

Sidelined Investors, Trading-Generated News, and Security Returns

H. Henry Cao

University of North Carolina

Joshua D. Coval

Harvard University

David Hirshleifer

Ohio State University

This article studies information blockages and the asymmetric release of information in a security market with fixed setup costs of trading. In this setting, “sidelined” investors may delay trading until price movements validate their private signals. Trading thereby internally generates the arrival of further news to the market. This leads to (1) negative skewness following price run-ups and positive skewness following price rundowns (even though the model is *ex ante* symmetric), (2) a lack of correspondence between large price changes and the arrival of external information, and (3) increases in volatility following large price changes.

The turn in sentiment was astonishing to watch. . . Everyone was waiting for the less-sophisticated investors to come in and panic and dump stock. They didn’t, and when there wasn’t that panicked sell response, there was a panicked *buy* response from these ‘sophisticated’ buyers.

—Kevin Bannon, Chief Investment Officer at Bank of New York,
WSJ 10/29/97

Recent empirical evidence (discussed below) has shown that after recent price run-ups stock returns are negatively skewed, so that stock prices are more likely to have large drops.¹ Conversely there is a tendency for positive

We thank Tony Bernardo, Kent Daniel, Allen Potesman, Tony Sanders, Tyler Shumway, Steve Slezak, Avandhar Subrahmanyam, Sheridan Titman, Pietro Veronisi, Jiang Wang, Guojun Wu, seminar participants at Ohio State University, Tel Aviv University, the 1997 Financial Economics and Accounting Conference at SUNY Buffalo, and the Society for Financial Studies 2000 Conference on Market Frictions and Behavioral Finance at Northwestern University, and Bob Korajczyk (the editor) and an anonymous referee for helpful comments and discussions. We are particularly grateful to Seongyeon Lim for insights which have significantly improved this article. This article was written while the first author was at the Fisher College of Business and the second author was at the University of Michigan Business School. Address correspondence to Joshua D. Coval, Harvard Business School, Soldiers Field, Boston, MA 02163, or e-mail: jcoval@hbs.edu.

¹ These are conditional *skewness* effects, and hence are distinct from the conditional mean “winner/loser” effects described by DeBondt and Thaler (1985, 1987).

skewness conditional upon a recent decline in stock price. Furthermore, there is evidence that volatility tends to cluster in asset returns, with large price changes following large changes.

These patterns of path-dependent shifts in volatility and in the asymmetry of return distributions pose a challenge to rational theories of securities pricing. It is far from clear why external news would arrive in a skewed pattern as a function of past external news. Nor is it obvious why external news should arrive more or less rapidly as a function of past news.

We argue in this article that such conditional patterns can arise naturally in an ex ante symmetric setting with fixed transaction costs,² which result in blockages in the pipeline that transfers individual information to the market. In our setting this pipeline intermittently clogs and then releases information in large lumps. Central to our approach is that information in the hands of *sidelined investors* is not immediately incorporated into the market price, and events in the trading process trigger participation and incorporation of the news such investors possess. Since trading events endogenously trigger the arrival of news, the distribution of future returns shifts predictably as a function of past price movements even when no new external news arrives. We offer a specific model that builds upon previous literature (discussed in Section 2) to show how such gradual revelation can generate the observed patterns in the conditional moments of security returns.

In principle, heterogeneous investors should constantly participate in all securities markets, instantaneously rebalancing as information arrives.³ If the world were really like this, market makers would see everyone's orders in every security at every moment, and so would have an excellent idea of what information is possessed by investors. As a result, prices would aggregate information very well.

In reality, most individuals and institutions trade any given stock only intermittently, if at all. This can be true even of stocks that many people have opinions about. There is much more water-cooler conversation about Yahoo than actual trading. Many people have strong views about the stock market as a whole, yet do minimal market timing. There is anecdotal support for the idea that traders consider the views of sidelined investors to be important, and that at certain specific points in time the market expects rapid resolution about the beliefs of sidelined investors. The quote heading this text is one example.⁴

The failure of most individuals to do at least some trading in stocks can be viewed as resulting from a fixed cost. The relevant cost could be broker-

² Throughout the article, in referring to transaction costs, we have in mind fixed setup costs of trading, which could include costs of attention or of mobilizing to make a trade.

³ Under certain special assumptions (such as normal distributions and existence of a riskfree asset) two-fund separation applies, but these assumptions are at best only approximations of reality.

⁴ Reuters (2/27/00) asserted that "Stock market bulls are praying for a 'buy-on-the-dip' rally this week after one of Wall Street's oldest and most beloved gauges, the Dow Jones industrial average, fell below the 10,000 watermark."

age fees for odd lots, or more importantly the time and attention required to explore, decide, execute, and track a trade. We provide a model which reflects the obvious but important fact that at any given time there are many sidelined investors not trading in a given stock. We explore here the premise that sidelined investors possess relevant information about security value,⁵ and how the entry of sidelined investors may be triggered by recent price movements.

Famous empirical findings of momentum and long-run return reversals involve path dependence in the *first* moments of stock return distributions [see, e.g., DeBondt and Thaler (1985) and Jegadeesh and Titman (1993)]. Several theoretical articles have derived such path dependence in first moments.⁶ In contrast, our model focuses on information aggregation and the path dependence of second and third moments.

A number of recent studies document remarkably strong variation in skewness of security returns as a function of past price movements. Following large price increases, return distributions become negatively skewed, with an increased probability of a large downward correction. Likewise, following substantial price declines, returns become positively skewed, with little chance of further large declines, and an increased probability of what traders sometimes refer to as a “dead-cat bounce.”

Using transactions data, Ederington and Lee (1996) provide evidence suggesting that such a phenomenon exists at a high frequency level.⁷ Using monthly data, Harvey and Siddique (1999) document considerable time variation in conditional skewness measures of the U.S. and world market portfolios.⁸ Lo and Wang (2000) (in their Table 6c) report evidence that is suggestive of weekly conditional skewness.⁹ In work which provides an ex post

⁵ The effects we model would still occur if many investors were mistaken in their belief that they have information.

⁶ Wang (1993) demonstrates theoretically that asymmetric information may partially account for the observed serial correlations in returns. He shows that less-informed investors depend heavily on past returns in forming their estimates of future returns, and as a result, amplify any existing mean reversion in underlying state variables. Llorente et al. (1999) show that return autocorrelations can depend on the nature of the trading activity with which they are associated. They find that positive autocorrelation tends to be generated by speculative trade, while negative autocorrelation is generated by allocational trade. Brown and Jennings (1989) derive conditions under which technical analysis is profitable; competitive traders develop more accurate estimates of a security's payoff by using a weighted average of first- and second-period prices than second-period prices alone.

⁷ They demonstrate that price changes of futures contracts on currency, bond, and stock indices become highly skewed following steady price rises or declines. For example, they show that if Treasury bond futures experience a sequence of four positive (negative) price changes, they are four times as likely to record a further price increase (decrease) (after adjusting for bid-ask bounce). Yet if the sequence is broken by a reversal, the price decline (rise) is 2.5 times as likely to be followed by a further decline (rise) as by a rise (decline).

⁸ The null of constant skewness is consistently rejected at the 1% significance level. Indeed, in an asset pricing framework in which investors have a preference for skewness, they show that conditional skewness helps account for the time-series variation in the expected U.S. and world market returns as well as the cross-sectional variation in the expected returns of individual securities.

⁹ They find that stocks whose returns one week are in the top (bottom) 10% are much more likely to end the following week in the bottom (top) 10% than any other decile. This very high probability for a shift to the

test of one of our model implications, Chen, Hong, and Stein (1999) recently reported that high monthly returns are associated with lower (more negative) skewness after controlling for other determinants of skewness.

Furthermore, there is evidence that volatility is clustered over time. Large stock price moves are associated with high future volatility. Schwert (1990) examines a century of U.S. stock index data and finds that volatility jumps following large market changes and typically persists for many weeks thereafter. GARCH (generalized autoregressive conditionally heteroscedastic) models, with their ability to capture conditional volatility clustering in asset returns, have been successfully applied to modeling risk premia in foreign exchange markets [see Engle and Bollerslev (1986)] and other securities [see Bollerslev, Engle, and Wooldridge (1988)].

This article constructs an information theoretic framework in which patterns of conditional skewness and heteroscedasticity in asset returns arise from the buildup and release of pockets of information temporarily hidden from most participants.¹⁰ As such, the model places greater emphasis than most securities market theory on the distribution of information across informed individuals rather than just the distinction between informed and uninformed. Shiller (1995) has emphasized that limitations in the effectiveness of conversation in conveying information may be a source of stock market volatility. Our model suggests that transfers of information between individuals may often have no effect, yet in some circumstances can trigger large stock price movements.

Since the analysis in this article includes a market maker, it is tempting to interpret the model in terms of high-frequency (intraday) events. This is one level of application, but we emphasize that in practice information can be sidelined for periods of days, weeks, or longer. We offer some basic stylized implications of sidelined information and do not attempt to calibrate the model according to different possible time horizons. However, we do explore the robustness of the results to alternative market-clearing mechanisms. We solve the model under both a microstructure market-clearing process (in which only a single order is cleared per period) and a batched order market-clearing process that may be more reflective of longer-horizon price-setting dynamics.

The remainder of the article is organized as follows. Section 1 describes the basic idea in relation to the literature. Section 2 lays out the model. Section 3 describes conditional patterns of trading and skewness when markets clear

extreme opposite decile contrasts with a much weaker tendency for switches to other opposing deciles (e.g., from decile 1 to decile 9). This evidence is suggestive of conditional skewness, although it is possible that a combination of conditional mean reversal and cross-sectional heteroscedasticity could explain this pattern.

¹⁰ There are alternative possible reasons for conditional skewness. Since limited liability equity is much like a call option on underlying firm value, a symmetric distribution for firm value changes would imply asymmetry in stock price changes. Such asymmetry should fluctuate over time as stock price and leverage changes [see Black (1972), Christie (1982), and Nelson (1991)]. A similar point applies to volatility shifts [see, e.g., Charoenruek (2000)].

based upon batched orders. Section 4 describes conditional patterns in trading, skewness, and volatility when orders are cleared sequentially. Section 5 concludes.

1. The Basic Idea and Previous Literature

Our approach is based on the idea that market imperfections, such as trading costs or a psychic cost of starting to trade in a stock affect trading strategies in a way that can conceal information. There is a probability that a set of investors receive a common informative signal. Observing price changes helps an individual infer whether his signal is valid. To be willing to trade an investor may sometimes need his information confirmed by the actions of other informed traders before he is confident that he can recover his fixed setup costs of trading.

After a price increase a favorably informed investor who is considering whether to buy faces a balance of effects. On the one hand the security is more expensive to buy, but on the other hand the validity of his information is confirmed. From the perspective of a favorably informed trader, the adverse price move may not eliminate the net gain from trading because the price is revised by an uninformed market maker. The informed investor concludes that the latest price move was probably a result of a similarly informed trader rather than a liquidity trader; the market maker is less sure of this. Thus the price rise triggers trading on the part of the favorably informed investor. In contrast, adversely informed traders become less confident that they have received correct signals, and may thereby be sidelined.

This informational overhang does not require investors to trade in opposition to their private signals (although our model accommodates this possibility). Those investors whose signals are consistent with the price move tend to trade in the direction of that move, and investors whose information opposes the price move tend to choose not to trade. Because of this body of sidelined investors, trading events can stimulate the arrival of news to the market.

This sidelining of investors causes conditional shifts in volatility and skewness as a function of past returns. After an upward price trend, most likely there are only a few sidelined investors with opposing signals, in which case the price is likely to rise moderately further. However, there is a modest probability that a large number of investors with adverse signals are sidelined. If so, their eventual entry will result in a major correction. The market maker, upon receiving a sell order, partially foresees this possibility and revises the price significantly downward. Thus security price changes have path-dependent distributions, becoming negatively skewed following price rises and positively skewed following declines.

In this way the combination of setup costs and uncertainty about the veracity of private signals produces highly uneven, path-dependent, and asymmetric price responses to symmetric arrival of new information. In particular,

a stock's prospects will depend heavily on the extent to which preceding trading has allowed prices to reflect investors' aggregate information set.

The variation in market participation rates has an important additional effect. Large corrections give rise to an increase in price volatility that lingers for several periods. Consider the case where prices have increased for several periods, successfully aggregating bullish viewpoints. At this point there will exist an "information overhang" of unaggregated bearish viewpoints in the market. Investors understand this, but they may later be unpleasantly surprised to discover there were many hidden bears, or relieved to learn that there were only a few. If a correction occurs, the market becomes highly uncertain as to how many "bears" will be prompted to enter the market—that is, how large the overhang is that the correction will release.

A further point of interest is that setup costs, on a per trade basis, can improve the market's aggregation of private information. This is because the setup costs act as a filter on market participation in a way that allows price movements to accurately test the market's information set.¹¹

This article extends the intuition of three important recent articles involving market imperfections: Romer (1993), Avery and Zemsky (1998), and Lee (1998). Romer (1993) offers a pair of models which demonstrate that, consistent with empirical evidence, even if investors are rational, substantial price movements can occur during periods when there is relatively little arrival of new information.¹² In his second model (the relevant one for our purposes), Romer shows that with relatively small transaction costs, prices can reflect fundamentals poorly. In his model an investor must pay a cost to have a positive probability of being permitted to trade immediately based on private information. The decision of whether to incur the cost is incurred *before* the investor receives his signal. Alternatively, he may opt to delay and then trade for free (if he is permitted to trade). Romer shows that the benefits of trading immediately can be quite small, resulting in slow impounding of information into price, and subsequent price movements without news. One difference of our article from Romer's (and the other articles) is that our focus is on the relation of volatility and skewness to preceding market events.

Analytically, whereas a trader in Romer's model must pay a fixed setup cost before receiving any information, in our model a trader must pay this cost whenever he chooses to trade. Thus an investor in our model not only

¹¹ Recall that informed investor participation rates respond sharply to price changes, with more purchasers participating as prices begin to rise and more sellers entering when prices begin to decline. This gives rise to significant imbalances in the participation rates of informed buyers and sellers and brings about the conditional patterns of skewness mentioned above. But since noise traders purchase and sell with equal probabilities, the imbalances in informed trader participation allow the market maker to infer better whether he is trading against informed or uninformed investors.

¹² The stock market crash of October 1987 provides a possible example. More generally, Cutler, Poterba, and Summers (1989) find that seldom are large index changes associated with clear news announcements. Roll (1988) finds that firm-specific price changes are also hard to explain based on public news events. Finally, French and Roll (1988) find price volatility to depend more on when the market is open than on the rate of information arrival.

learns from earlier traders about whether his signal is correct, he also *weighs what he has learned against the cost of trading* on his improved information. As a result, in our model only an investor with significantly improved confidence that his private signal is correct will choose to trade. This creates path dependence in conditional skewness—returns become negatively skewed after a stock price run-up and positively skewed after a rundown.

Avery and Zemsky (AZ) show that a price rise can encourage even an investor with an adverse signal to buy when there is a transaction cost (or, as in AZ, a bid-ask spread). This is because the price rise can persuade the investor that others possess favorable information, and the market maker may react sluggishly to this good news. AZ provide the important insight that this can occur if the investor has an informational advantage over the market maker along a second dimension of uncertainty—uncertainty over *whether* informative signals were sent.¹³ The informed traders realize that information signals were sent so that the previous order probably came from a favorably informed trader, whereas the market maker places greater weight on the possibility of a liquidity trade.

Along a sample path in which even investors with adverse signals buy in the face of rising prices, eventually price rises are not accompanied by significant favorable revisions by informed traders about underlying value. As market volume dries up, the market maker shrinks bid ask spreads and locked-up adverse information can be released by sell orders. If the amount of revealed adverse information is greater than expected, a crash can result. In contrast with AZ, in our model the possibility of asymmetric price drops does not require a prelude in which investors trade in opposition to their own signals. Instead, it is caused by the sidelining (nontrading) of informed investors. In addition, we explicitly analyze how shifts in volatility and skewness are related to past price movements.

Lee (1998) examines a setting in which transaction costs limit information aggregation, leading to possible market crashes. In his setting, large stock price drops are triggered by rare, highly adverse signals. He demonstrates that a sequence of trades by traders with positive signals can push a security's price up to the point where subsequent traders with high probability choose not to trade. This drying up of trade makes the market vulnerable. As a result, it is possible that large numbers of subsequent adverse signals accumulate yet do not get incorporated into the price. A crash is triggered when a trader with an intense, rare realization is selected to trade. This can cause the release of hidden adverse information.

Lee's article does not, however, directly examine the nature of and extent to which market imperfections cause conditional patterns of skewness and volatility. Even frictionless models are consistent with conditional skewness and volatility arising from the assumed discreteness of states [see, e.g.,

¹³ Easley and O'Hara (1987) term this second dimension of uncertainty "event uncertainty."

Veronesi (1999)]. To examine the effects of market imperfections, in this article we compare the cases of positive and zero setup costs to derive conditions under which setup costs do indeed increase conditional skewness. In our article the signal distribution is symmetric, and disproportionate market movements are triggered by normal signal values.

More broadly, the preceding literature has focused primarily on the possibility of slow incorporation of information and of market crashes, whereas we attempt to derive specific predictions for shifts in skewness and volatility in relation to past price moves and volatility.

2. The Model

Consider a setting with a continuum of risk-neutral investors. Each period each investor makes a decision of whether to submit a buy order for one share, a sell order for one share, or not to submit an order for the single risky security. (Since there is a continuum of investors, each share represents an infinitesimal quantity.) The orders are completed by a competitive risk-neutral market maker. At the beginning of time investors receive noisy signals about the security's true value. The market maker quotes a price at which he is willing to purchase or sell the security. The market maker receives no signals and trades costlessly. In contrast, investors face positive setup costs.

We analyze two alternative market-clearing processes. In the first process, all informed trader demand is pooled with liquidity demand and the market maker clears the net demand at the preannounced price. This approach closely resembles a multiperiod version of Kyle (1985) except that the informed traders here are nonstrategic. In the second process, the market maker randomly draws a single order from the total pool of informed and liquidity demand. This setup parallels that of Glosten and Milgrom (1985), except that we do not model the bid and ask spread here.

2.1 Payoff and signal structure

The risky asset has terminal payoff $V = 0$ or 1 (see Figure 1). In order to illustrate how conditional skewness effects arise in an *ex ante* symmetric setting, we assume these outcomes have equal probability. With probability γ investors receive no information prior to trading. With probability $1 - \gamma$, at date 0 a subset of investors (informed traders) receive conditionally independent signals about the terminal value. Each signal is conveyed as either 0 or 1 and is equal to the true value V with probability p , the signal accuracy.

This signal structure has the effect that when an investor receives a signal, he can infer that signals have been sent to other traders as well. However, the signals received by the others may differ from his. His knowledge about whether signals were sent gives him an advantage over the market maker. The fact that different informed traders see different signals implies that he can learn new things by observing market price.




Figure 1

Model summary

With probability $1 - \gamma$, a subset of investors (informed traders) receive signals about the terminal value. Each signal is conveyed as either 0 or 1 and is equal to the true value V with probability p . We label the state in which the terminal value is 1 and signals are sent as state h_1 . The state where the terminal value is 0 and signals are sent is labeled state l_0 . When the terminal values are 0 or 1 and no signals have been sent, the states are labeled h' and l' , respectively.

To motivate this signal structure, consider first the feature that an investor who receives a signal can infer that signals were sent to others. This is reasonable since in reality there is some commonality in information arrival. Examples may be a subset of traders who have access to analysts who have recently met with management, or to suppliers of or buyers from the firm who have useful information about its performance. Another example is a subset of traders who subscribe to a private analyst's newsletter.

The further point is that the signals received are not necessarily identical. The motivation is that information signals can sometimes rationally be interpreted differently by different investors. (Implicitly we have in mind a setting where there is interaction between different kinds of information and investors begin with a diversity of prior information.) For example, suppose that an investor learns that Pawheld Inc. is developing a new personal digital assistant for dogs. The investor may understand that a small set of other investors have also learned about the firm's planned investment initiative. But some investors can rationally interpret this as positive news (if they think the project is promising), and others that it is negative news. Thus each informed investor will understand that there is a set of other informed investors out there, but he can't be sure whether they regard the information as favorable or unfavorable. As a second example, consider a subset of investors who have just learned that the firm is planning on laying off workers. Some investors may interpret this as a "red flag" for low product demand, others as an indication that the firm has discovered a good way to improve productive efficiency.

The behavior of the system depends on whether signals are sent as well as on terminal payoffs. We label the state in which the terminal value is 1 and signals are sent as state h . The state where the terminal value is 0 and signals are sent is labeled state l . When the terminal values are 0 or 1 and no signals have been sent, the states are labeled l' and h' , respectively.

2.2 Traders

There is a continuum of risk-neutral investors. M is the proportion of information-based traders, some of whom receive signals about V if signals are sent. Let b be the proportion of the M traders who actually receive a signal when it is sent (whom we will call "informed traders").

We assume that if an information-based trader does not receive a signal, he does not submit any order at all.¹⁴ At time t , N_t is the proportion of liquidity traders; they do not receive signals.

¹⁴ This assumption in effect places an upper bound on b , so that a trader who does not receive a signal is virtually no better informed than the market maker, and hence does not trade. When no signals are sent, the market maker will overadjust the price in response to order flow (from the perspective of information-based traders who know that signals were not sent). Such traders may trade against price movements based on the information that no signals were sent. We rule this out for tractability, but the intuition for the result does not rely on this assumption.

Informed traders must pay a fixed setup cost in order to trade. These costs vary uniformly across the population from 0 to 1, so the fraction of the population with transaction costs $C \leq c$ is

$$F(c) = c. \quad (1)$$

Liquidity traders, on the other hand, are compelled to place orders in each round. We compare outcomes to a benchmark scenario in which informed traders face zero costs of participating.

Finally, the trades are executed in a market of competitive, risk-neutral market makers. We model this as a single, representative risk-neutral market maker. For simplicity we assume that the market maker quotes a single price at which he is willing to buy or sell the security. Clearly this assumption results in the market maker earning negative average profits (since he breaks even against the liquidity traders and loses, on average, to the informed traders). Similar results would apply, with less tractability, under the assumption that the market maker charges the minimum bid-ask spread needed for him to break even on average on his trades.¹⁵

2.3 Equilibrium trade and prices

Before trading takes place, the terminal payoff is determined and any signals are sent to investors. The security is then traded for T rounds. Each round of trade consists of four stages:

1. The market maker announces a price at which he is willing to purchase or sell per share of the security.
2. Each investor decides whether to submit a buy or sell order for a share of the security.
3. We consider two alternative market-clearing processes. In the first, the market maker trades with the batched pool of informed and liquidity orders. In the second version, the market maker randomly draws a single order from the pool of all submitted orders.
4. Traders and the market maker observe the net demand (in the first version) or the selected order (in the second version) and update their beliefs accordingly.

Let π_{st} be the market-maker's assessed probability that an information event has occurred and the state is s at the beginning of round t . Then the price P_t announced at the beginning of round t will be

$$P_t = \pi_{ht} + 0.5(1 - \pi_{ht} - \pi_{lt}), \quad (2)$$

¹⁵ Alternatively, we can think of the transaction cost as being paid as compensation to market makers for their services.

where with probability π_{ht} the terminal payoff is 1 and with probability $1 - \pi_{ht} - \pi_{lt}$ no signals have been sent, so that the expected terminal payoff is equal to its unconditional value 0.5.

Let π_{ht}^σ be the probability assessed by an investor at the beginning of date t who has received signal σ that the state is h . Since the asset pays either 0 or 1, an investor's expectation of the terminal value will be his perceived probability of state h , π_{ht}^σ . When faced with a fixed setup cost C , a risk-neutral investor with signal 1 will place a buy order if $C \leq \pi_{ht}^1 - P_t$. Likewise, an investor with signal 0 will sell if $C \leq P_t - \pi_{ht}^0$. Among investors who have received signal σ , the proportion who place orders in round t is defined as θ_t^σ . By Equation (1), this can be expressed as

$$\theta_t^\sigma = |\pi_{ht}^\sigma - P_t|. \quad (3)$$

In the no-transaction cost case, since all traders participate, $\theta_t^\sigma = 1$.

We will see that the sidelining of informed investors leads to three main effects. First, informed traders will on average tend to buy in response to price rises and sell in response to declines. Second, as noted in the introduction, price changes exhibit path-dependent, conditional skewness. Third, there is conditional heteroscedasticity.

Intuitively, price movements that confirm only one set of investors' signals create another group of sidelined investors. The question facing market participants is how many "bears" and "bulls" are on the sidelines. Sidelined investors are deterred by the setup costs and await a possible subsequent confirming price move (or moves) before returning to the market. The potential for a greater-than-expected number of sidelined investors to return to the market at some point creates the possibility of a large correction. As a result, future price changes become highly skewed. (The intuition for conditional heteroscedasticity is provided in Section 4.)

3. Market Clearing with Batched Orders

In the batched orders version of the market-clearing process, all participating traders' orders are executed each period and the market maker observes the net order flow. For tractability in this section we consider a three-period version of the model. Recalling that p is signal accuracy, net informed trader demand in date t can be written as

$$I_{ht} = p(\pi_{ht}^1 - P_t) - (1 - p)(P_t - \pi_{ht}^0) \text{ if } s = h, \text{ and} \quad (4)$$

$$I_{lt} = (1 - p)(\pi_{lt}^1 - P_t) - p(P_t - \pi_{lt}^0), \text{ if } s = l. \quad (5)$$

With zero transaction costs, informed demand is simply

$$I_{ht} = pS(\pi_{ht}^1 - P_t) - (1 - p)S(P_t - \pi_{ht}^0), \text{ if } s = h, \text{ and} \quad (6)$$

$$I_{lt} = (1 - p)S(\pi_{lt}^1 - P_t) - pS(P_t - \pi_{lt}^0), \text{ if } s = l, \quad (7)$$

where $S(\cdot)$ is the sign function which takes value 1 when the argument is positive, -1 when the argument is negative, and 0 if the argument is zero.

Next, let D_t represent the actual net demand observed in date t :

$$D_t = I_t + u_t, \quad (8)$$

where u_t is the net liquidity demand in date t . We assume that the proportion of liquidity traders is distributed with probability density $2\phi(x)$, $x > 0$, where ϕ is the normal density function. In addition, we assume that the liquidity traders' trades are perfectly correlated, and they buy or sell one share per trader with equal probability. This implies that the net liquidity demand u_t is normally distributed with mean zero and variance one. This approach is similar to that of Kumar and Seppi (1992).¹⁶

Finally, we determine how investors and arbitrageurs update assessments of state probabilities. At the end of each round of trade the only new information being revealed to market participants is the net demand. If the date t demand is D_t , investors update their beliefs about states based on their observation of demand using Bayes' rule, yielding

$$\pi_{h,t+1} = \frac{\pi_{ht}f(D_t - I_{ht})}{\pi_{ht}f(D_t - I_{ht}) + \pi_{lt}f(D_t - I_{lt}) + (1 - \pi_{ht} - \pi_{lt})f(D_t)}$$

$$\pi_{l,t+1} = \frac{\pi_{lt}f(D_t - I_{lt})}{\pi_{ht}f(D_t - I_{ht}) + \pi_{lt}f(D_t - I_{lt}) + (1 - \pi_{ht} - \pi_{lt})f(D_t)},$$

where $f(x)$ is the normal pdf.

Let us define *price chasing* by a group of investors as higher average net buying (or lower net selling) when the price has risen more.

Proposition 1. *If the market maker is sufficiently uninformed (i.e., γ is sufficiently small) and transaction costs are uniformly distributed, informed traders chase the date 1 price change.*

Proof. See the appendix.

Further illustration of the relation of price movements to informed demand is provided in Figure 2, which plots the path of net informed trader demand that results from a sequence of liquidity demands when the true state is h . The solid line depicts informed demand when the aggregate liquidity demand in the first two periods is $(2, -2)$, whereas the dotted line depicts informed demand when liquidity demand is $(-2, 2)$. The graph illustrates how informed trader demand responds to price changes. In the graph, net

¹⁶ Equivalently we could fix the mass of the liquidity traders at one but have each liquidity trader i submit an order of $u_t + \epsilon_{it}$, where u_t is the common component with normal distribution and ϵ_{it} is independently and normally distributed across liquidity traders. By the law of large numbers, the net liquidity demand is u_t . See Diamond and Verrecchia (1982), Carmel (1995) and Cao (2001) for such an approach.




Figure 2

Informed demand with positive transaction costs

The graph plots the net informed demand in state h with liquidity demand in the first two periods is set equal to 2 then -2 (solid line) and -2 then 2 (dotted line).

informed trader demand increases after a price increase and decreases after a price drop. When the price rises, traders with favorable signals increase their demand substantially, whereas those with adverse signals tend to be sidelined. Traders with favorable signals interpret the price increase as a verification of their signals, whereas traders with adverse signals view it to be disconfirming. Conversely, when prices decline, traders with favorable signals move to the sidelines, while those with adverse signals submit sell orders more aggressively. When there are no transaction costs, informed traders' demand is unrelated to past price moves.

The two panels of Figure 3 plot the price path that results from a sequence of liquidity shocks with (left panel) and without (right panel) transaction costs when the true state is h . The solid line shows the price path when the sequence of liquidity demands in the first two dates is $(2, -2)$; the dotted line shows the price path with a liquidity demand sequence of $(-2, 2)$. Figure 3 shows that the price process is path dependent with respect to the arrival of




Figure 3

Price paths with positive and zero transaction costs

The two graphs plot the price path in state h under positive (left panel) and zero (right panel) transaction costs. Liquidity demand in the first two periods is set equal to 2 then -2 (solid line) and -2 then 2 (dotted line).

liquidity demand. With transaction cost, the arrival of positive demand sidelines informed traders with adverse signals. Similarly the arrival of negative liquidity demand causes informed traders with favorable signals to be sidelined. When the order of the liquidity demand is reversed, the net effect does not reverse completely. This is because the information contained in each round of trade changes, depending on the relative participation rates of the informed traders. In contrast, when there is no transaction cost, the ultimate price is not affected by the order of arrival of the liquidity demand. This is because traders participate in the market all the time, and therefore each transaction is equally informative about the underlying distribution of signals.

To examine conditional skewness, we investigate the relation between the date 1 price change (which is perfectly correlated with the date 1 price) and the skewness of the date 2 price change. The following proposition provides the date 2 conditional skewness calculation.

Proposition 2. *When γ is near zero, the conditional skewness of the date 2 price change can be expressed as*

$$\frac{E[(P_2 - P_1)^3 | P_1]}{(E[(P_2 - P_1)^2 | P_1])^{3/2}}, \quad (9)$$

where

$$\begin{aligned} E[(P_2 - P_1)^3 | P_1] = & (\pi_{h1})^3 e^{3(I_{h1})^2} - (\pi_{l1})^3 e^{3(I_{l1})^2} + 2(\pi_{h1} - \pi_{l1})^3 \\ & - 3(\pi_{h1})^2 \pi_{l1} e^{[(2I_{h1} + I_{l1})^2 - 2(I_{h1})^2 - (I_{l1})^2]/2} \\ & - 3(\pi_{h1})^2 (\pi_{h1} - \pi_{l1}) e^{(I_{h1})^2} \end{aligned}$$

$$\begin{aligned} &+ 3(\pi_{l1})^2 \pi_{h1} e^{[(2I_{l1}+I_{h1})^2-2(I_{l1})^2-(I_{h1})^2]/2} \\ &- 3(\pi_{l1})^2 (\pi_{h1} - \pi_{l1}) e^{(I_{l1})^2} \\ &+ 6\pi_{h1} \pi_{l1} (\pi_{h1} - \pi_{l1}) e^{I_{l1}I_{h1}}, \text{ and} \\ E[(P_2 - P_1)^2 | P_1] &= (\pi_{h1})^2 e^{(I_{h1})^2} + (\pi_{l1})^2 e^{(I_{l1})^2} - (\pi_{h1} - \pi_{l1})^2 - 2\pi_{h1} \pi_{l1} e^{I_{l1}I_{h1}}. \end{aligned}$$

Proof. See the appendix.

Figure 4 plots the difference between the skewness of date 2 price changes [as given by Equation (2)] under uniformly distributed transaction costs and




Figure 4
Conditional skewness with batched orders

The probability of receiving an accurate signal in state h or l is graphed on the x -axis as p and the standardized period 1 price change is graphed on the y -axis as $\Delta P_1'$. The z -axis graphs the difference in skewness of the period 2 price change between uniformly distributed transaction costs and zero transaction costs. The figure displays the negative relationship between period 1 price changes and period 2 skewness in the presence of transaction costs.

under zero transaction costs.¹⁷ The difference in skewness is plotted against the date 1 price level standardized by the date 1 standard deviation (denoted as $\Delta P_1'$) and the signal precision (p). Figure 4 indicates that following a date 1 price declines, skewness is uniformly positive. As the price change increases, skewness shifts from positive to negative, with skewness uniformly negative following date 1 price increases. In other words, there is a decreasing relationship between date 1 price changes and the skewness of date 2 price changes.

Intuitively, conditional on a price increase, it becomes more likely that the true state is high. Thus if the true state is actually low, when the true state is revealed there is a large correction. Such a large downward correction is more likely when there is a positive transaction cost, because the transaction cost sidelines many investors with adverse information, while accommodating more participation by investors with favorable information. The market price is set rationally, but there remains an asymmetry between a low probability that the market is very unfavorably surprised and a high probability that there is some slight favorable resolution. To put this another way, the sidelining of investors makes the expected imbalance between buys and sells subsequent to a price rise greater. This makes negative order flow more informative, so that with positive transaction costs a large negative price correction is more likely. As a consequence, the conditional skewness is more negative.

The difference in skewness subsequent to high versus low past returns is most pronounced for high levels of the signal precision (p). This pattern suggests that higher precision increases the amount of sidelining of informationally disfavored traders (i.e., it sidelines optimists when the price declines, and sidelines pessimists when the price increases). Why does high precision increase this sidelining? It would seem that there is a trade-off. On the one hand, to the extent that high precision is reflected in market prices, investors need to place more weight upon the market price, which increases sidelining. On the other hand, a higher precision directly encourages an investor to place more weight on his own signal, which should encourage participation. It is not obvious which effect must dominate. However, high precision also makes prices less noisy, since the proportion of participating informed traders increases relative to that of liquidity traders. While the high precision discourages trading by investors whose information is opposed to the price move, it encourages trading by those whose information is consistent with the price move, so that overall high precision still encourages trading. Thus, as signal precision rises, the effect of more precise market prices (which

¹⁷ Since the payoff space is discrete, and the order flow has a mixed normal distribution, after observing the first period order flow the conditional distribution of the payoff is no longer symmetric, even when the transaction cost is zero. Thus there will be some conditional skewness even with zero transaction costs. However, the skewness induced by discreteness is small. More importantly, we examine the incremental effect of introducing transaction costs on skewness, above and beyond the pure effect of discreteness.

discourages trading in opposition to prices) outweighs the direct effect of greater individual precision (which encourages trading).

The skewness calculations we focus on in the batched order version of the model relate to date 2 price changes. Intuitively, similar skewness-inducing forces should operate at longer horizons as well. If the price continues to rise over several trading dates, it would, up to a point, sideline an even higher fraction of traders with adverse information signals, and accommodate the participation of an even higher fraction of favorably informed traders. Thus the probability of a price drop next period is even lower, but in the event that it occurs, it should be correspondingly more severe.¹⁸ Thus the intuition underlying the date 2 result seems to extend to multiple periods. The next section examines a setting with sequential orders in which we examine the effects of nonparticipation over many periods and patterns of conditional heteroscedasticity as well as skewness.

4. Market Clearing with Sequential Orders

In this section we consider a second market-clearing setting where the market maker draws a single order each period from the pool of informed and liquidity orders. The market maker only observes the order he draws; he obtains no further information regarding the composition of the order pool.¹⁹

4.1 Beliefs and trades

The indicator variable ϕ_t^σ is set equal to one when a trader's expected valuation of the risky asset based upon signal σ is larger than the current price, and is set to zero otherwise:

$$\phi_t^\sigma = \begin{cases} 1 & \text{if } \pi_{S_t}^\sigma > P_t \\ 0 & \text{if } \pi_{S_t}^\sigma \leq P_t. \end{cases}$$

¹⁸ However, once prices increase to the point where they became too expensive for traders with favorable signals to participate, the skewness will drop as the imbalance between participation rates subsides. How soon this occurs will depend on the sluggishness of the market maker in updating prices (i.e., the level of γ). This suggests that longer-horizon skewness will be associated with information events that are more infrequent, from the perspective of the market maker.

¹⁹ In reality a market maker observes more than the executed order. In our setting, this would provide the market maker with more information, which would help him infer whether signals were sent. In the limit, if a market maker observes an infinite set of orders, he may be able to infer perfectly whether signals were sent, which would eliminate the effects analyzed here. Realistically, however, in most markets only a small fraction of the population makes orders at any moment in time, which limits the information available to the market maker. This stylized fact corresponds in our model to a nonnegligible average fixed setup cost of trading. Our assumption that market makers know perfectly the conditional probability that each investor will receive a favorable signal (given the true value) also makes the inference problem for the market maker easier in our model than in a realistic setting. If instead these probabilities are not known perfectly by the market maker, then the market maker would not be able to infer perfectly from order flow whether signals were sent, even if he observed many orders. It is also worth noting that in markets that have multiple market makers, a given market maker does not observe all orders submitted to the market.

This implies that for a trader with signal σ , proportion $\theta_t^\sigma \phi_t^\sigma$ will purchase. Let q_{st} be the probability that the market maker executes a buy order in round t if the state is s . This probability is the fraction of all orders submitted to the market at a given time that are buy orders. This is

$$\begin{aligned} q_{ht} &= \frac{[\theta_t^1 \phi_t^1 p + \theta_t^0 \phi_t^0 (1-p)]Mb + 0.5N}{[\theta_t^1 p + \theta_t^0 (1-p)]Mb + N} \\ q_{lt} &= \frac{[\theta_t^1 \phi_t^1 (1-p) + \theta_t^0 \phi_t^0 p]Mb + 0.5N}{[\theta_t^1 (1-p) + \theta_t^0 p]Mb + N}. \end{aligned} \quad (10)$$

Since no information traders exist in states h' and l' , $q_{h't}$ and $q_{l't}$ are both 0.5; the only market participants are liquidity traders who purchase or sell with equal probability. The fraction of liquidity traders can be made arbitrarily small, as long as the market maker remains unable to distinguish between the signal and the nonsignal states. In the numerical simulations, we assume that the number of liquidity traders $N \ll Mb$.²⁰ This simplifies Equation (10) to

$$\begin{aligned} q_{ht} &= \frac{\theta_t^1 \phi_t^1 p + \theta_t^0 \phi_t^0 (1-p)}{\theta_t^1 p + \theta_t^0 (1-p)} \\ q_{lt} &= \frac{\theta_t^1 \phi_t^1 (1-p) + \theta_t^0 \phi_t^0 p}{\theta_t^1 (1-p) + \theta_t^0 p}. \end{aligned} \quad (11)$$

We next determine how investors and arbitrageurs update assessments of state probabilities. At the end of each round of trade the only new information being revealed to market participants through price is whether the executed trade was a buy or sell order. Hence participants will use Bayes' rule to update their beliefs. If the period- t trade is a buy, the updated probabilities are

$$\begin{aligned} \pi_{h,t+1} &= \frac{\pi_{ht} q_{ht}}{\pi_{ht} q_{ht} + \pi_{lt} q_{lt} + (1 - \pi_{ht} - \pi_{lt})/2} \\ \pi_{l,t+1} &= \frac{\pi_{lt} q_{lt}}{\pi_{ht} q_{ht} + \pi_{lt} q_{lt} + (1 - \pi_{ht} - \pi_{lt})/2} \\ \pi_{h,t+1}^\sigma &= \frac{\pi_{ht}^\sigma q_{ht}}{\pi_{ht}^\sigma q_{ht} + (1 - \pi_{ht}^\sigma) q_{lt}}. \end{aligned}$$

²⁰ The trades of liquidity traders are dominant in nonsignal states. The effect of this assumption is to make informed traders remain dominant in signal states. This assumption is for algebraic simplicity, and is not necessary for the main results derived here.

If the period- t trade is a sell, the updated probabilities are

$$\begin{aligned}\pi_{h,t+1} &= \frac{\pi_{ht}(1 - q_{ht})}{\pi_{ht}(1 - q_{ht}) + \pi_{lt}(1 - q_{lt}) + (1 - \pi_{ht} - \pi_{lt})/2} \\ \pi_{l,t+1} &= \frac{\pi_{lt}(1 - q_{lt})}{\pi_{ht}(1 - q_{ht}) + \pi_{lt}(1 - q_{lt}) + (1 - \pi_{ht} - \pi_{lt})/2} \\ \pi_{h,t+1}^\sigma &= \frac{\pi_{ht}^\sigma(1 - q_{ht})}{\pi_{ht}^\sigma(1 - q_{ht}) + (1 - \pi_{ht}^\sigma)(1 - q_{lt})}.\end{aligned}$$

Hence the above updating rules, in conjunction with the expressions for prices, Equation (2), participation rates, Equation (3), and buy probabilities, Equation (10), uniquely determine how the market evolves in response to a sequence of orders.²¹

4.2 An illustrative example

In Figure 5, each node shows the price P , the state probabilities π_h and π_l , and market participation rates θ^0 and θ^1 (t subscripts omitted) which follow a particular sequence of buy and sell orders. Prior probabilities are set to unconditional levels [i.e., $\pi_{h0} = \pi_{l0} = (1 - \gamma)/2$]. The probability of nonsignal states, γ , is set to 0.9 and the fraction of signal recipients who receive correct signals, p , is set to 0.75. The first tree depicts the case with fixed setup costs; in the second they are absent. To see the conditional skewness of price changes that arises in the presence of setup costs, compare the sequence {BUY, BUY, BUY} to that of the sequence {BUY, BUY, SELL}. At $t = 2$, after two BUYs, the price has increased to .5483. At this point, traders who have received a signal of 1 are more than 40 times as likely to participate in the market as those who received a signal of 0 (.365 versus .008). Hence the probability of a subsequent increase in price is above 58% [.317/ (.317 + .227)]. On the other hand, if a SELL order does execute, the decline in price is substantial—the price drops below its unconditional mean to .4983. Moreover, further declines now become likely, as sellers are more than six times as likely as buyers to participate at this point.²²

A second effect of introducing fixed setup costs is that prices become more volatile in earlier trading rounds, meaning that the market aggregates information *more* rapidly per order. For example, consider the order sequence

²¹ In order to distinguish the effects that we focus on from those which derive from investors trading against their own private signals, in addition to our basic model we also examine an artificial scenario in which the trades of participating investors are constrained to be consistent with their signals. Specifically, if an investor finds it optimal to trade in the opposite direction of his signal, in the artificial scenario he is instead constrained at that point in time to trade zero. Liquidity traders are assumed always to participate in the market and place buy and sell orders with equal probability. We have verified that the results in the artificial scenario are similar to those of the actual model.

²² Because of the discrete state space, even in the zero setup cost case the market exhibits some conditional skewness, but it is much less. The probability of a BUY after two previous BUY orders is only 52.5% [.269/ (.269 + .243)].




Figure 5
Order flow tree diagram ($\gamma = 0.9, p = 0.75$)

{BUY, SELL, BUY}. In the zero transaction cost setting, this sequence of orders does little to resolve uncertainty. State probabilities remain similar to their unconditional levels. In contrast, with fixed setup costs, uncertainty is largely resolved following such sequences. The probability of a nonsignal state is now 99.7%, and the probability of states h or l has fallen from 10% to 0.3%. Setup costs advance resolution and volatility to earlier periods, and improve the market's per trade aggregation of information. This is surprising since the setup costs lock some pockets of informed traders out of the market.

Since fixed trading costs have the effect of sidelining investors and thereby temporarily concealing their information, it is perhaps to be expected that the

overall effect would be a reduction in the efficiency with which the security market aggregates private information. On a per-trade basis, in this setting the opposite is the case. By creating temporary imbalances between informed versus noise trading, fixed setup costs actually allow the market maker to make better inferences, increasing market efficiency.

The reason for this is that setup costs filter market participation in a way that allows for accurate tests of the market information set.²³ Following a BUY, investors receiving signals of 1 are 15 times as likely to participate as those receiving signals of 0. Hence a subsequent SELL is a strong indicator that the true state is not 1. If the true state is h (so signals are sent), the probability that the next order is a BUY is 97.8% $((.75)(.375)/[(.75)(.375) + (.25)(.025)])$. Hence when a SELL occurs in period 2, market participants become fairly certain that the true state is not h . As a result, investors who possess a signal of 1 move to the sideline, while investors with signals of 0 reenter the market. Now, a further reversal (BUY) will be a strong indicator that the true state is not l , as the possibility of seeing a BUY in state l is 3.6% $((.25)(.043)/[(.25)(.043) + (.75)(.382)])$. As a result, in period 4, after observing the sequence {BUY, SELL, BUY}, the market maker becomes highly certain (99.7%) that he is not facing informed investors and that signals have not been sent.

The final effect of time-varying market participation is to create conditional heteroscedasticity. Figure 6 depicts the price paths that follow a sequence of five BUY orders and then a SELL. Parameter values remain equal to those used in Figure 5 ($\gamma = 0.9$ and $p = 0.75$). Note that with setup costs, the SELL order triggers a market crash of 29.7% (from .7088 down to .4981), whereas the no-transaction cost decline is 9.5% (from .6476 to .5858). At this point, in the no-transaction cost setting, the price reflects a reduction of probability of state h from .296 to .174. The probability of state l remains low, and further possible price declines are of limited size and probability.

With setup costs, the crash is much larger, as the price decline reflects not only a great lowering of the probability of state h , but also a large increase in likelihood of state l . At this point there is considerable uncertainty regarding the true state. In fact, even though only one SELL has been observed in the first six rounds, state l is now more likely than state h . This is because the initial BUY, having banished bearish traders with 0 signals to the sidelines, makes the sequence of subsequent BUY orders highly likely. The surprising arrival of the SELL order, however, casts severe doubt on the true state being h , since most traders with 0 signals (who would be scarce if the true state were h) are still on the sidelines. At this point investors are highly uncertain as to how many bears the SELL order will attract back into the market. The probability of two further SELL orders is quite high [.295] and the

²³ See Fishman and Hagerty (1992) for a model that involves a somewhat analogous filter benefit from restrictions on disclosure.




Figure 6
Order flow tree diagram following five buys ($\gamma = 0.9$, $p = 0.75$)

price decline these would bring is substantial (from .4981 to .3963). This uncertainty about how many sidelined investors had contrary signals induces the increase in variance following a reversal.

4.3 Simulation results

To see that these results generalize, Figures 7 and 8 graph the conditional skewness and conditional variance that accompany different levels of setup costs in a market with 100 rounds of trade. For each parameterization, the market is simulated 1000 times. Each of the graphs plot the changes relative to a benchmark case that controls for changes in conditional moments which will arise simply due the discrete state space or the zero transaction cost




Figure 7
Conditional skewness with sequential orders

The probability of receiving an accurate signal in state h or l is graphed on the y -axis as p and the probability of nonsignal states is graphed on the x -axis as γ . The z -axis graphs the change in skewness of returns conditional on an increase in price during the prior trading round (less the change in the no-transaction cost case). The variance of liquidity trades is arbitrarily small.

conditional skewness effects we have assumed.²⁴ Finally, the simulations are conducted under the general scenario in which traders are permitted to trade against their signals. Highly similar results are obtained when traders are restricted from trading against their signals.

4.3.1 Conditional skewness as a function of past price moves. Figure 7 shows the change in conditional skewness, that is,

$$[Skew_t - Skew_{t-1}] - [Skew_t^0 - Skew_{t-1}^0],$$

²⁴ Specifically, in the benchmark case the market maker receives purchases with probability p in state h , $1-p$ in state l , and 0.5 otherwise.




Figure 8
Conditional variance

The probability of receiving an accurate signal in state h or l is graphed on the y -axis as p and the probability of nonsignal states is graphed on the x -axis as γ . The z -axis graphs the increase in variance that arises if five like price changes are followed by a reversal (relative to that of the no-transaction cost case). The variance of liquidity trades is arbitrarily small.

where a superscript of 0 refers to the zero transaction cost case, conditional on the prior price change being positive. The x -axis (right axis) of the graphs measures the probability of a nonsignal state (γ)—that no signals have been sent and the terminal payoff is either 0 or 1 with expected value 0.5. The y -axis (left axis) reflects the accuracy of the signals when signals have been sent (p). When setup costs are positive, asset returns uniformly exhibit increased negative skewness following a price rise. The increase in skewness is most pronounced when p is large and when γ is high. This mirrors the result for the conditional skewness under the batched order version of the model depicted in Figure 4.

This is because conditional skewness tends to be encouraged by low correction probabilities. If investors view the probability that their signals are

informative as low, and they view the amount of information in informative signals to be low, participation rates will not respond drastically to price changes. Investors will not shift this assessment much in response to price changes. For higher levels of p , when few investors have received incorrect signals, corrections occur with lower frequency (and involve correspondingly large price movements). For instance, with a large p if a sequence of purchases were observed, the arrival of a sell order would be highly unlikely. Since few adverse signals are sent if the true state is h (in which signals are sent), a sell order would virtually rule out the possibility of state h . As a result, in the unlikely event that a sell order was observed, the price would necessarily fall below one-half. When γ is high, the imbalance between buyers and sellers at any given time can be large. Because the order flow is far less informative to the market maker than to the informed traders, he does not update prices aggressively. As a result of the sluggishness of price moves, investors with signals that were confirmed by price moves participate extensively. Likewise, investors with unconfirmed contrary signals move to the sidelines even if their setup costs are fairly modest. Hence the arrival of informative signals makes orders highly autocorrelated. Thus the decline in price that a contrarian order brings is substantial.

Another point of interest is that setup costs can cause uncertainty to be resolved more quickly. In simulations not reported here, we find that per-trade resolution of uncertainty (as measured by return variance over the first 20 periods of trade) is more rapid when there are positive setup costs.

The reason that setup costs can advance per-trade resolution is that during states when signals are sent, the sidelining of investors with contrary signals helps the market maker identify that an imbalance exists between buyers and sellers and that signals have been sent. In the absence of setup costs, all investors remain in the market, regardless of how confident they are in their signals. Hence when there are no setup costs, and the imbalance between buyers and sellers is moderate, the market maker requires many observations of trade to determine the underlying distribution of information. On the other hand, when setup costs influence trader participation rates, the market maker can learn more quickly whether an information event has occurred. Since the sidelining of traders increases the autocorrelation in order flow, this helps the market maker learn that an information event has occurred and hastens information aggregation.

4.3.2 Postcorrection volatility. We next examine how volatility changes after a market reversal.²⁵ Figure 8 plots the increase in variance which follows a significant correction in the presence of transaction costs. The change in

²⁵ Our finding that volatility rises after a price reversal is developed in the sequential order version for the model. A similar intuition applies in the batched order version of the model, but the analysis is intractable as it involves numerical analysis of triple integrals.

variance is measured relative to the benchmark case and is conditional on the prior five price changes being of identical sign followed by a reversal, that is,

$$[var_t - var_{t-6}] - [var_t^0 - var_{t-6}^0].$$

Here we see that after corrections, price changes tend to be far more volatile in the presence of transaction costs. This increase in volatility is also most pronounced for intermediate signal quality p . With intermediate signal quality, there exists considerable potential for a large number of bears to reenter the market following a crash. This possibility increases uncertainty following corrections, and this uncertainty lingers several periods. When signals are of extremely high quality, traders with incorrect signals rarely push prices in the wrong direction and sideline traders with accurate information. Therefore reversals are rarely associated with the threat of large numbers of contrarians reentering the market. Likewise, when signals are of extremely low quality, five consecutive price changes of like identical sign will not compel a sufficient number of traders with contrary signals to move to the sideline for their reentry following a reversal to be important.

Our implication that volatility increases after market reversals has not been directly tested. The evidence of volatility clustering mentioned in the introduction is suggestive, but GARCH models generally do not distinguish between volatility increases which follow reversals of large run-ups (or run-downs) and those that simply follow periods of large changes in either direction.²⁶

4.3.3 Sample path realizations. Figure 9 plots price paths and participation rates for simulations of the market. In all four panels the state is h and the accuracy of signals (p) is set to 0.75. In the upper two panels, the probability of nonsignal states (γ) is set to 0.9, while in the latter two it is set to 0.999. The first panel plots a price path which is common when the state is h . The initial buy order deters sellers from joining the market and encourages buyers to enter. Further purchase orders arrive, as the price begins to rise, and the fraction of buyers remains high. At a certain point even buyers begin to leave the market rather than purchase at an increasingly expensive price. Some sellers reenter during periods 4–10 in order to short the asset at a high price that retains some possibility of correcting. After a while they become convinced that the market has correctly aggregated information, and they exit the market along with the buyers.

The second upper panel demonstrates how corrections, and their induced shifts in market participation rates, can cause wide swings in the asset price. In period 5, when a sell order arrives, the price not only corrects, but induces

²⁶ The success of exponential GARCH models [see Nelson (1991)], which allow volatility to be an exponential function of past shocks, does highlight the strong effect particularly large shocks have on future return volatility.




Figure 9
Simulations of price paths and participation rates

The above panels simulate prices and participation rates in the market when the state is h . Participation rates are scaled ($\times 2$) to match the price scale. The amount of noise trading is arbitrarily small; $p = 0.75$; in the upper two panels $\gamma = 0.9$ and in the lower two $\gamma = 0.999$.

a dramatic shift in the composition of market participants from buyers to sellers. This leads to a period of sustained selling and a continued decline in the price. This pattern repeats every few periods until a sustained sequence of purchases begins to move the price toward the terminal value.

The third panel (on bottom) depicts the situation where a sell order is initially executed even though the true state is h and therefore sellers only comprise 25% ($1 - p$) of market participants. This sale causes a massive exit of buyers and entry of sellers. As a result, many subsequent sell orders are executed, and the price declines to less than 10% of its terminal value. Eventually when the price drops so far that sufficient sellers have exited the

market and the pool of orders has become more balanced, a buy order is executed. This brings about a sustained period of rising prices and buyers returning to the market. After 15 purchases, the price has converged to the terminal value.

The final panel shows that even when prices have neared the terminal value, unexpected orders can bring about sustained periods of divergence. Subsequent to period 28, a sequence of sell orders causes the price to drop from more than 0.95 to less than 0.4. This is because the initial sell orders were so unexpected that they brought on a sustained period of reentry on the part of sellers.

5. Summary, Extensions, and Empirical Testing

This article examines a setting in which fixed setup costs of trading cause investors' market participation to vary through time depending on the past history of prices. Low participation blocks the information pipeline from individuals to the market. In our setting this pipeline temporarily clogs and then releases information intermittently in large lumps. Sidelined investors wait for confirming price changes before they are sufficiently confident in their information to enter the market. Thus the trading process itself endogenously triggers the arrival of further news.

As a result, the arrival of trading-generated news triggers the aggregation of investors' information into price in a way that is related to past price movements. Specifically, return distributions exhibit conditional skewness, becoming more negatively skewed following positively trending prices and more positively skewed following downward-trending prices. As mentioned in the introduction, this prediction was subsequently tested by Chen, Hong, and Stein (1999), who have confirmed such an effect.^{27,28} Time-varying market participation in our model also implies a period of increased volatility following corrections. This implication is consistent with the findings of Schwert (1990). Surprisingly, even though fixed setup costs conceal the information of sidelined investors, by creating a filter on participation the overall effect can be to aggregate more information per trade.

²⁷ An ancillary implication of our approach is that such effects will be stronger in securities for which setup costs of trading are higher (since the effects are driven by the sidelining of investors). This suggests a stronger effect in small firms than in large firms. Chen, Hong, and Stein find no such effect of size. For reasons of tractability, the implications of our model are stylized. We do not attempt to calibrate based upon realistic parameter values to predict the magnitude of skewness effects at different time horizons. Thus the findings of Chen, Hong, and Stein should be viewed as confirmation of the model's key stylized implication, but there is no indication as to whether these results fit a realistically calibrated multiperiod version of either of our models. This remains a topic for further exploration.

²⁸ Campa, Chang, and Reider (1998) document a positive correlation between the skewness of the risk-neutral exchange rate distribution and the exchange rate level. To the extent that exchange rate levels are correlated with past changes in exchange rates, this finding is suggestive that exchange rate changes may be correlated with shifts in the skewness of actual future exchange rate changes. Alternatively, however, this finding may reflect shifts in state prices of consumption as a function of the exchange rate level.

The particular market imperfection we have focused on as a source of information blockage and conditional patterns of volatility and skewness is a fixed setup cost of transacting. An interesting further direction in which our approach can be taken is to explore how conditional moments of return distributions are affected by other market imperfections that can sideline informed investors. In a behavioral setting, Hong and Stein (2001) pursue such a route to explore the informational effects of short-sale constraints. More generally, it would be of interest to consider fixed setup costs that are finite for short sales but larger than for purchases. It is also possible that limits to participation differ across investor classes (e.g., individuals and different classes of institutions). Such an analysis may offer implications for trading strategies of different groups in relation to past returns and to shifts in variance and skewness.

Empirically it would be interesting to examine which classes of investors tend most often to carry bursts of sidelined information to the market. Such an investor class should display especially heavy trading at the time of a reversal of a price trend. We are not aware of any existing studies of this issue, but the data is available to examine it.

A market in which conditional volatility and skewness arise from blockage and release of the information pipeline has some other features which merit further research. For example, it would be interesting to examine how volume of trade is related to future volatility and skewness. Another interesting direction is to explore the effects of communication between individuals. In our framework, price changes are driven by market participants' gradual learning of each others' information rather than by the observation of new signals. In our model such a process makes price movements highly sensitive to communication among traders.²⁹ What if the "silent majority" of sidelined investors is not so silent? The introduction of conversation among a subset of market participants may have large effects on the equilibrium. A sidelined investor who learns that another sidelined investor shares a similar signal may decide to participate. Such participation may trigger rapid emptying of the sideline, and thereby large sudden price moves.

Similarly, if a market "guru" such as Abbie Cohen receives a private signal, it will not be her immediate trading on this signal that moves the price so much as the announcement of her information by means of a market forecast. Hence gurus or investment newsletters may have effects that seem disproportionate to the accuracy or timeliness of the information they reveal.

²⁹ Bloomfield and Libby (1996) found securities prices in an experimental market to be strongly influenced by the cross-sectional distribution of information, holding constant the aggregate information available to market participants. For an analysis of verbal communication and influence and its effects on financial markets when individuals are "near rational," see DeMarzo, Vayanos, and Zwiebel (1998).

Appendix A: Proof of Proposition 1

We show that informed investors' demand is positively related to the period 1 price. Let I_t denote informed investors' demand. We need to show that $\text{corr}(P_1, I_t) > 0$ when γ is close to zero. Let I_{st}^σ denote the total demand of informed investors who have received signal σ when the true state is s at time t . We first determine period 1 price and informed investors' demand and then the correlation of the two.

Since $\pi_{h0} = \pi_{l0} = \gamma/2$, the initial price will be 0.5. From Equation (2), after seeing the initial demand, D_0 , the market maker will announce a period 1 price of

$$P_1 = 0.5 + \frac{\gamma}{4} \left[\frac{f(D_0 - I_{h0}) - f(D_0 - I_{l0})}{\gamma f(D_0 - I_{h0})/2 + \gamma f(D_0 - I_{l0})/2 + (1 - \gamma)f(D_0)} \right]. \quad (12)$$

It is easy to show that price will be an increasing function of total demand, D_0 . When γ is close to zero, the price at time 1 can be simplified to

$$P_1 = \frac{1}{2} + \frac{\gamma}{4} \frac{f(D_0 - I_{h0}) - f(D_0 - I_{l0})}{f(D_0)} = \frac{1}{2} + \frac{\gamma}{4} \left[e^{D_0 I_{h0} - \frac{I_{h0}^2}{2}} - e^{D_0 I_{l0} - \frac{I_{l0}^2}{2}} \right], \quad (13)$$

where $f(\cdot)$ is the normal density function.

The informed investors' demand depends on their expectations of the value of the asset. Since informed traders know the true state is either h or l , and thus $\pi_{h0}^1 = \pi_{l0}^0 = p$, their perception of state probabilities in period 1 is

$$\pi_{h1}^1 = \frac{pf(D_0 - I_{h0})}{pf(D_0 - I_{h0}) + (1 - p)f(D_0 - I_{l0})}, \quad (14)$$

and

$$\pi_{h1}^0 = \frac{(1 - p)f(D_0 - I_{h0})}{(1 - p)f(D_0 - I_{h0}) + pf(D_0 - I_{l0})}. \quad (15)$$

Since I_{h0} is positive and I_{l0} is negative, it is clear that informed investors' expectation is an increasing function of the total demand in period 1. The total demand of informed investors with positive signals at state h is

$$I_{h1}^1 = p(\pi_{h1}^1 - P_1), \quad (16)$$

while the total demand of informed investors with negative signals at state h is

$$I_{h1}^0 = (1 - p)(\pi_{h1}^0 - P_1). \quad (17)$$

We first show that the conditional correlation between informed investors' demand and price is positive. Given state h , to prove that $\text{cov}(P_1, I_{h1}^1) > 0$, it suffices to show that $\text{cov}(P_1, \pi_{h1}^1) > 0$, because γ is close to zero, $\text{cov}(P_1, \pi_{h1}^1)$ is of the same order as γ , and $\text{var}(P_1)$ is of the order of γ^2 . Since both P_1 and π_{h1}^1 are increasing functions of D_0 , $\text{cov}(P_1, \pi_{h1}^1) > 0$. Consequently, $\text{cov}(P_1, I_{h1}^1) > 0$. Similarly, it can be shown that $\text{cov}(P_1, I_{h1}^0) > 0$. Using similar logic, there is a positive correlation given state l as well.

Now we are ready to show that the total demand and period 1 price are positively correlated. To show that trader demand is positively related to period 1 prices, we need to demonstrate that $\text{corr}(P_1, I_t) > 0$. Thus we need to show that

$$\begin{aligned} & \int_{D_0} (P_1 - 0.5) \{ [p(\pi_{h1}^1 - P_1) + (1 - p)(\pi_{h1}^0 - P_1)] 0.5f(D_0 - I_{h0}) \\ & + [(1 - p)(\pi_{h1}^1 - P_1) + p(\pi_{h1}^0 - P_1)] 0.5f(D_0 - I_{l0}) \} dD_0 \end{aligned}$$

$$= \int_{D_0} (P_1 - 0.5) \{ [p\pi_{h1}^1 + (1-p)\pi_{h1}^0] 0.5f(D_0 - I_{h0}) \\ + [(1-p)\pi_{h1}^1 + p\pi_{h1}^0] 0.5f(D_0 - I_{l0}) \} dD_0 + \text{var}[P_1].$$

Since the variance of P_1 is of the order of γ^2 , we can ignore the last term $\text{var}[P_1]$, so we obtain

$$\int_{D_0} \frac{\gamma}{4} \frac{[f(D_0 - I_{h0}) - f(D_0 - I_{l0})][pf(D_0 - I_{h0}) + (1-p)f(D_0 - I_{l0})]}{(\gamma/2)f(D_0 - I_{h0}) + (\gamma/2)f(D_0 - I_{l0}) + (1-\gamma)f(D_0)} dD_0 \\ = \int_{D_0} \frac{\gamma}{4} \frac{(p - \frac{1}{2})[f(D_0 - I_{h0}) - f(D_0 - I_{l0})]^2 + \frac{1}{2}[f(D_0 - I_{h0})^2 - f(D_0 - I_{l0})^2]}{(\gamma/2)f(D_0 - I_{h0}) + (\gamma/2)f(D_0 - I_{l0}) + (1-\gamma)f(D_0)} dD_0 \\ = \int_{D_0} \frac{\gamma}{4} \frac{(p - 0.5)[f(D_0 - I_{h0}) - f(D_0 - I_{l0})]^2}{(\gamma/2)f(D_0 - I_{h0}) + (\gamma/2)f(D_0 - I_{l0}) + (1-\gamma)f(D_0)} dD_0 > 0. \quad \blacksquare$$

Appendix B: Proof of Proposition 2

From the expression of price in Equation (2), we can determine the skewness using direct integration:

$$E[(P_2 - P_1)^3 | P_1] = \int_{D_1} (P_2 - P_1)^3 [\pi_{h1}f(D_1 - I_{h1}) + \pi_{l1}f(D_1 - I_{l1}) + (1 - \pi_{h1} - \pi_{l1})f(D_1)] dD_1 \\ = \int_{D_1} [(\pi_{h1})^3 e^{3D_1 I_{h1} - 3(I_{h1})^2/2} - (\pi_{l1})^3 e^{3D_1 I_{l1} - 3(I_{l1})^2/2} - (\pi_{h1} - \pi_{l1})^3 \\ - 3(\pi_{h1})^2 e^{2D_1 I_{h1} - (I_{h1})^2} \pi_{l1} e^{D_1 I_{l1} - (I_{l1})^2/2} - 3(\pi_{h1})^2 e^{2D_1 I_{h1} - (I_{h1})^2} (\pi_{h1} - \pi_{l1}) \\ + 3(\pi_{l1})^2 e^{2D_1 I_{l1} - (I_{l1})^2} \pi_{h1} e^{D_1 I_{h1} - (I_{h1})^2/2} - 3(\pi_{l1})^2 e^{2D_1 I_{l1} - (I_{l1})^2} (\pi_{h1} - \pi_{l1}) \\ + 3(\pi_{h1} - \pi_{l1})^2 \pi_{h1} e^{D_1 I_{h1} - (I_{h1})^2/2} - 3(\pi_{h1} - \pi_{l1})^2 \pi_{l1} e^{D_1 I_{l1} - (I_{l1})^2/2} \\ + 6\pi_{h1}\pi_{l1}(\pi_{h1} - \pi_{l1})e^{D_1(I_{h1} + I_{l1}) - [(I_{l1})^2 + (I_{h1})^2]/2}] f(D_1) dD_1 \\ = (\pi_{h1})^3 e^{3(I_{h1})^2} - (\pi_{l1})^3 e^{3(I_{l1})^2} + 2(\pi_{h1} - \pi_{l1})^3 \\ - 3(\pi_{h1})^2 \pi_{l1} e^{(2I_{h1} + I_{l1})^2 - 2(I_{h1})^2 - (I_{l1})^2/2} \\ - 3(\pi_{h1})^2 (\pi_{h1} - \pi_{l1}) e^{(I_{h1})^2} + 3(\pi_{l1})^2 \pi_{h1} e^{[(2I_{l1} + I_{h1})^2 - 2(I_{l1})^2 - (I_{h1})^2]/2} \\ - 3(\pi_{l1})^2 (\pi_{h1} - \pi_{l1}) e^{(I_{l1})^2} + 6\pi_{h1}\pi_{l1}(\pi_{h1} - \pi_{l1}) e^{I_{l1} I_{h1}}. \quad (18)$$

$$E[(P_2 - P_1)^2 | P_1] = \int_{D_1} (P_2 - P_1)^2 [\pi_{h1}f(D_1 - I_{h1}) + \pi_{l1}f(D_1 - I_{l1}) + (1 - \pi_{h1} - \pi_{l1})f(D_1)] dD_1 \\ = (\pi_{h1})^2 e^{(I_{h1})^2} + (\pi_{l1})^2 e^{(I_{l1})^2} - (\pi_{h1} - \pi_{l1})^2 - 2\pi_{h1}\pi_{l1} e^{I_{l1} I_{h1}}. \quad \blacksquare$$

References

- Avery, C., and P. Zemsky, 1998, "Multi-Dimensional Uncertainty and Herd Behavior in Financial Markets," *American Economic Review*, 88, 724–748.
- Black, F., 1972, "Capital Market Equilibrium with Restricted Borrowing," *Journal of Business*, 44, 444–455.
- Bloomfield, R., and R. Libby, 1996, "Market Reactions to Differentially Available Information," *Journal of Accounting Research*, 34, 183–207.

- Bollerslev, T., R. F. Engle, and J. M. Wooldridge, 1988, "A Capital Asset Pricing Model with Time Varying Covariances," *Journal of Political Economy*, 96, 116–131.
- Brown, D. P., and R. H. Jennings, 1989, "On Technical Analysis," *Review of Financial Studies*, 2, 527–551.
- Campa, J. M., P. H. K. Chang, and R. L. Reider, 1998, "Implied Exchange Rate Distributions: Evidence from the OTC Options Markets," *Journal of International Money and Finance*, 17, 117–160.
- Cao, H. H., 2001, "Endogenous Liquidity Trading: A Partially Revealing Rational Expectations Model for Multi-Asset Securities Markets," working paper, University of North Carolina.
- Carmel, J. M., 1995, "Information Aggregation without Exogenous Liquidity Trading," working paper, University of Michigan.
- Charoenrook, A., 2000, "The Role of Capital Structure in Tests of Asset Pricing Models: Theory and Empirical Evidence," working paper, University of Michigan.
- Chen, J., H. Hong, and J. C. Stein, 1999, "Forecasting Crashes: Trading Volume, Past Returns, and Conditional Skewness in Stock Prices," working paper, Stanford University.
- Christie, A. A., 1982, "The Stochastic Behavior of Common Stock Variances: Value, Leverage, and Interest Rate Effects," *Journal of Financial Economics*, 23, 407–432.
- Cutler, D. M., J. M. Poterba, and L. H. Summers, 1989, "What Moves Stock-Prices," *Journal of Portfolio Management*, 15, 4–12.
- DeBondt, W. F. M., and R. H. Thaler, 1985, "Does the Stock Market Overreact?," *Journal of Finance*, 40, 793–808.
- DeBondt, W. F. M., and R. H. Thaler, 1987, "Further Evidence on Investor Overreaction and Stock Market Seasonality," *Journal of Finance*, 42, 557–581.
- DeMarzo, P., D. Vayanos, and J. Zweibel, 1998, "A Near-Rational Model of Persuasion—With Implications for Financial Markets," working paper, Stanford University.
- Diamond, D. W., and R. E. Verrecchia, 1982, "Optimal Managerial Contracts and Equilibrium Security Prices," *Journal of Finance*, 37, 275–287.
- Easley, D., and M. O'Hara, 1987, "Price, Trade Size, and Information in Securities Markets," *Journal of Financial Economics*, 19, 69–90.
- Ederington, L., and J. H. Lee, 1996, "Persistent Futures Price Patterns: Market Making, Limit Orders, and Time," working paper, University of Oklahoma.
- Engle, R. F., and T. Bollerslev, 1986, "Modelling the Persistence of Conditional Variances," *Economic Reviews*, 5, 1–50.
- Fishman, M. J., and K. M. Hagerty, 1992, "Insider Trading and the Efficiency of Stock Prices," *Rand Journal of Economics*, 23, 106–122.
- French, K. R., and R. Roll, 1988, "Stock Return Variances: The Arrival of Information and the Reaction of Traders," *Journal of Financial Economics*, 17, 5–26.
- Glosten, L. R., and P. R. Milgrom, 1985, "Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71–100.
- Harvey, C. R., and A. Siddique, 1999, "Conditional Skewness in Asset Pricing Tests," forthcoming in *Journal of Finance*.
- Hong, H., and J. C. Stein, 2001, "Differences of Opinion, Rational Arbitrage, and Market Crashes," working paper, Stanford University.
- Jegadeesh, N., and S. Titman, 1993, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *Journal of Finance*, 48, 65–91.

- Kumar, P., and D. J. Seppi, 1992, "Futures Manipulation with Cash Settlement," *Journal of Finance*, 47, 1485–1502.
- Kyle, P., 1985, "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1335.
- Lee, I. H., 1998, "Market Crashes and Informational Avalanches," *Review of Economic Studies*, 65, 741–759.
- Llorente, G., R. Michaely, G. Saar, and J. Wang, 1999, "Dynamic Volume-Return Relation of Individual Stocks," working paper, Cornell University.
- Lo, A. W., and J. Wang, 2000, "Trading Volume: Definitions, Data Analysis, and Implications of Portfolio Theory," Working Paper no. 7625, NBER, Cambridge, MA.
- Nelson, D. B., 1991, "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 59, 347–370.
- Roll, R., 1988, " R^2 ," *Journal of Finance*, 43, 541–566.
- Romer, D., 1993, "Rational Asset-Price Movements Without News," *American Economic Review*, 83, 1112–1130.
- Schwert, G. W., 1990, "Stock Volatility and the Crash of '87," *Review of Financial Studies*, 3, 77–102.
- Shiller, R. J., 1995, "Conversation, Information, and Herd Behavior," *American Economic Review*, 85, 181–185.
- Veronesi, P., 1999, "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model," *Review of Financial Studies*, 12, 975–1007.
- Wang, J., 1993, "A Model of Intertemporal Asset Prices Under Asymmetric Information," *Review of Economic Studies*, 60, 249–282.

Copyright of Review of Financial Studies is the property of Society for Financial Studies and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.