Underestimation of Portfolio Insurance and the Crash of October 1987

Charles J. Jacklin
Allan W. Kleidon
Paul Pfleiderer
Stanford University

We examine market crashes in the multiperiod framework of Glosten and Milgrom (1985). Our analysis shows that if the market's prior beliefs underestimate the extent of dynamic hedging strategies such as portfolio insurance, then the price will be greater than that which would be implied by fundamentals if the extent of portfolio insurance were with certainty. Over time, the market learns of the amount of portfolio insurance, and consequently reevaluates the previous inferences drawn from purchases that were erroneously regarded as based on favorable information. The result is that the price falls when the amount of portfolio insurance is revealed.

One-time events allow for many possible explanations, and the stock market crash of October 1987 is no exception. Many attempts to explain the behavior of prices on October 19 and the surrounding days emphasize strains on liquidity in the major markets

Helpful comments have been received from Gene Fama, Gerard Gennotte, Bruce Grundy, Hayne Leland, Mark Rubinstein. Bill Sharpe, Chester Spatt, and an anonymous referee. We are also grateful for research assistance from Howard Corb, Stephen Gray, Rui Kan, and especially Robert Whitelaw. This research is supported by the Stanford Program in Finance. The first author acknowledges the support of the Fletcher Jones Faculty Fellowship, the second author acknowledges the support of Batterymarch Financial Management and the Business School Trust Faculty Fellowship, and the third author acknowledges the support of the Robert M. and Anne T. Bass Fellowship. Address correspondence and reprint requests to Charles J. Jacklin, Graduate School of Business, Stanford University, Stanford, CA 94305.

The Review of Financial Studies 1992 Volume 5, number 1, pp. 35-63

[see, e.g., the Presidential Task Force or Brady Report (1988)]. While it is quite likely that liquidity problems created by the extremely large volume of transactions played an important role, it seems doubtful that the crash can be understood solely in such terms. First, prior to the crash there was an extended rise in stock prices, and, second, if the precipitous price fall on October 19 were due solely to intense selling pressure and illiquidity, then it seems difficult to explain why prices did not return to their precrash levels until 1989.[1] In this article, we argue that both an extended price rise and a permanent price fall can be explained in terms of imperfect information about the extent of dynamic hedging strategies, which we collectively label as "portfolio insurance."

Prior to the crash, Grossman (1988) examined the liquidity implications of imperfect information about portfolio insurance. He argued that if individuals follow portfolio insurance programs that are implemented through dynamic trading strategies, no direct information about the amount of portfolio insurance is received by the market, and hence coordinated selling by a number of insured investors can take the market by surprise. Grossman argues that the use of dynamic strategies means that "the stocks' future price volatility can rise because of a current lack of information about the extent to which dynamic hedging strategies are in place" (p. 278). The key feature of Grossman's model is that the fraction of security holders who choose a dynamic hedging strategy is not public knowledge prior to their trading.

In this article, we also assume that the amount of dynamic hedging is not public knowledge. However, we focus not on the potential liquidity problems caused by coordinated selling-which we acknowledge as real——but rather on the inferential problem caused prior to such selling. We look at the period before October 19, and ask this: If the market was surprised at the amount of portfolio insurance selling on that day (or more correctly on the previous Friday, October 16), then what are the implications for the price level *before* the amount of portfolio insurance was revealed? We are particularly interested in the problem caused for those investors who rationally condition their demands on the demands of others because they are imperfectly informed.

We assume that some uninformed investors follow dynamic strategies, ranging from informal stop-loss strategies (which have long existed) to the recently introduced formal portfolio insurance con-

[1] The (monthly) Dow Jones Industrial Average was 2224.59 at the beginning of October 1987, and remained below 2200 until January 1989. However. see Amihud, Mendelson, and Wood (1990). discussed below, for one explanation of a permanent price decline linked to liquidity.

tracts issued by portfolio insurers who hedge their commitments by following a dynamic trading strategy that synthesizes a put on the portfolio. It should be stressed that innovative formal portfolio insurance is only one (small) part of the strategic behavior we characterize here as "portfolio insurance." Given the development of stock index futures and their use by formal portfolio insurers, other investors are likely to have been informally relying on the use of dynamic trading strategies similar to the strategies being employed by formal portfolio insurers. Whether formal or informal, the dynamic hedging strategy entails increased buying of stocks (or futures on stocks) as prices rise and increased selling as prices fall. In addition, having obtained insurance, investors may then buy more aggressively (or sell less) than they would have in the absence of insurance.[2]

If there is imperfect information about the fraction of traders undertaking formal or informal insurance, then informationless trades can be misinterpreted as informed trading. We find that when the realized proportion of portfolio-insured investors is greater than the ex ante expectation, this can lead to the incorrect inference that purchases by such investors are based on favorable information. As the market learns over time about the proportion of insured investors, a reassessment is made of the apparent information contained in the previous trades.

We take the sales by insured investors, particularly on October 16, as revealing an important piece of information to the market, namely, the extent of dynamic strategies *and consequently the extent of prior uninformed purchases.*[3] Abstracting from the liquidity issues raised by Grossman, our model implies that the price prior to the crash was too high relative to the price that would have prevailed if the extent of portfolio insurance were common knowledge, and that it fell when information about the extent of portfolio insurance was revealed. There is nothing irrational about this; it is simply a consequence of imperfect information about the trading strategies of some uninformed traders.

This argument resolves one of the puzzles of the crash, namely, the source of the information if the crash was induced by a fundamental reevaluation of future prospects. In our model, there need be no change in market fundamentals per se, which is consistent with

[2] Federal Reserve Chairman Alan Greenspan testified before a Senate committee that portfolio insurance can encourage investors "to hold larger permanent equity positions"; see Miller et al. (1988, note 4). See also the Presidential Task Force on Market Mechanisms Report (1988, p. 9).

[3] The SEC (1988. p. 2-10) reports that on October 16, "portfolio insurers ended the day with a significant 'overhang' of sales of futures or stocks that were called for by the insurance parameters but could not be executed on the 16th." and that "this overhang, and investor knowledge of it, may have contributed to the rapid decline on the afternoon of the 16th and undoubtedly added to the selling pressure on the 19th."

the often-noted absence of identifiable news items of sufficient magnitude to explain the crash. Rather, what we need is something to "smoke out" the amount of portfolio insurance and show that it is higher than expected. There need be no specific information event that causes the original price decline that in turn unleashes the coordinated selling of dynamic strategies and hence reveals their quantity; even an unusual run of sales by uninformed liquidity traders that gets translated into a price decline will suffice. The price decline causes some selling by insured investors, which reveals to the market some information about the amount of portfolio insurance; the market then revises downward its valuation of the risky asset, which causes further selling by insured investors. On the other hand, some bad news that correctly causes a price decline can perfectly well serve as the catalyst that starts the downward plunge.[4]

Several other articles are closely related to ours. Amihud, Mendelson, and Wood (1990) argue that the liquidity problems of October 1987 caused by portfolio insurance sales showed that the market was not as liquid as previously believed, and that consequently the market value fell permanently. Black (1988) argues that expected returns changed on October 19, as investors "became aware that the typical investor was more dynamic than they thought" (p. 5). Brennan and Schwartz (1989) show that portfolio insurance is unlikely to create the requisite degree of illiquidity to explain October 1987 if the amount of such insurance activities is known. They add, however, that if the amount is not known, their results may be misleading either because of the liquidity arguments of Grossman (1988) or because of "possible misinterpretation of the information content of portfolio-insurance-induced transactions" (p. 471). Miller (1990) also regards portfolio insurance as potentially destabilizing if the absence of information in such trades is not understood.

Gennotte and Leland (1990) examine the kind of incorrect inferences alluded to by Brennan and Schwartz (1989) and Miller (1990) and exploited by our model, and make similar assumptions to ours with respect to ex post underestimation of the extent of portfolio insurance. However, our analyses differ in several important respects. They use a one-period model and assume that the resulting incorrect inferences are made on October 19 when large amounts of selling are carried out. Their argument is that on, say, October 16, the market was correctly priced relative to market fundamentals; that portfolio insurance selling on October 19 was widely and incorrectly interpreted as selling by informed traders, rather than by dynamic hedgers;

[4] The SEC Report (1988. p. 2-7) notes that some negative news items occurred in the week prior to the crash. on October 14-16. Mitchell and Netter (1989) point to antitakeover legislation proposed on October 13 as the fundamental cause of the 10 percent price decline by October 16.

and that the price consequently crashed discontinuously to a new low level. Our model is explicitly multiperiod in nature, and it is the rational learning over time about the extent of dynamic strategies that produces crashes.[5] Gennotte and Leland emphasize that a crash in their model is *not* due to a discontinuous change in the underlying information; in our model, the price level immediately prior to a crash is too high when judged by the admittedly unrealistic benchmark of what would have prevailed had the extent of dynamic hedging been common knowledge, and it falls when the extent is revealed. By contrast, no crash would have occurred in the model of Gennotte and Leland had the amount of portfolio insurance been revealed just prior to October 19.

We proceed as follows. In Section 1, we set out the model, which modifies the Glosten and Milgrom (1985) trading model to allow uncertainty about the extent of portfolio insurance. We use the Glosten and Milgrom model to focus on the difficulties of inference about asset value in the presence of unknown amounts of portfolio insurance, since the risk neutrality of the market maker in this model abstracts from liquidity effects. In Section 2, we present the results of the model, and demonstrate the behavior we link to the 1987 crash. We examine several variants to isolate the effects of underestimation of the amount of portfolio insurance. In particular, we contrast the properties of models in which there is a given amount of portfolio insurance, but where the market does or does not know the amount. In making these comparisons, we demonstrate that portfolio insurance per se does not lead to crashes in our model; the price behavior in terms of both means and variances is similar whether there is no portfolio insurance or whether there is a large amount of it *but the market knows this*. Finally, although we concentrate on mean effects in the price level, our model may also help to explain the dramatic rise in volatility experienced during and shortly after the crash. Caveats and a discussion of our model's empirical implications and its applicability to other market crashes are contained in Section 3. Concluding remarks are given in Section 4.

1. The Model

Our model of price information and the trading process is based on Glosten and Milgrom (1985). In each period a trader chosen at ran-

[5] The work by Brown and Jennings (1989) and Grundy and McNichols (1989) is also relevant here, since multiple rounds of trade reveal information about the previous amounts of informed versus uninformed trading. Our model similarly hinges on the revelation of the extent to which previous trading is split between informed and a particular class of uninformed investors. See also Easley and O'Hara (1991). who examine the effects of uninformed stop orders on securities prices.

dom is given the opportunity to trade with a risk-neutral market maker. At the beginning of period t the market maker posts a bid price B_t at which he is willing to buy one unit of the risky asset and an ask price A_t at which he is willing to sell one unit. The trader who arrives in that period observes these prices and then decides whether or not to trade. The market maker then revises his bid and ask prices for the next period (B_{t+1} and A_{t+1}) and bases this revision on whether he made a trade in period t and whether the trade was a sale or a purchase. In period T the risky asset has an exogenously determined payoff $\tilde{V}$. To keep the analysis as simple as possible we assume that $\tilde{V}$ can take one of three values, V_H, V_M, and V_L, with probabilities y_H, y_M, and y_L, respectively. As in Glosten and Milgrom we assume that the market maker is driven by competitive forces to post bid and ask prices that yield zero profit in expectation. This means that B_t is equal to the expectation of $\tilde{V}$ given the trading history (the sequence of trades observed by the market maker) up to period t and given the arrival of a sell order in period t. Similarly A_t is the expectation of $\tilde{V}$ given the trading history and the arrival of a buy order in period t.

The stochastic process governing the arrival of orders depends on the amount of portfolio insurance. Figure 1 shows how the order arrival process is determined each period. We assume that in each period an informed trader arrives with probability f and that this probability is independent of the amount of portfolio insurance. An informed trader knows the realization of $\tilde{V}$, that is, he knows whether $\tilde{V}$ is equal to V_H, V_M, or V_L. If an informed trader arrives at the market maker's desk in period t, he submits a buy order if $\tilde{V} > A_t$, submits a sell order if $\tilde{V} < B_t$, and does not trade if $B_t \leq \tilde{V} \leq A_t$.[6]

With probability $1 - f$ an uninformed trader arrives at the market maker's trading desk and is given the opportunity to trade. That uninformed trader can be either an insured investor (an investor who is either initiating or following a portfolio insurance strategy) or a liquidity trader. We characterize the amount of portfolio insurance by q, the probability that an uninformed trader is an insured investor. We assume (again to keep the analysis as simple as possible) that q can take on two possible values, q_H and q_L where $q_H > q_L$. Initially (in period 1) the market maker's prior probability that the level of portfolio insurance is high $(q = q_H)$ is q. Over time the orders that arrive at the market maker's desk will reveal some information about the amount of portfolio insurance. This means that as the market maker observes the sequence of orders, he will revise the probability that there is a high level of portfolio insurance from his initial prior of q.

[6] The informed trader is indifferent to trading or not trading if either $\tilde{V} = B_t$ or $\tilde{V} = A_t$. We assume that in this case of indifference the trader does not submit an order.

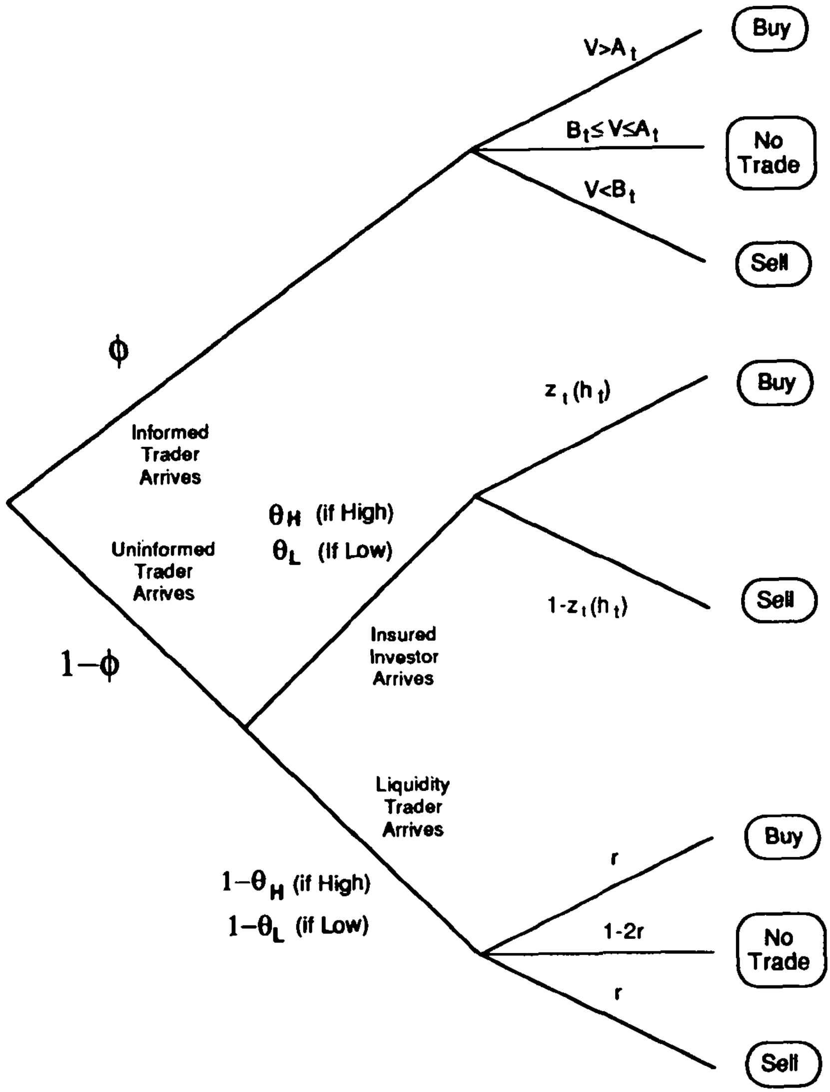

Figure 1
Determination of order arrival process each period

If an uninformed liquidity trader arrives, then that trader submits a sell order with probability *r*, a buy order with probability *r*, and does not trade with probability 1 - *2r*. We assume that liquidity traders trade for purely exogenous reasons. For example, a liquidity trader may be motivated to trade by an exogenous shock to his holdings of a nontraded asset that changes his optimal position in the risky asset. Given this view of liquidity trading, the assumption that liquidity traders are equally likely to submit buy or sell orders is a natural one to make.

Our specification of the collective behavior of uninformed investors who are following insurance strategies (insured investors) is more complex than the specification for liquidity traders. This is because in any period the order flow from insured investors depends on the path of prices before that period. Since portfolio insurance can be thought of as a long position in a bond and a call, an insured investor following a dynamic portfolio insurance strategy is essentially following the same strategy as used to replicate a call option. As the price of the risky asset rises, more of the risky asset is purchased. As the value falls, shares in the risky asset are sold. Within the context of the simple trading model we are analyzing it is clearly not possible to specify precisely the order flow created by insured investors following dynamic strategies. Even if it were possible to model precisely the order flow coming from one insured trader, we would still have difficulties in specifying the aggregate trading of all insured investors, because the actual execution of the dynamic strategy depends on the level of the insurance policy's "floor" or, equivalently, on the exercise price of the call. The aggregate trading therefore depends on the distribution of floors among the insured traders and this distribution is not easily observed, especially since there was a tendency to readjust the level of floors as portfolio values changed. In modeling portfolio insurance in the Glosten-Milgrom trading environment, we can therefore only hope to capture qualitatively the effects of these dynamic trading policies on the order flow. What we intend to develop here is a probabilistic representation that accounts for the most important consequences of portfolio insurance.

As mentioned in the introduction, the availability of portfolio insurance seems to have caused some investors to demand more of the risky asset than they otherwise would have demanded since insurance reduced the perceived risk of holding risky assets. This suggests that as portfolio insurance strategies are initiated, the insured investors are likely to be submitting buy orders. However, once a policy is in place, an insured investor submits a buy order only if the value of the risky asset has risen and a sell order only if it has fallen.

To represent these effects we let h_t be the history of trading up to, but not including, period t and let $\hat{V}_t = E(\tilde{V} | h_t,$ order flow at time $t)$. Note that if there is a trade in period t, the transaction price will equal $\hat{V}_t$. However, in those cases where there is no trade in period t, the conditional expectation $\hat{V}_t$ does not correspond to an observable price and it generally differs from previous transaction prices. When an insured trader arrives in the market, that trader can either be initiating a portfolio insurance policy or carrying out the dynamic trading strategy required by a policy that is already in effect. We let λ_t be the probability that the insured investor who is arriving in the market has already initiated a portfolio insurance program. It follows that $1 - \lambda_t$ is the probability that the insured investor is initiating a policy. We assume that λ_t is an increasing function of t since as time goes by more investors from the total pool of potential insured investors will have had the opportunity to establish insurance programs.

Given these definitions for $\hat{V}_t$ and λ_t, we can define $z_t(h_t)$, the probability that an insured investor submits a buy order if given the opportunity to trade in period t, to be

$$z_t(h_t) = (1 - \lambda_t) + \lambda_t \, \delta_{\{\hat{V}_{t-1} > \hat{V}_{t-2}\}}, \tag{1}$$

where $\delta_{\{\hat{V}_{t-1} > \hat{V}_{t-2}\}}$ is equal to 1 if $\hat{V}_{t-1} > \hat{V}_{t-2}$ is equal to 0 otherwise. In our specification of $z_t(h_t)$ we are implicitly assuming that if the trader arriving in the market is initiating a portfolio insurance policy, a buy order is submitted with probability equal to 1. This captures the notion that when a portfolio insurance program is put on, the trader tends to increase the position in risky assets. If the trader who arrives has already initiated portfolio insurance, then we assume a buy order is submitted only if the value of the risky asset has increased from period $t - 2$ to $t - 1$. This, of course, accounts in a rough way for the fact that the hedging strategy requires that a bigger position be taken in the risky asset when the value has risen and a smaller position when it has fallen. The function $z_t(h_t)$ therefore represents the stylized effects of portfolio insurance on the order flow.[7]

We now consider the market maker's determination of bid and ask prices in each period. Since there are two possible levels of portfolio insurance and three possible realizations of $\tilde{V}$, there are six possible states of the world. Let states 1, 2, and 3 be those in which the amount of portfolio insurance is high (the probability that an uninformed trader is an insured investor is q_H). In states 1, 2, and 3, the terminal value $\tilde{V} = V_H$, V_M, and V_L, respectively. Define states 4, 5, and 6 similarly for the case where the amount of portfolio insurance is low. Let $\pi_t^i(h_t)$ be the probability that the state is i given the trading history up to period t. Initially $\pi_1^1(h_1) = q\psi_H$, $\pi_1^2(h_1) = q\psi_M$, and $\pi_1^3(h_1) = q\psi_L$. The initial probabilities for states 4 through 6 are the same except that q is replaced by $(1 - q)$.

The bid and ask prices posted in period t are determined by the zero-profit conditions that $B_t = E(\tilde{V} \mid h_t$, a sell order is submitted in period $t)$ and $A_t = E(\tilde{V} \mid h_t$, a buy order is submitted in period $t)$. We

[7] We have considered several other possible specifications for $z_t(h_t)$, with virtually no effect on the qualitative nature of the simulation results reported below.

consider here the determination of B_t; the determination of A_t is very similar and is not discussed in detail. Define $S_t^i(h_t, B_t)$ to be the probability that a sell order is placed in period t given that (a) the state of the world is i, (b) the trading history is h_t, and (c) the posted bid price in period t is B_t.[8] Assume first that $B_t \geq V_M$. It can be easily verified that

$$S_t^1(h_t, B_t) = (1 - \phi)(\theta_H(1 - z_t(h_t)) + (1 - \theta_H)r), \tag{2a}$$

$$S_t^2(h_t, B_t) = \phi + (1 - \phi)(\theta_H(1 - z_t(h_t)) + (1 - \theta_H)r), \tag{2b}$$

$$S_t^3(h_t, B_t) = \phi + (1 - \phi)(\theta_H(1 - z_t(h_t)) + (1 - \theta_H)r), \tag{2c}$$

$$S_t^4(h_t, B_t) = (1 - \phi)(\theta_L(1 - z_t(h_t)) + (1 - \theta_L)r), \tag{2d}$$

$$S_t^5(h_t, B_t) = \phi + (1 - \phi)(\theta_L(1 - z_t(h_t)) + (1 - \theta_L)r), \tag{2e}$$

$$S_t^6(h_t, B_t) = \phi + (1 - \phi)(\theta_L(1 - z_t(h_t)) + (1 - \theta_L)r). \tag{2f}$$

Given the market maker's priors based on the trading history observed up to period t, the probability of a sell order being placed in period t is $\Sigma_{i=1}^6 \pi_t^i(h_t)S_t^i(h_t, B_t)$. If a sell order is received in period t, then the market maker's posterior assessments of the probabilities of the six states become (using Bayes' rule)

$$(\pi_{t+1}^i \mid h_t, \text{sell order in period } t) = \frac{\pi_t^i(h_t)S_t^i(h_t, B_t)}{\Sigma_{j=1}^6 \pi_t^j(h_t)S_t^j(h_t, B_t)}. \tag{3}$$

Then

$$\mathrm{E}(\tilde{V} \mid h_t, \text{a sell order is submitted in period } t)$$

$$= \sum_{i=1}^{6} \frac{\pi_t^i(h_t)S_t^i(h_t, B_t)V(i)}{\Sigma_{j=1}^6 \pi_t^j(h_t)S_t^j(h_t, B_t)}, \tag{4}$$

where we define $V(i)$ to be the realization of $\tilde{V}$ in state i. The zero-profit bid price is equal to the right-hand side of (4) as long as this is greater than or equal to V_M, since the values of $S^i(h_t, B_t)$ were predicated on the assumption that $B_t \geq V_M$. If the right-hand side of (4) is less than V_M, then $\mathrm{E}(\tilde{V} \mid h_t,$ a sell order is submitted in period $t)$ must be calculated as above except with

$$S_t^2(h_t, B_t) = (1 - \phi)(\theta_H(1 - z_t(h_t)) + (1 - \theta_H)r)$$

and

$$S_t^5(h_t, B_t) = (1 - \phi)(\theta_L(1 - z_t(h_t)) + (1 - \theta_L)r).$$

The value obtained then is the zero-profit bid price.[9]

[8] Given our description of the determinants of the behavior of the traders, the probability of a sell order is clearly not affected by the ask price the market maker posts, nor is the probability of a buy order affected by the posted bid price.

[9] In the general formulation of the Glosten and Milgrom model, there may be multiple ask and multiple bid prices that satisfy the zero-profit conditions. However, under our assumptions, the

2. Implementation and Results

In Section 2.1, we describe our implementation of the basic model in Section 1, including a discussion of parameter values and four different scenarios concerning the amount of portfolio insurance and the market maker's prior beliefs about portfolio insurance. We are primarily interested in the average behavior of the markets under the four scenarios we discuss. In Sections 2.2 and 2.3, we present results based on cross-sectional distributions over 10,000 replications for each of the four models. Each replication represents 100 periods, with the market maker optimally updating beliefs on all relevant parameters as information is revealed through the order sequence. In Section 2.2, we give the average price paths for each model, without conditioning on the terminal value. However, it is convenient to contrast results across the four models for the case when the realization of the terminal value is the mean of the terminal distribution, and these results, which are the major focus of our discussion, are presented in Section 2.3. We conclude, in Section 2.4, with examples of individual sample price paths for the models we discuss.

2.1 Implementation of the basic model

We compare four scenarios.

Model 1. The market maker is uncertain about the amount of portfolio insurance ($0 < q < 1$) and the actual amount is higher than the initial expectation ($q = q_H$).

Model 2. The market maker is uncertain about the amount of portfolio insurance and the actual amount is lower than the initial expectation ($0 < q < 1$, $q = q_L$).

Model 3. The market maker knows the amount of insurance with certainty and there is high portfolio insurance ($q = 1$, $q = q_H$).

Model 4. There is no portfolio insurance, and the market maker knows this with certainty.

In each case, we set $V_H = 100$, $V_M = 50$, and $V_L = 0$, and the market maker's prior distribution is $y_H = y_M = y_L = \frac{1}{3}$. For each case we take the probability of arrival of an informed investor to be $f = .20$, and we assume equal probability of a buy, sell, or no trade by an uninformed trader without portfolio insurance. The probability of a buy order by an insured investor is given by $z_t(h_t)$ in Equation (1) above,

pair of ask and bid prices that satisfies these conditions is unique. In particular, since the demand of the liquidity traders is inelastic, widening spreads would increase the market maker's expected profits from trade with liquidity traders and decrease his expected losses from trade with informed traders resulting in positive expected profit overall.

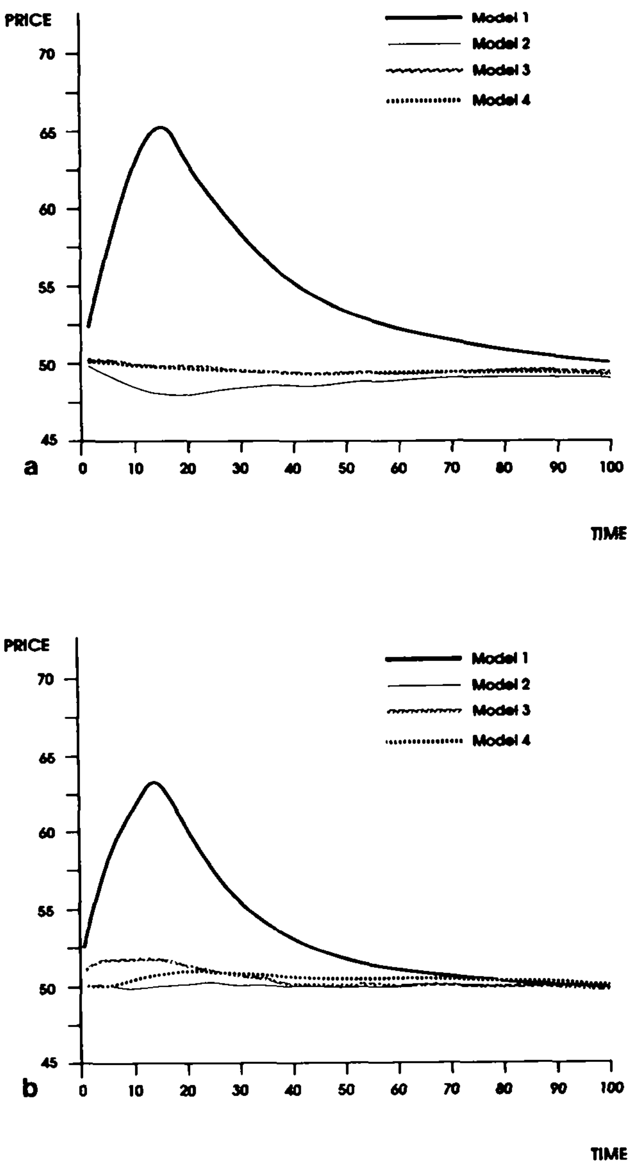

Figure 2
Plots of means of the distributions
Taken across 10,000 replications of the mean of the market maker's posterior distribution of terminal value at each date $t = 1,\ldots,100$ for Model 1 ($\theta = \theta_H = .4$, $q = .1$), Model 2 ($\theta = \theta_L = .1$, $q = .1$), Model 3 ($\theta = \theta_H = .4$, $q = 1.0$), and Model 4 (no insurance), where q is the market maker's prior assessment that the amount of portfolio insurance, θ, equals θ_H, with (a) no conditioning on terminal value or (b) conditioning on mean terminal value.

with λ, (the probability that an insured trader has already initiated a portfolio insurance program) given by $F(t - 15)/2)$, where $F(\cdot)$ is the standard normal distribution function, One can interpret this parameterization to mean that by period 15 half of the total pool of insured investors has initiated a portfolio insurance policy.

The models differ in the fraction of the uninformed investors that follow portfolio insurance strategies, and in whether the market maker knows the extent of such insurance. In Models 1 and 2, there is portfolio insurance and the market maker is uncertain about its extent. In Model 1, which is the major focus of this article, there is a large amount of portfolio insurance: the probability that the uninformed investor has portfolio insurance is 40 percent ($q = q_H = .4$). In Model 2, the corresponding probability is 10 percent ($q = q_L = .1$). The market maker's prior belief in both cases is 90 percent probability of low insurance and 10 percent probability of high insurance, that is, $q = .1$.

There is high portfolio insurance in Model 3, which the market maker knows with certainty. Hence, the only difference in parameters between Models 1 and 3 is that for Model 3 the prior probability of high insurance is given by $q = 1$. In Model 4, no uninformed investors have portfolio insurance, which the market maker knows with certainty.

2.2 Average price effects: No conditioning on terminal value

In Figure 2, we give the average across 10,000 replications of the mean of the market maker's posterior distribution of terminal value at each date $t = 1,...,100$, for each of the models we examine. Since the market maker sets bids and asks as the means of the posterior distribution of terminal value conditional on a trade at that bid or ask, we refer to these means as transactions prices.[10] For each of the replications underlying Figure 2a, the terminal value was chosen randomly from the distribution of $\tilde{V}$. Hence Figure 2a shows the average price paths for each of the four models, with no conditioning on the terminal value.[11] Figure 2b shows comparable results conditional on the terminal value equaling its expectation.

The most important result from Figure 2a is that uncertainty on the part of the market maker about the amount of portfolio insurance causes deviations of prices from the mean of the terminal distribution. This is seen in the average price paths for Models 1 and 2, which

[10] This involves a slight shorthand, since in some periods there will be no trade and hence no market transaction price; the mean is then conditional on no trade.

[11] To maintain comparability, for each replication the random draws that determine whether an informed or uninformed trader has the opportunity to trade at date t are identical across the four models; similarly, whether the uninformed trader is using portfolio insurance is identical for Models 1 and 3.

have high and low levels of portfolio insurance, respectively (although the market maker does not know the amount). In Model 1, the market maker confuses initial purchases by insured investors with informed purchases, so that the average price rises to a high of 65.22 in period 15. Over time, however, price declines (due to informed or liquidity sales) trigger sales from insured investors. These sales provide information about the amount of portfolio insurance and cause the market maker to revise beliefs about the terminal value downward. This is seen in Figure 2a by the subsequent decline in the average price to 50.00 by period 100.

The opposite behavior is seen in Model 2, in which there is a low amount of portfolio insurance. Since the market maker initially places some positive probability (here, .1) on high insurance, there will initially be fewer purchases from insured investors than the market maker expects, which will be misinterpreted as bad news. This is seen in Figure 2a as a decline in average price for Model 2 to 47.80 in period 18. Over time, the trading sequence reveals information about the true amount of portfolio insurance, and the average price rises to 49.15 at $t = 100$.

It is clear from Figure 2a that the issue is not portfolio insurance per se, but rather uncertainty about the quantity of insurance, since the average price paths are virtually identical for Model 3, in which there is high portfolio insurance and this is known with certainty, and Model 4, in which there is no portfolio insurance. In both cases, the average price is the mean of the terminal distribution for each date $t = 1,...,100$, with none of the average deviation from this value caused by the uncertainty in Models 1 and 2.[12]

Finally, note that since the market maker's prior belief places a 10 percent chance on high insurance, and since the price each period is set as the market maker's conditional expectation (given the sequence of trades) of the terminal value, we should expect to see a greater average price rise if there is high insurance than the average price fall if there is low insurance. If we were to average across realizations of the amount of portfolio insurance, and if the low insurance realization occurred more frequently than the high realization (commensurate with the market maker's prior beliefs), then average prices would be set appropriately even if there were portfolio insurance. However, conditioning on realizations of high and low levels of insurance reveals the greater variance caused by uncertainty about the extent of insurance.

For those replications underlying Figure 2a in which the terminal value was $V_H = 100$, the effect of portfolio insurance will be to take

[12]Since there were slightly more draws of V_t than of V_H in our 10,000 replications (3407 versus 3270), the (empirical) mean of the terminal distribution in Figure 2a is slightly below 50.

the price to its ex post value faster than in its absence—there will be no average "crash," since the terminal value is the maximum possible. We are interested in the possibility of the price rising above the ex post value, and then falling to that value as information about the extent of portfolio insurance is revealed. This is represented in our framework by Model 1, in which there is more insurance than initially expected, and by the case in which the terminal value is equal to the initial unconditional expectation. This case is given in Figure 2b, to which we now turn.

2.3 Cross-sectional results: Conditioned on mean terminal value

All results in this section condition on $\tilde{V} = V_M = 50$. We present results on the distributions across the 10,000 replications of following variables: the mean of the market maker's posterior distribution of terminal value at each date $t = 1,...,100$; the price maximum and minimum for each replication; the variance of the market maker's posterior distribution of terminal value at each date $t = 1,...,100$; the posterior probability of high portfolio insurance at each date $t = 1,...,100$; and whether the trade at each date was a purchase or sale. These are discussed in turn.

Mean of posterior distribution. Figure 2b shows results corresponding to those in 2a, namely, the cross-sectional average of the mean of the market maker's posterior distribution of terminal value at each date $t = 1,...,100$ (for which we again use the shorthand "price"), for each of the four models. However, we now condition on the mean terminal value. There is very little to distinguish between prices for the very different behavior represented by Models 2, 3, and 4. For example, Model 3 has high portfolio insurance while Model 4 has none, yet there is little difference between the prices shown in Figure 2b.[13] The market maker is able to interpret portfolio-insurance-based trades very accurately if the amount of insurance is known with certainty.

By contrast, the average price path of Model 1 is very different from that of the other models, indicating the effects of underestimation of the quantity of insurance based trades. While Models 2, 3, and 4 have an average price of approximately 50 at each date $t = 1,...,100$, the average price for Model 1 deviates from this "fundamental value." The average price rises steadily from the unconditional prior value

[15]We have seen from Figure 2a that when there is no conditioning on terminal value, the mean price path for Model 3 is virtually identical to that of Model 4. We see in Figure 2b that there are some slight effects of conditioning on $\tilde{V} = V_M$, since there is a very slight average price rise for Model 3 in Figure 2b, which is ignored in subsequent analyses. Percentiles of the cross-sectional distributions are also very similar for these three models.

of 50 to a maximum of 62.85 at t = 14, then declines to about 50 by period 70.[14]

Glosten and Milgrom's model implies that the market maker will eventually learn all parameters for each replication. However, while on average the mispricing caused by lack of knowledge of the amount of portfolio insurance dies out by about period 70 as shown in Figure 2b, the effects are still evident in the 1st and 99th price percentiles, which have a larger range for Model 1, even by t = 100, than for Models 2, 3, and 4. This increased uncertainty is discussed more fully below with respect to the variance of the posterior distributions of terminal value.

Price maximum and minimum. Figure 3 shows percentiles of the distributions of the price maximum and minimum for each replication across the 10,000 replications. Figure 3a shows a greater maximum induced by underestimation of the amount of portfolio insurance. For example, the median price maximum for Model 1 is 81.04, while that for each of the other models is about 65. Since the price will ultimately go to the terminal value of 50, the maximum for Model 1 represents a 62 percent premium over the full information value, while the maximum for the other models is only 30 percent above 50. The comparable values for the distributions of the minimum are given in Figure 3b. There is little difference across each of the four models between the medians, which are about 35 (or 30 percent below 50) .[15]

Variance of posterior distribution. A well documented feature of the 1987 crash was that uncertainty in the market, as reflected for example in the implied volatilities of options, increased significantly around the crash [see, for example, Schwert (1990)]. We now show that, for our model, the variance of the posterior distribution of the terminal value is greater when inference is complicated by uncertainty about portfolio insurance.

The variances for Models 1, 2, and 3 are scaled by the variance of Model 4 (no portfolio insurance) at each date $t = 1,...,100$, since in the Glosten and Milgrom model the variance of the posterior distribution decreases as time passes; in effect, there is no renewal of the

[14]Although smooth mispricing relative to fundamentals appears in the averages across the 10,000 replications, such behavior is not shown in any individual replication (see Section 2.4). Since in our models there is no information release at a particular date that acts as a catalyst for insurance selling (such as the bad news on October 14-16, 1987), price falls occur at different dates across the replications, resulting in the smooth average behavior of Figure 2.

[15]The average price behavior in Figure 2 suggests that for Model 1 the maximum should occur before the minimum, and this is indeed the case; the average difference between the time to the maximum and the time to the minimum is negative (-15) for Model 1, and close to 0 for the other models.

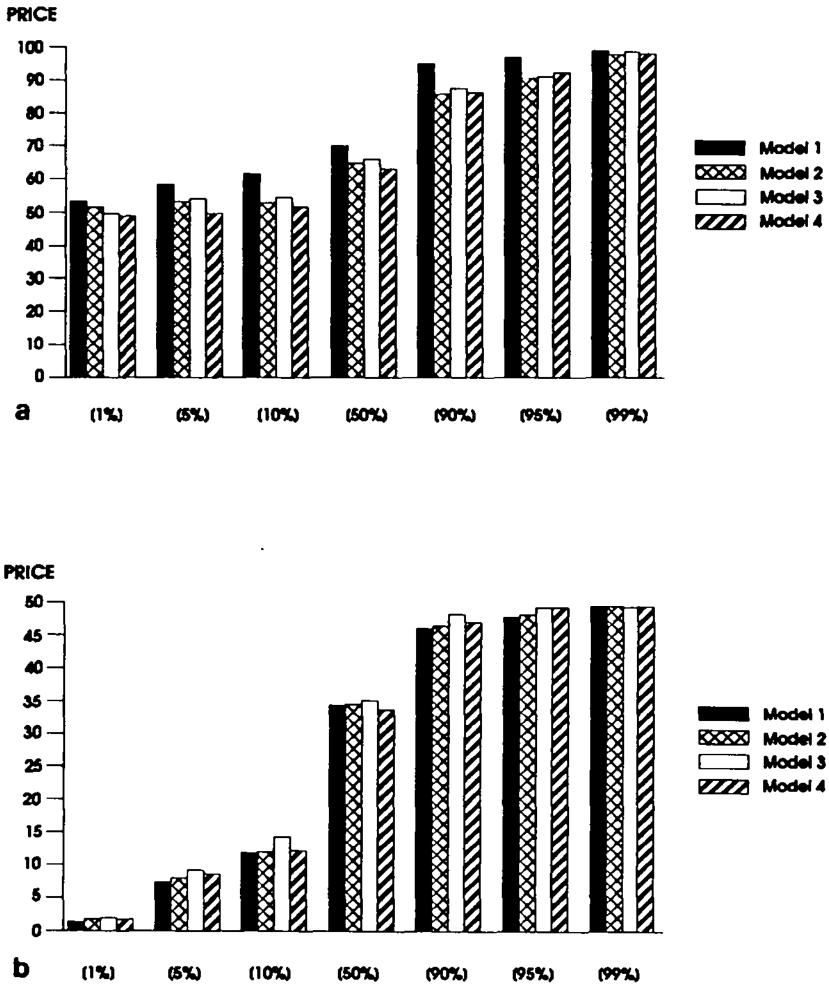

Figure 3
Percentiles of the distributions of the prices
The (a) maximum and (b) minimum of the prices for each replication across 10,000 replications, for Model 1 ($\theta = \theta_H = .4$, $q = .1$), Model 2 ($\theta = \theta_L = .1$, $q = .1$), Model 3 ($\theta = \theta_H = .4$, $q = 1.0$), and Model 4 (no insurance), where q is the market maker's prior assessment that the amount of portfolio insurance, θ, equals θ_H.

fundamental uncertainty. If there is no difference between the variance of the posterior distributions of terminal value across models, this variance ratio will equal 1.0 for each of Models 1, 2, and 3.

The (median) variance results are given in Figure 4. For each date from 1 to 100, the median ratios for Models 2 and 3 are close to 1, which means that there is little difference in the variance of the posterior distributions of terminal value for Models 2, 3, and 4. However, the variance ratio of Model 1 grows over time. By t = 100, the

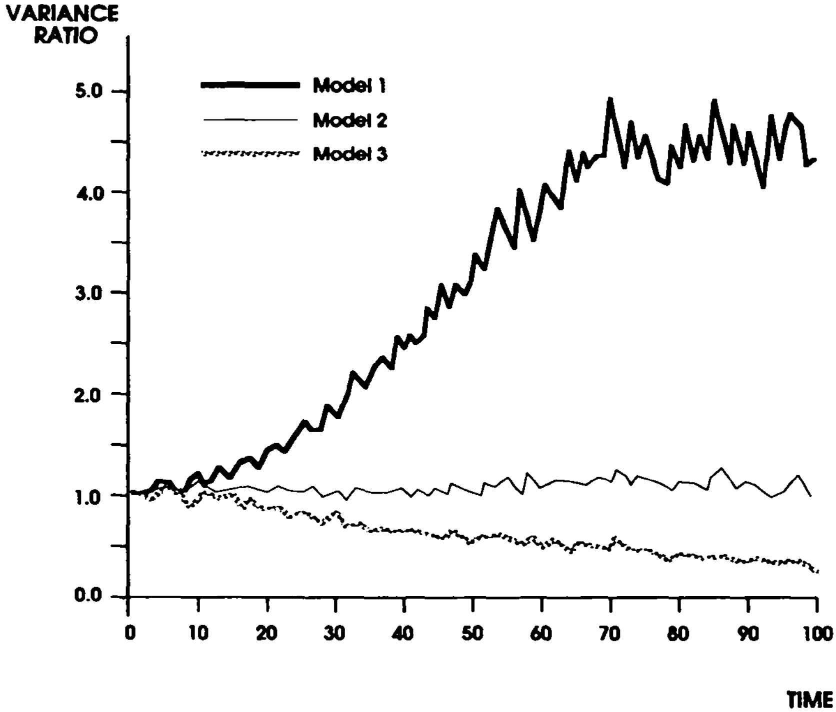

Figure 4
Plots at each date t = 1,...,100 of the 50th percentiles of the distributions
Taken across 10,000 replications of the ratio of the variance of the market maker's posterior distribution of terminal value for Model 1 ($\theta = \theta_H = .4$, $q = .1$), Model 2 ($\theta = \theta_L = .1$, $q = .1$), and Model 3 ($\theta = \theta_H = .4$, $q = 1.0$) to the corresponding variance for Model 4 (no insurance), where q is the market maker's prior assessment that the amount of portfolio insurance, θ, equals θ_H.

average variance of the posterior distribution for Model 1 is between four and five times that of each of the other models.[16]

Posterior probability of high portfolio insurance. Figure 5 shows the posterior probability of a high amount of portfolio insurance in the market. Ultimately the posterior probabilities in Models 1 and 2 will converge to 1 and 0, respectively; Figures 5a and 5b show the cross-sectional distributions for these models at t = 1,...,100.[17] Figure 2b above shows that on average most of the mispricing in Model 1

[16]The difference between Model 1 versus the rest is less (more) pronounced in higher (lower) percentiles than the medians shown in Figure 4. For example, for those cases in which the market maker is unable to resolve much uncertainty about the terminal value in any of the four models (say, the 99th percentiles), the relative variance is not strikingly different from 1 (although the variance of Model 1 is substantially higher than that of the other models in absolute terms).

[17]Recall that in Model 3, the prior belief that there is high portfolio insurance is 1 since the market maker knows this with certainty, and in Model 4 the market maker knows of the absence of insurance with certainty.

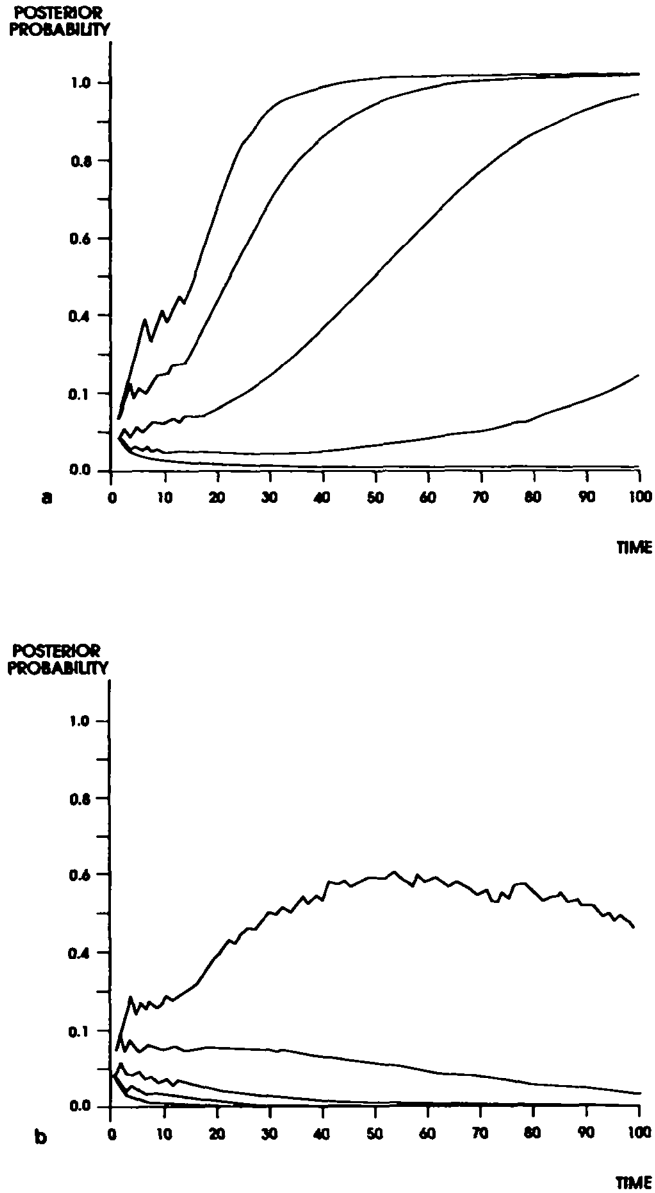

Figure 5
Plots of the distribution of the market maker's posterior probability of a high amount of portfolio insurance
From bottom to top, plots of 1st, 10th, 50th, 90th, and 99th percentiles of the distribution across 10,000 replications, at each date $t = 1,...,100$, for (a) Model 1 ($\theta = \theta_H = .4$, $q = .1$) and (b) Model 2 ($\theta = \theta_L = .1$, $q = .1$), where q is the market maker's prior assessment that the amount of portfolio insurance, θ, equals θ_H.

has disappeared by about $t = 70$; but Figure 5a shows that the median posterior probability of high insurance is only about .75 at $t = 70$. This uncertainty about whether the previous order flow is caused by informed traders or by insured traders is reflected in the higher variance of the posterior distribution of terminal value shown in Figure 4.

The closer correspondence between initial beliefs and reality for Model 2 (prior probability of high insurance .10, when in fact there is low insurance) translates in Figure 5b to much less uncertainty by date 70 than in Figure 5a. For example, the median posterior probability is close to 0; even the 90th percentile is less than .10.

Proportion of purchases versus sales. We end our comparison of the average differences in behavior of the four models by comparing the proportions across the 10,000 replications of purchases versus sales at each date $t = 1,...,100$. The informed will trade only if it is profitable for them to do so, given the market maker's quotes and the terminal value. The uninformed investors without portfolio insurance have equal probability of buying, selling, or not trading, while insured investors have a price-path-dependent probability of purchase or sale that implies purchases of stock when the price rises, and conversely.

Figures 6a to 6d give the proportions of purchases and sales for each of Models 1 to 4, respectively. It is verified, in Figure 6d, that in the absence of portfolio insurance, about one third of the potential trades are purchases and one third sales, once the market maker has inferred the terminal value from the order flow (that is, after about 20 trades in Figure 6d). Figures 6a and 6c show that our model of portfolio insurance results in larger purchases than sales initially, which then settle down to roughly equal proportions. However, Figure 6a shows an interim period from $t = 17$ to $t = 40$ in which sales exceed purchases, corresponding to the average price declines in Model 1.

Perhaps most revealing from Figure 6 are not the differences in trading patterns between Models 1 and 3 seen in Figure 6a and 6c, but rather the differences between the trading behavior of Models 2, 3, and 4 seen in Figures 6b, 6c, and 6d. There are very different patterns of buying and selling in these models; yet, as we have seen, there is little to distinguish between their resulting prices, in either means or variances. What is relevant in this framework is not the existence or quantity of portfolio insurance, but rather what is known by the market.

2.4 Individual price paths

In this section, we complement the average behavior shown in Section 2.3 by giving plots of the first 3 of the 10,000 replications for Models

1, 3, and 4.[18] In each case, the sequence of trading opportunities between informed versus uninformed is identical, with the uninformed trading as an insured investor as dictated by the model. Thus, for Models 1 and 3, exactly the same sequence of informed, uninformed, and insured traders has the opportunity to trade; the differences in price paths are driven solely by the different knowledge possessed by the market maker about the extent of portfolio insurance.

The replications are presented in Figure 7. We do not suggest that any of these plots is identical to the crash in 1987, and clearly there is much individual variation, depending to a large extent on the particular sample of buys and sells of the uninformed liquidity traders. Nevertheless, the plots serve to illustrate several of the results discussed above.

First, the most important result of our model is the average mispricing due to uncertainty about the extent of portfolio insurance, resulting in the average rise and fall in price for Model 1 shown in Figure 2 above. Replication 2 illustrates this behavior, since the price for Model 1 rises from 50 at $t = 1$ to 91.53 at $t = 13$, then falls to 56.64 at $t = 27$. Over the same interval, the highest prices for Models 2, 3, and 4 are 53.20, 67.76, and 53.15, respectively. There is another rise and fall for Model 1 from $t = 27$ to $t = 41$, from 56.64 to a high of 88.57 at $t = 34$ before falling to 48.24 at $t = 41$. The difference caused by uncertainty about the extent of portfolio insurance is illustrated by the comparable behavior of Model 3, which shows a similar but much more tempered price path over this interval.

Second, Model 1 results in higher uncertainty on the part of the market maker even beyond the period in which average price effects die out, and this is also illustrated by replication 2, since the price falls to a low of 12.97 at $t = 82$. This phenomenon of higher variance of the posterior distribution of terminal value for Model 1, which results in the greater reaction of the price path to purchases and sales in Model 1 versus Model 3, is also seen in replication 3 from, say, $t = 60$ to $t = 100$.

Third, replication 3 also illustrates the possibility that depending on the random purchase sequence of uninformed liquidity investors, the market maker may be mistaken about the terminal value even after 100 periods. In replication 3, the price at $t = 100$ is close to 100 for Model 4 (no insurance), showing that the market maker's incorrect inference about terminal value is clearly affected by the trades of uninsured liquidity traders. However, although it is possible to see

[18] Since the replications are independent, these three replications are chosen at random. We do not present the results for Model 2 to avoid cluttering the plots; however. they do not differ in significant ways from those of the zero-insurance Model 4.

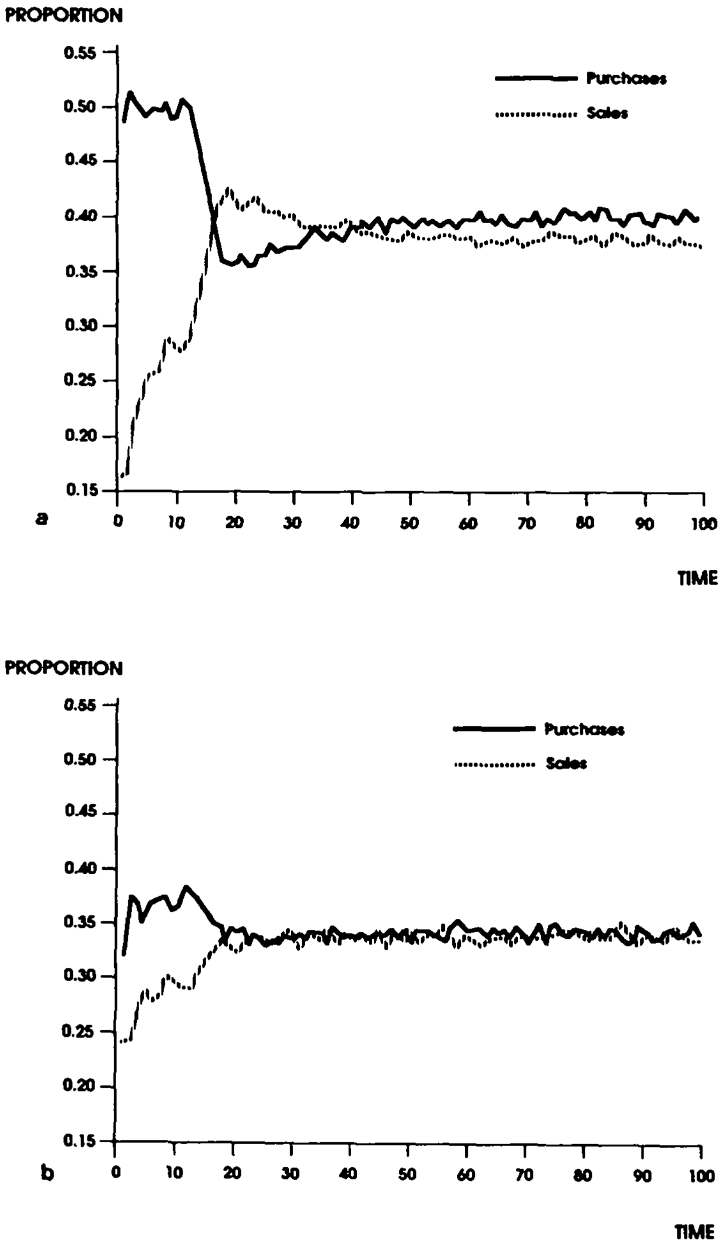

Figure 6
Proportions of purchases and sales
At each date $t = 1,\ldots,100$ for (a) Model 1 ($\theta = \theta_H = .4$, $q = .1$) and (b) Model 2 ($\theta = \theta_L = .1$, $q = .1$), where q is the market maker's prior assessment that the amount of portfolio insurance, θ, equals θ_H.

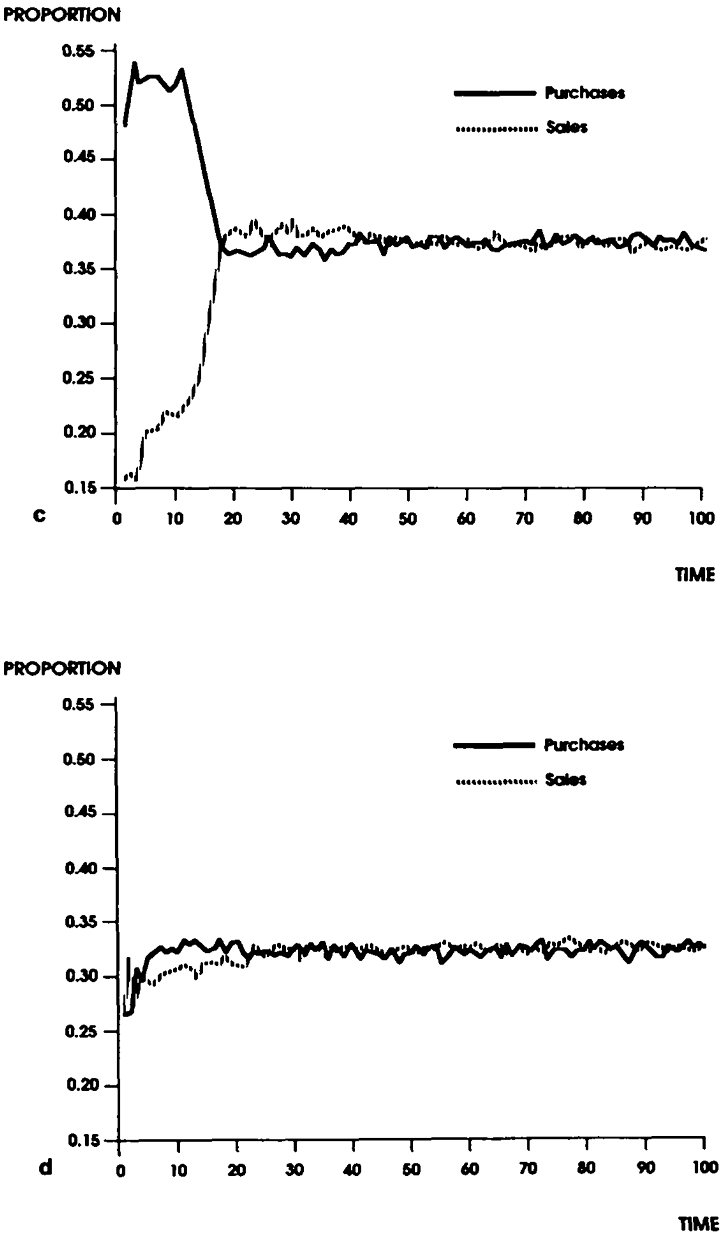

Figure 6 (continued)
Proportions of purchases and sales
At each date $t = 1,\ldots,100$ for (c) Model 3 ($\theta = \theta_H = .4$, $q = 1.0$) and (d) Model 4 (no insurance), where q is the market maker's prior assessment that the amount of portfolio insurance, θ, equals θ_H.

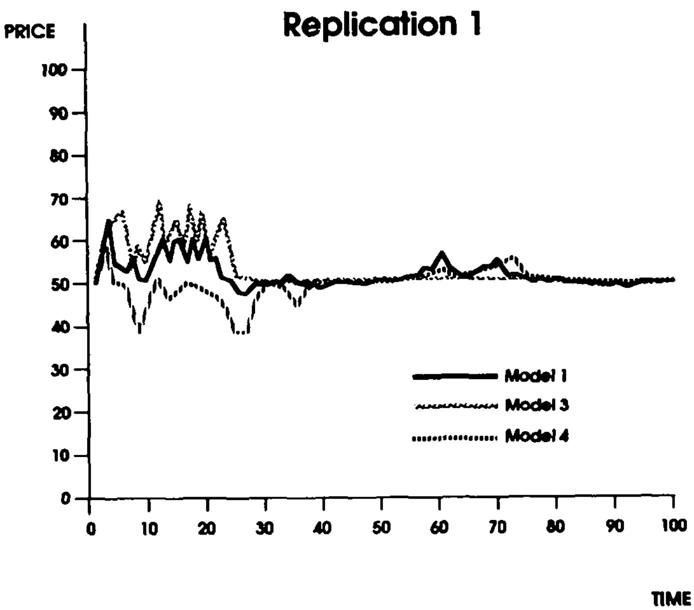

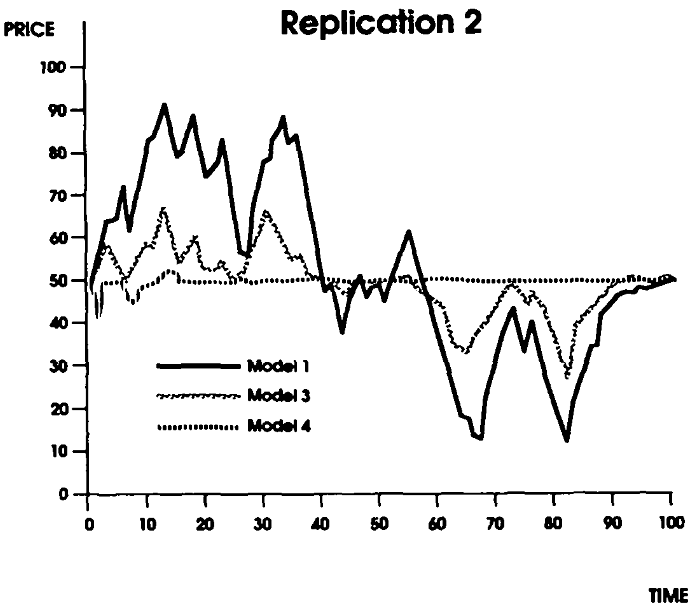

Figure 7
Plots of three replications
Each plot gives 100 corresponding observations for Model 1 ($\theta = \theta_H = .4$, $q = .1$), Model 3 ($\theta = \theta_H = .4$, $q = 1.0$), and Model 4 (no insurance), where q is the market maker's prior assessment that the amount of portfolio insurance, θ, equals θ_H.

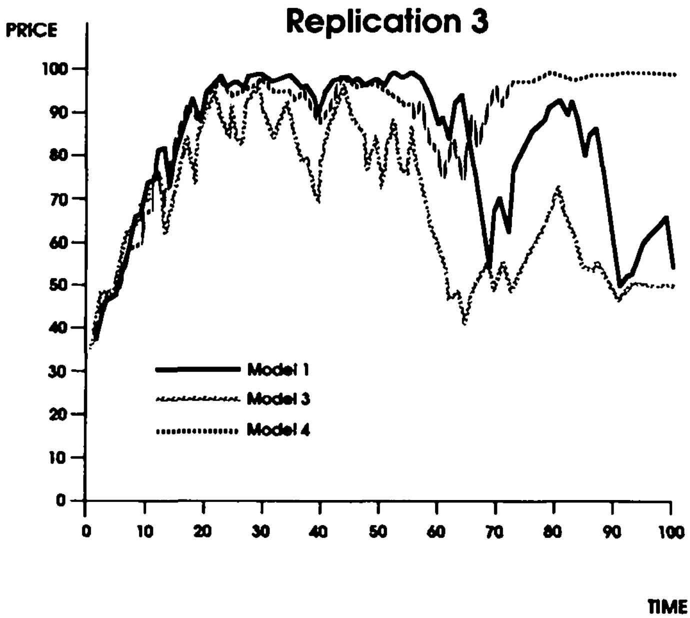

Figure 7 (continued)

this degree of error at t = 100 across all models, examination of the 99th price percentiles shows that it occurs more frequently for Model 1.

These plots illustrate the basic features of our model, but the reader should not interpret them as indicative of the relative frequency of the possible price paths illustrated there. Rather, the cross-sectional results in Section 2.3 are to be relied on for the likely behavior caused by underestimation of portfolio insurance within our model.

3. Caveats, Empirical Implications, and Previous Crashes

We now consider some limitations of our application of Glosten and Milgrom's model and then discuss the empirical implications of our findings and the applicability of our model to previous crashes.

3.1 Limitations of our model

The market maker in Glosten and Milgrom's model plays the role of aggregator of information and must infer the relevant information from the order sequence. In our application, this information includes the proportion of uninformed traders who are using portfolio insurance strategies, as well as the terminal value that the informed know with certainty. The limitation that the informed can trade only one share at a time is necessary to protect the market maker and give the market maker the incentive to provide his services.

While this model is useful for understanding a possible role for portfolio insurance in the crash of 1987, it is not perfect in two main aspects. First, the timing of the relevant trades is not accurate, since if each trading period is equally long then we fail to capture the huge relative volume on October 19, 20, and 21 in particular. We could make an assumption concerning the number of trading opportunities per calendar time interval, and accelerate the trading after portfolio insurance selling begins, but there seems little point in doing this given the arbitrary nature of the exercise. The cost is that our results show "crashes" in the sense that the price rises above the level that would be justified by fundamentals if the extent of portfolio insurance were known with certainty, and subsequently falls to that level; but the speed with which that fall occurs is limited in our model by the assumption of a uniform trading rate over time.

Second, the assumption that the informed know the terminal price with certainty and the uninformed know nothing is, of course, stylized. Alternatively, informed traders could observe the terminal value with noise, and submit orders that condition on their personal information and the information revealed in the price. In this case, purchases by insured traders will cause similar confusion as in our model, but the mechanism will be slightly different. On average, informed individuals will receive information that is bad relative to what appears to be revealed by the price, but given their beliefs about the amount of portfolio insurance, they will rationally infer that other informed traders have received better information than they. The analysis then follows our model: price declines reveal the amount of portfolio insurance, informed investors rationally change their inferences about previous purchases, and the price now correctly aggregates their information. This analysis would be consistent with claims by investors that they believed prior to the crash that the market price was high, given their personal information.

3.2 Empirical implications

The closest analysis of the crash to ours is that of Gennotte and Leland (1990), but the mechanism and empirical implications of the two models are quite different. First, our model predicts that if the extent of portfolio insurance were revealed by the large portfolio-insurance-based sales on Friday, October 16, then investors would infer at that time that the previous prices were too high. Roll (1989) notes that prices fell in overseas markets prior to the U.S. open on October 19, and the SEC (1988, p. 2-10) reports that the overhang of portfolio insurance trades on October 16, "*and investor knowledge of it,* may have contributed to the rapid decline on the afternoon of the 16th and undoubtedly added to the selling pressure on the 19th" (emphasis added).[19] In Gennotte and Leland's model, it is important that the amount of portfolio insurance not be known prior to trade on October 19.[20]

Second, our model implies that the price just prior to the crash was higher than the price that would have existed, given fundamentals, had the extent of portfolio insurance been known with certainty. Gennotte and Leland's model implies that the price just prior to October 19 was correct, given the fundamentals, although the amount of portfolio insurance was underestimated by the market. As reported by the Brady Commission (1988, chap. 3), traditional measures of market value such as price/earnings ratios indicated that prior to the crash the market was priced high relative to normal levels of such measures. Although not conclusive evidence for our approach, the relatively high level of these measures is clearly consistent with our analysis.

3.3 Previous market crashes

The introduction of new investor strategies that are not fully recognized by the market underlies our analysis of the October 1987 crash. Below, we discuss the possible application of our model to previous market crashes in 1929 and 1962.

Gennotte and Leland (1990) identify a 1929 counterpart to portfolio insurance, namely, stop-loss strategies, and assume a misestimation of the extent of those stop-loss strategies. The 1929 crash shares several characteristics with the 1987 crash that are consistent with our analysis under this assumption. First, both crashes followed periods of extended market rises. Second, large price declines immediately preceded the crash in both 1929 and 1987, but there was no direct information event of sufficient magnitude to explain the crash itself. Third, both market crashes were permanent (that is, prices did not return to precrash levels for months).

The market also "broke" in May 1962, with the same search for a cause that followed 1987.[21] However, the NYSE report on the 1962 crash focuses on the role of sales to meet margin calls and makes no reference to the introduction of new products or strategies that would be required in our model. We note that the 1962 crash did not follow an extended period of price rises but was "[p]receded by a severe five-month decline in stock prices" [NYSE (1963), p. 2]. Further, the

[19] If there is positive correlation across economies, then investors in countries other than the United States will infer information from the order flow in U.S. markets. In this setting, a straightforward extension of our model would imply that it is quite possible to have contemporaneous market breaks worldwide even though portfolio insurance was concentrated in the United States.

[20] See, for example, their chronology of events on October 16 to October 19 (1990, pp. 1011-1012).

[21] Much of the main report by the NYSE (1963), "The Stock Market under Strew The Events of May 28, 29 and 31, 1962," is strikingly similar to the reports of 1987.

crash was not permanent-the price recovered to its previous level three days (two trading days) after the decline of May 28.

4. Concluding Remarks

In this article, we examine the role played by formal and informal dynamic hedging strategies, or "portfolio insurance," in the market crash of October 1987. We show that if the market's prior beliefs underestimate the extent of portfolio insurance, then the price will be greater than that which would be implied by fundamentals if the extent of portfolio insurance were known with certainty. Over time, the market learns of the amount of portfolio insurance, and consequently reevaluates the previous inferences drawn from purchases that were erroneously regarded as based on favorable information. The result is that the price falls when the amount of portfolio insurance is revealed.

Our analysis strongly indicates that it is not portfolio insurance per se that causes problems, but rather a lack of information by the market about the extent of trading based on such strategies. A similar conclusion is reached by Gennotte and Leland (1990). This is quite different from an assumption of irrationality on the part of some or all investors, such as made by Shleifer and Summers (1990) who assume as "an alternative to the efficient markets hypothesis" that "some investors are not fully rational" (p. 19). It is ironic that their motivation for this approach is that "the stock in the efficient markets hypothesis-at least as it has traditionally been formulated-crashed along with the rest of the market on October 19, 1987" (p. 19). Our results suggest that lack of perfect information can explain the salient features of the 1987 crash, without recourse to assumptions of investor irrationality.

References

Amihud, Y., H. Mendelson, and R.A. Wood, 1990, "Liquidity and the 1987 Stock Market Crash," *Journal of Portfolio Management,* 16, 65-69.

Black, F., 1988, "An Equilibrium Model of the Crash," in S. Fischer (ed.), *NBER Macroeconomics Annual 1988,* MIT Press, Cambridge, Mass., pp. 269-275.

Brennan, M.J., and E.S. Schwartz, 1989, "Portfolio Insurance and Financial Market Equilibrium," *Journal of Business,* 62, 455-472.

Brown, D.P., and R.H. Jennings, 1989, "On Technical Analysis." *Review of Financia1 Studies,* 2, 527-551.

Easley, D., and M. O'Hara. 1991, "Order Form and Information in Securities Markets," *Journal of Finance,* 46, 905-927.

Gennotte. G.. and H. Leland, 1990, "Market Liquidity, Hedging and Crashes," *American Economic Review,* 80, 999-1021.

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.

Grossman, S.J., 1988. "Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies," *Journal of Business,* 61, 275-298.

Grundy, B.D., and M. McNichols, 1989, "Trade and the Revelation of Information through Prices and Direct Disclosure," *Review of Financial Studies,* 2. 495-526.

Miller, M.H.. 1990, "Financial Markets: Some Current Problems and Future Prospects," manuscript, Graduate School of Business, University of Chicago.

Miller, M.H., et al., 1988. *Final Report of the Committee of Inquiry Appointed by the Chicago Mercantile Exchange to Examine the Events Surrounding October 19, 1987.*

Mitchell, M.L., and J.M. Netter, 1989, "Triggering the 1987 Stock Market Crash: Antitakeover Provisions in the Proposed House Ways and Means Tax Bill?" *Journal of Financial Economics,* 24, 37-68.

New York Stock Exchange, 1963. *The Stock Market under Stress: The Events of May 28, 29 and 31. 1962.*

Presidential Task Force on Market Mechanisms (Brady Commission), 1988. *Report,* U.S. Government Printing Office. Washington, D.C.

Roll, R.W., 1989, "The International Crash of October 1987," in R.W. Kamphuis, Jr., R.C. Kormendi, and J.W.H. Watson (eds.), *Black Monday and the Future of Financial Markets,* Mid-America Institute for Public Policy Research Inc., Irwin, Homewood, Ill.

Schwert, G.W., 1990, "Stock Volatility and the Crash of '87," *Review of Financial Studies,* 3, 77-102.

Securitiesand Exchange Commission, 1988. *The October 1987 Market Break: Report by the Division of Market Regulation,* U.S. Government Printing Office, Washington, D.C.

Shleifer, A.. and L.H. Summers, 1990, "The Noise Trader Approach to Finance," *Journal of Economic Perspectives,* 4, 19-33.