There are two other important phenomena that arise in a model with a finite number of agents. The first is the possibility of mistakes. In the current model, the crash always reveals the truth. With a finite number of agents, there is always the risk that there will be a preponderance of bad signals even in the good state, so that firms might

⁸In certain circumstances, it is possible that one way around this issue may be for agents to write contracts containing incentives for optimal gathering and sharing of information.

incorrectly infer that demand is low and leave the market. The second phenomenon is the possibility that the business-as-usual phase may be self-reinforcing. In our model, there is no learning from others until the first suspensions occur. In many situations, however, the fact that the other firms are continuing their projects may convey a positive signal, which may make firms more reluctant to suspend or may encourage new entry. With a finite number of firms, the probability that other firms have received good signals will grow the longer they remain in the market, thereby raising the probability that demand is high.⁹

V. Conclusion

We have presented a model of irreversible investment that connects the potential for a general collapse of investment to the difficulties of aggregating private information. The model illustrates a three-stage market dynamic. At the first stage, routine behavior tends to keep information of common interest trapped in private hands. At the second stage, this private information reaches a threshold that triggers some agents to take an action that releases their information to the market. The final stage involves a pattern of market evolution as other participants react to the initial departure from routine behavior.

We believe that this market dynamic sheds a common light on many apparently distinct economic phenomena. For example, we outlined applications of the framework to the international debt crisis, to bank runs, and to political upheavals.

APPENDIX

PROOF OF THEOREM 1:

Let $t \geq 1$ be the unique positive integer such that

$$(A1) \quad \lambda^{t-1} \leq \mu\kappa < \lambda^t.$$

⁹This kind of positive feedback is stressed by Charles Kindleberger (1989) and is present in the models of Rob (1993) and Caplin and Leahy (1993).

When $t > T - 1$ we show that there is an equilibrium with no suspension. In cases where $t \leq T - 1$ and either condition (a) or (b) is satisfied, we show that there is an equilibrium with first-suspension time t . The key simplifying feature of (A1) is that it implies that the only firms that suspend in the first suspension period are those that have received all low signals.

We first address the case in which $t > T - 1$. We look for an equilibrium with no suspension. A necessary and sufficient condition for such an equilibrium to exist is that all firms always find it in their interest to continue given that others are continuing. From Subsection II-E we have $p_H = 2(\kappa + T)$ and the conditions for continuation in all periods:

$$\frac{q^s}{p^s + q^s} 2(\kappa + T) \geq T + 1 - s \quad \forall s \in [1, T - 1].$$

The lower κ , the more likely this equality is to be violated. As $t \geq T - 1$ implies that $\mu\kappa \geq \lambda^{T-1}$, we replace κ with λ^{T-1}/μ , in which case the conditions can be rearranged as

$$(A2) \quad \lambda^{T-1} \geq$$

$$\max_{1 \leq s \leq T-1} \frac{\mu}{2} [(\lambda^s + 1)(T - s + 1) - 2T].$$

It suffices to show that if either condition (a) or (b) holds, then (A2) is satisfied for all $T \geq 1$. We delay the proof that (A2) holds for the appropriate range of parameter values until after we have taken up the case with suspension at some time $1 \leq t \leq T - 1$.

With $t \leq T - 1$, the zero-profit condition from the text is

$$(A3) \quad \frac{1}{2}P_H - \kappa - (t - 1) - \frac{1}{2}(T + 1 - t) \\ - \mu\kappa \cdot \Pr(\{h(\sigma^t) < \bar{h}\} \cap \{H\}) \\ - \Pr(\{h(\sigma^t) \geq \bar{h}\} \cap \{L\}) = 0.$$

The condition that a firm with $1 \leq s \leq t - 1$

bad signals be willing to continue to t is

$$(A4) \quad \frac{q^s}{p^s + q^s} [P_H - T - 1 + t] \\ \geq t - s + \mu\kappa \cdot \Pr_s(\{h(\sigma^t) \leq \bar{h}\} \cap \{H\}) \\ + \Pr_s(\{h(\sigma^t) > \bar{h}\} \cap \{L\}).$$

Substitution of (A3) into (A4) to eliminate $P_H - T - 1 + t$ and substitution for the probabilities gives

$$(A5) \quad \frac{q^s}{p^s + q^s} [2(\kappa + t - 1) + q'\mu\kappa + (1 - p')] \\ \geq t - s + \frac{q'\mu\kappa}{p^s + q^s} + \frac{p^s - p^t}{p^s + q^s}$$

The worst case for these inequalities is $\mu\kappa = \lambda^{t-1}$. Substitution and rearrangement reduces (A5) to

$$2\lambda^{t-1} + 2\mu t \\ \geq \max_{1 \leq s \leq T-1} \mu(\lambda^s + 1)(t - s + 1) + M$$

where the constant M is

$$M = (p - q)p^{t-s-1} \left(1 - \frac{1}{q^s}\right) \\ + (1 - p^{t-s})q^s - \frac{1}{q^s} < 0.$$

But now the series of inequalities follow if we establish that, for $t \geq 1$,

$$(A6) \quad \lambda^{t-1} \\ \geq \max_{1 \leq s \leq t-1} \frac{\mu}{2} [(\lambda^s + 1)(t - s + 1) - 2t].$$

Note that the inequality in (A6) is of precisely the same form as the inequality in (A2). Overall, we conclude that a time that satisfies all the necessary conditions from Section II surely exists if inequality (A6) is valid for all $1 \leq t \leq T - 1$. It remains to show

that either condition (a) or (b) in Theorem 1 is sufficient to ensure the validity of this series of inequalities.

There are several preliminary comments. Note that the inequality is trivially satisfied for $t = 1, 2$, so that equilibrium exists with $T = 2$ for any p, μ , and κ . More generally, the inequalities are simple to check computationally, and brute force was used to establish the assertion in part (a) of the theorem that there is always an equilibrium for $t \leq 16$. The conditions in part (b) are established analytically as follows.

The first observation is that, for any given data μ and λ , the set of necessary conditions for an equilibrium [that (A6) be valid for all $1 \leq t \leq T - 1$] gets more and more stringent as T is increased. This means that, if an equilibrium exists some given value of T , it also exists for all $T' \leq T$. Now consider the first time \bar{T} such that the series of inequalities in (A6) fails. We can make two relevant inferences:

$$(A7) \quad \lambda^{\bar{T}-1} \\ < \frac{\mu}{2} \max_{1 \leq s \leq \bar{T}-1} [(\lambda^s + 1)(\bar{T} - s + 1) - 2\bar{T}]$$

and

$$(A8) \quad \lambda^{\bar{T}-2} \\ \geq \frac{\mu}{2} \max_{1 \leq s \leq \bar{T}-2} [(\lambda^s + 1)(\bar{T} - s) - 2(\bar{T} - 1)].$$

Inequality (A7) follows from the fact that the first failure of (A6) occurs at time $\bar{T} - 1$, while (A8) follows from the fact that the (A6) is valid at $\bar{T} - 2$.

When the right-hand side of (A7) is evaluated at $s = (\bar{T} - 1)$, we have

$$\mu[\lambda^{\bar{T}-1} + 1 - \bar{T}] < \lambda^{\bar{T}-1}.$$

Therefore we can restrict the maximization in (A7) to $1 \leq s \leq \bar{T} - 2$. We can then substi-

tute (A8) into (A7) and derive

$$(A9) \quad \lambda^{\bar{T}-1} < \lambda^{\bar{T}-2} + \frac{\mu}{2} \max_{1 \leq s \leq \bar{T}-2} (\lambda^s - 1).$$

Noting that inequality (A9) is maximized at $s = \bar{T} - 2$, we finally derive the following necessary condition for \bar{T} to be the first time at which there is an existence problem:

$$(A10) \quad \lambda^{\bar{T}-1} < \left(1 + \frac{\mu}{2}\right) \lambda^{\bar{T}-2} - \frac{\mu}{2}.$$

Inequality (A10) immediately yields the condition $\lambda \geq 1 + \mu/2$ as sufficient to imply the existence of a time that satisfies all the necessary and sufficient conditions of Section II.

REFERENCES

- Banerjee, Abhijit V.** "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, August 1992, 107(3), pp. 797-817.
- Bikhchandani, Sushil; Hirshleifer, David and Welch, Ivo.** "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy*, October 1992, 100(5), pp. 992-1026.
- Blinder, Alan S.** "Retail Inventory Behavior and Business Fluctuations." *Brookings Papers on Economic Activity*, 1981, (2), pp. 443-505.
- Bulow, Jeremy and Klemperer, Paul.** "Rational Frenzies and Crashes." National Bureau of Economic Research (Cambridge, MA) Technical Paper No. 112, September 1991.
- Caballero, Ricardo and Engel, Eduardo.** "Dynamic (S,s) Economics." *Econometrica*, November 1991, 59(6), pp. 1659-86.
- Calomiris, Charles and Gorton, Gary.** "The Origins of Banking Panics: Models, Facts, and Bank Regulation," in R. Glenn Hubbard, ed., *Financial markets and financial crises*. Chicago: University of Chicago Press, 1991, pp. 109-74.
- Caplin, Andrew.** "The Variability of Aggregate Demand with (S,s) Inventory Policies." *Econometrica*, November 1985, 53(6), pp. 1395-1409.
- Caplin, Andrew and Leahy, John.** "State-Dependent Pricing and the Dynamics of Money and Output." *Quarterly Journal of Economics*, August 1991, 106(3), pp. 683-708.
- _____. "Sectoral Shocks, Learning, and Aggregate Fluctuations." *Review of Economic Studies*, October 1993, 60(4), pp. 777-94.
- Caplin, Andrew and Spulber, Daniel.** "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics*, November 1987, 102(4), pp. 703-25.
- Chamley, Christophe and Gale, Douglas.** "Information Revelation and Strategic Delay in a Model of Investment." Boston University Working Paper No. 10, 1992.
- Gorton, Gary.** "Banking Panics and Business Cycles." *Oxford Economic Papers*, December 1988, 40(4), pp. 751-81.
- Grossman, Sanford J. and Stiglitz, Joseph E.** "On the Impossibility of Informationally Efficient Markets." *American Economic Review*, June 1980, 70(3), pp. 393-408.
- Kindleberger, Charles.** *Manias, panics and crashes: A history of financial crisis*. New York: Basic Books, 1989.
- Lee, In Ho.** "Market Crashes and Informational Cascades." Mimeo, University of California-Los Angeles Graduate School of Management, 1992.
- Lohmann, Susanne.** "Rationality, Revolution and Revolt: The Dynamics of Informational Cascades." Mimeo, Stanford University Graduate School of Business, 1992a.
- _____. "The Dynamics of Regime Collapse: A Case Study of the Monday Demonstrations in Leipzig, East Germany, 1989-1991." Stanford University Graduate School of Business Research Paper No. 1225, 1992.
- Mishkin, Frederic.** "Asymmetric Information and Financial Crises: A Historical Perspective," in R. G. Hubbard, ed., *Financial markets and financial crises*. Chicago: University of Chicago Press, 1991, pp. 69-108.
- Rob, Rafael.** "Learning and Capacity Expansion under Demand Uncertainty." *Review of Economic Studies*, July 1991, 58(4), pp. 655-75.

- Romer, David.** "Rational Asset Price Movements without News About Fundamentals." Mimeo, University of California-Berkeley, 1990.
- Sachs, Jeffrey.** "Developing Country Debt," in M. Feldstein, ed., *International economic cooperation*. Chicago: University of Chicago Press, 1988, pp. 233-77.
- Welch, Ivo.** "Sequential Sales, Learning, and Cascades." *Journal of Finance*, June 1992, 47(2), pp. 695-732.
- Zeria, Joseph.** "Investment as a Process of Search." *Journal of Political Economy*, February 1987, 95(1), pp. 204-10.
- _____. "Informational Cycles." Mimeo, Hebrew University, 1990.