LIQUIDITY AND FINANCIAL MARKET RUNS*

Antonio E. Bernardo and Ivo Welch

We model a run on a financial market, in which each risk-neutral investor fears having to liquidate shares *after* a run, but *before* prices can recover back to fundamental values. To avoid having to possibly liquidate shares at the *marginal* postrun price—in which case the risk-averse market-making sector will already hold a lot of share inventory and thus be more reluctant to absorb additional shares—each investor may prefer selling today at the *average* in-run price, thereby causing the run itself. Liquidity runs and crises are not caused by liquidity shocks per se, but by the *fear of future* liquidity shocks.

In contrast to the *financial institutions* literature (e.g., Diamond and Dybvig [1983]), runs on *financial markets* have not been a prime subject of inquiry. Our paper offers such a model, in which investors fear (but do not necessarily experience) future liquidity shocks. This creates two scenarios. In the good scenario, a risk-neutral public holds most of the risky shares. Investors hit by a liquidity shock in the future will sell to the risk-averse market-making sector at a "low-inventory price," which will be close to the risk-neutral value of the asset. In the good scenario, the market-making sector provides the public with low-cost insurance against liquidity shocks.

In the bad scenario, every investor conjectures that other investors intend to sell today, thus causing a "run." By joining the pool of selling requests today, an individual investor can expect to receive the average price that is necessary to induce the market-making sector to absorb all tendered shares today. The investor's alternative is to not enter the pool and instead to hold onto the shares. In making this decision, this investor is better off *if* he can wait out the storm and realize the eventual expected asset value. However, if he were randomly hit by the possible liquidity shock, this investor would need to sell his shares *behind* the rest of the public. But, with the market-making sector already holding the shares of other investors who had joined the run, this postrun price will be worse than the average in-run price today. If the average in-run price is greater than the expected payoff achieved by waiting, this investor will join the herd and also sell into the run. If other investors act alike, the conjecture that other inves-

tors sell today ends up true. In the bad scenario, the market-making sector inefficiently holds too many shares and provides the public with high-cost insurance against liquidity shocks.

Our bad scenario relies on two critical assumptions. **Our first critical assumption is that execution order is not perfectly sequential.** If, instead, execution were perfectly sequential, investors could not expect to avoid a place in the rear by joining the selling pool. Thus, the last investors (who would know they are last) would be better off waiting rather than joining the herd, and the bad scenario would unravel.

In reality, financial markets lack perfectly sequential execution in at least three circumstances. First, there is often no sequential execution after a market closure: for example, at the stock market opening or after a trading halt, markets are often conducted in a "batch" mode where all orders are crossed at the same price—and, indeed, fears of stock market runs seem higher around market closures. Second, there is anecdotal evidence that sequential execution broke down in the three most recent stock market crashes (1929, 1987, 1997). Describing the 1987 crash, Greenwald and Stein [1988, p. 15f] state that "investors cannot know with any precision at what prices their orders are executed . . . trades consummated only minutes apart were executed at wildly different prices, so that an investor submitting a market order had virtually no idea where it would be completed." Not knowing his place in the queue, a tendering investor would expect to receive some average price. The chain of perfect sequentiality may not just be broken on the exchange itself, but also in the communication of brokers with the exchanges and with their investors. In September 9, 1998, the S.E.C. Staff Legal Bulletin No. 8 describes the situation during the crash of 1997:

> The Commission received several complaints from customers regarding broker-dealer operations during the heavy trading volume effected on October 27 and 28. In particular, customers complained about receiving poor or untimely executions from broker-dealers . . . Numerous customers of online broker-dealers were unable to gain timely access to their accounts on October 28 . . . .
> Some broker-dealers experienced disruptions in their trading system operations that prevented them from routing customer orders to the designated market center for execution on a timely basis. Although disruptions occurred throughout the day, capacity problems were particularly high just after the opening of the market and prior to the close of trading on October 28. An unusually high numbers of orders that queued up in broker-dealers'

internal order handling systems prior to the opening of trading on October 28 helped precipitate these problems.

Third, in many over-the-counter financial markets, counterparties need to be found, and when multiple sellers are searching for counterparties, there is randomness as to who will find the potential buyers first.

Our second critical assumption is that the market-making sector is risk-averse and cannot expand infinitely in an instant. Consequently, selling pressure from individual investors causes prices to fall in the short run. If, instead, the market-making sector could expand infinitely, there would be no price impact if a large part of the investing public simultaneously wanted to exit. "Coming later" would then not impose a penalty on liquidity-shocked investors, and, again, the bad scenario would unravel.

Our assumption precludes the presence of enough standby investors who could eliminate any time lag between the exit of liquidity-shocked investors and the entry of more market-making capacity. Indeed, our bad scenario ends with the inflow of such investors—who *do* earn a positive rate of return *commensurate* with their willingness to provide liquidity and bear risk. The important question is only whether the time to reentry is nontrivial.

Consistent with our assumption of short-run "price pressure," there is evidence that stock price changes were negatively correlated with market-maker buying activity at half-hour intervals during the crash of October 1987 [Gammill and Marsh 1988]. Even in normal markets, Madhavan and Smidt [1993] and Hasbrouck and Sofianos [1993] find inventory effects on prices for NYSE specialists *after* controlling for the information content of order flow.

There is also anecdotal evidence that the lag can be nontrivial in situations in which the dynamics of our model can reasonably apply—a "zag" following the "zig" of a sharp stock-market-wide drop. Schwert [1998] reports all large daily S&P500 increases and decreases from 1885–1997. The aforementioned stock market crashes of 1929, 1987, and 1997 feature prominently. The single largest daily S&P500 point increase (+44.86) occurred on 10/28/1997, one day after the single largest daily S&P500 point decrease (−64.65). His daily returns even understate the rapid crash-and-recovery dynamics. Figure I shows the
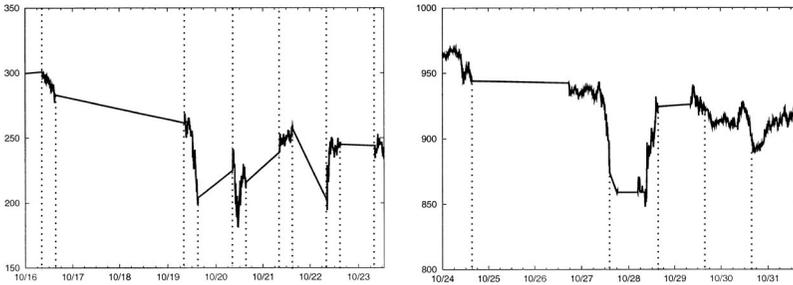
FIGURE I

Intraday S&P500 Futures during the 1987 and 1997 Stock Market Crashes

The data are described in Fair [2002]. The full vertical line marks 8:30 a.m. The half vertical line marks 16:00. There is no trading on weekends. By 1997 overnight futures trading provided data while the ordinary stock market was closed.

intraday behavior of the S&P500. Half of this 1997 crash occurred during the last hour of trading on 10/27, and half of the recovery occurred during the first hour of trading on 10/28.

The biggest daily percentage loss (−20 percent) occurred on 10/19/1987–and was soon followed by the seventh-best gain (+9 percent) on 10/21. Again, Figure I shows that the daily statistics do not fully convey the rapid crash-and-recovery dynamics. At 9:30 a.m. on 10/20 the S&P stood at 228. At 11:01 a.m. it had dropped to 181.5. At 12:04 it had returned to 212. Similarly, on 10/21 the market closed at 257. On 10/22, at 8:30 a.m. the market opened at 202, and immediately fell to 195. But by 8:38 a.m. it had recovered to 229, and by 9:47 a.m. it had further increased to 243.

There are no intraday data for 1929. However, Schwert [1998] reports that the second-worst and third-worst percentage drops in one day, −12 percent and −10 percent, occurred on 10/28 and 10/29/1929—and were promptly followed by the second-best percentage rise, +13 percent on 10/30.

It is not easy to get rich in crash situations. It is costly to create buffer stock or standby liquidity and the uncertainty in execution makes exploiting the bottom difficult. A large arbitrageur might have to lurk for years with ample financial slack to profit from the rare crash [Greenwald and Stein 1988, p. 19].[1]

1. Gammill and Marsh [1988] describe the broader market-making sector during the 1987 crash in great detail. Dennis and Strickland [2002] find that institutions are more likely than retail investors to sell into a dropping market. Indeed, portfolio insurers even precommit to such strategies. Amihud, Mendelson, and Wood [1990] document the liquidity decline during the 1987 crash, and

Thus, the market-making sector may be smaller than often assumed. In sum, we believe liquidity runs and crashes to be sufficiently rare phenomena that moderating market forces may not be sufficiently profitable to take effect *instantly*.

Our model produces a short-run accelerator effect, whereby small changes in the likelihood of a liquidity shock can have big effects on the allocation of risk and the equilibrium prices. It is important to point out that in our model prices and market-making inventories are driven by the fear of future liquidity shocks, not by the liquidity shocks themselves. Liquidity shocks might loom in the *future* but cause a run *today*. If underlying exogenous parameters change, high volatility and runs (low prices, high market-making inventory) can appear and disappear many times before the liquidity shocks themselves. An empiricist might not even necessarily recognize the relevance of actual liquidity constraints.

We also examine a model of liquidity shocks determined endogenously via margin constraints. In this model, fears of a margin call in the future can be self-fulfilling because they prompt investors to sell shares today in advance of the crowd. In sum, we believe that our model's intuition that investors fear "coming in last" is solid, and resonates with many who witnessed recent market crashes.

I. THE BASIC MODEL SETUP

Date		
Investors Trade.	Possible Liquidity Shock.	Unshocked investors return.
Focus of our Paper.	Shocked Investors	Asset value is revealed.
	Forced to Trade.	
$t = 0$	$t = 1$	$t = 2$

Our model has three dates ($t = 0, 1, 2$) and two assets paying off at date 2: a risk-free bond in infinitely elastic supply with a gross payoff of \$1 and a risky "stock" in finite supply with gross random payoff of $\bar{Z}$. We normalize the date 0 and date 1

---

describe that "Orders could not be executed, and information on market conditions and on order execution was delayed. Consequently, much of the burden of responding to the unexpected order flow fell on the exchange specialists, market-makers, and other traders with immediate access to the trading floor." For example, Goldstein and Kavajecz [2003] document that ordinary liquidity providers tend to withdraw during extreme market movements.

price of the bond to be $1, and solve for the price of the stock. The stock trades at date 0 and date 1. There are two types of traders in our model.

Market-makers constitute an entire sector, not just the specialist, but all traders willing to absorb shares upon demand and without fear of liquidity shocks. The market-making sector is assumed to be risk-averse in aggregate. Consequently, the price at which this sector is willing to hold shares is decreasing in its inventory [Garman 1976; Ho and Stoll 1981]. It holds zero inventory of shares at date 0.

Individual investors are atomistic, identical, risk-neutral, and endowed with shares which sum to the total supply of shares (normalized to one). They can suffer a liquidity shock at date 1. In Section II each investor may be forced to liquidate her shares with an exogenous probability s. In Section III margin constraints endogenize the date 1 liquidation probability s to depend on the date 0 stock price.

To recap, there are two important differences. First, individual investors are risk-neutral, and the market-making sector is risk-averse. This captures the fact that the investing public has more risk-absorption capacity than the market-making sector, and that, in a Pareto-efficient outcome, shares should be held by the investing public. Second, only individual investors face a potential liquidity shock at date 1. To purchase shares during a run, an investor must be classified as a member of the market-making sector. In subsection IV.A we allow a limited number of individual investors who learn that they are not subject to any future forced liquidation to join and thereby deepen the market-making sector. In one sense, the question as to why the market-making sector does not expand infinitely in an instant is similar to why banks in Diamond and Dybvig [1983] do not quickly find additional backers to avoid inefficient liquidation.

Our investors submit market orders and are unsure of the exact price at which the order will be executed.[2] The price is determined by assuming that the risk-averse market-makers earn zero expected utility at each trading date. Again, to prevent

---

2. Limit orders could potentially deepen the market-making sector but are not considered in our model. Greenwald and Stein [1988, footnote 16] also note that "limit orders do not represent an especially attractive alternative under the conditions of October 19 and 20. An investor's threshold price should depend on his most current information, which includes the current market price. Under very volatile conditions, this can mean resubmitting limit orders on an almost continuous basis, which would have been extremely difficult to accomplish."

an infinitely deep market-making sector and for simplicity, we assume that the market-makers ignore future trading opportunities when setting prices (i.e., behave myopically).[3] There are two interpretations for this price determination—and both lead to the same price functions and thus identical solutions in our model. In the first interpretation, individual investors submit sell orders at each trading date that are batched and bought by a single market-maker at a single average zero-expected-utility price. The zero expected utility condition is plausible, e.g., if the market is contestable [Baumol, Panzar, and Willig 1982]. In this case, a single market-maker may find it profit-maximizing to set prices competitively to preclude entry. For ease of exposition, our paper proceeds under this interpretation. In the second interpretation, individual investors submit sell orders at each trading date which are filled sequentially but randomly by N competitive market-makers with identical risk aversion. Being myopic, each market-maker fills $1/N$ of every order at a price that makes him indifferent between buying and not buying the last share of each order. Again, later orders are filled at lower prices. If the order size is arbitrarily small, the competitive market-makers achieve zero expected utility (i.e., do not earn inframarginal rents) and the "average" price received by the individual investors is identical to that in our first interpretation.[4] In our basic model, we assume no entry into the market-making sector between date 0 and date 1. In reality, we would expect that large deviations from the fundamental value of the stock caused by limited market-making depth would attract a new supply of liquidity. Arguably, it is precisely such a process that brings a financial market run to

---

3. Diamond and Verrecchia [1991] offer a price-setting mechanism that differs from ours in two respects: market-makers earn surplus on inframarginal trades, and they set prices by solving a dynamic optimization problem. The latter implies that market-makers forecast future buys and sells when setting today's price. Whether or not market-makers earn surplus is not important for the qualitative results of our model. Solving the market-makers' price function in a dynamic optimization problem is not tractable in our setup. Nevertheless, our results would still obtain, because the price at $t = 1$ in the event of a liquidity shock would still be lower than the price at $t = 0$.

4. This is an accounting identity, independent of the market-maker's utility (and price) function. Let $p(\theta)$ represent the price that makes the market-maker indifferent to the marginal order θ. If orders are arbitrarily small, executed sequentially (at price $p(\theta)$) and randomly, the individual investor receives the average price $p(\alpha) = 1/\alpha \cdot \int_0^\alpha p(\theta) \, d\theta$ when α proportion of investors submit sell orders. If the α orders are batched and purchased by the market-maker at price $\bar{p}(\alpha)$ which makes the market-maker indifferent between holding α shares and not holding these shares, then the market-maker's utility will be identical (in every future state) to the sequential execution case if he chooses $\bar{p}(\alpha) = 1/\alpha \cdot \int_0^\alpha p(\theta) \, d\theta$.

an end. We extend our basic model to allow for limited entry into the market-making sector in subsection IV.A. The important necessary feature of our model is that the market-making sector is exposed to fundamental stock price risk in the time between the short-run price pressure created by widespread investor selling and the entry of new liquidity.

II. Equilibrium with Exogenous Liquidity Shocks

II.A. Equilibrium Definition

Consider an individual investor who conjectures that a total of α shares will be sold by individual investors to the market-making sector at date 0, and let $p_0(\alpha)$ denote the date 0 price set by the market-makers when α sell orders arrive at date 0. If this investor also sells her shares at date 0, she will expect to receive the price $p_0(\alpha)$. However, if this investor chooses not to sell her shares at date 0, then at date 1 she will either (i) be forced to liquidate her shares, with probability s, or (ii) not be forced to liquidate her shares, with probability $1 - s$, in which case she will optimally wait to receive the expected value of the stock, μ, at date 2. If she suffers a liquidity shock, and fraction $q_1(\alpha)$ investors will also sell at date 1, her price will be $p_1(q_1(\alpha); \alpha)$, set by the market-makers when they already hold α shares of inventory and $q_1(\alpha)$ new sell orders arrive at date 1. Therefore, she optimally sells at date 0 iff

$$(1) \qquad p_0(\alpha) \geq s \cdot p_1(q_1(\alpha); \alpha) + (1 - s) \cdot \mu.$$

A risk-averse market-making sector implies that $p'(\cdot) < 0$, so the price will be lower when more investors are dumping their shares. The price at date 0 is higher than the price at date 1 (applicable if the investor suffers a liquidity shock), but lower than the expected value μ (applicable if she can hold on). Each individual investor's decision to sell depends critically on how many other investors are selling at date 0.

Define $F(\alpha)$ to be the expected net benefit of selling shares at date 0 (compared to not selling) when the investor conjectures that α shares are sold at date 0:

$$(2) \qquad F(\alpha) = \underbrace{p_0(\alpha)}_{\substack{\text{if tender} \\ \text{today}}} - \underbrace{s \cdot p_1(q_1(\alpha); \alpha)}_{\substack{\text{if forced to liquidate} \\ \text{tomorrow}}} - \underbrace{(1 - s) \cdot \mu}_{\substack{\text{if liquidation} \\ \text{not necessary}}}.$$

Then (i) waiting ($\alpha^* = 0$) is a pure strategy Nash equilibrium iff $F(0) \leq 0$; (ii) selling ($\alpha^* = 1$) is a pure strategy Nash equilibrium iff $F(1) \geq 0$; and (iii) $\alpha^* \in (0,1)$ is a mixed strategy Nash equilibrium iff $F(\alpha^*) = 0$.

As it turns out, $\alpha^* = 0$ is *never* a Nash equilibrium if $s > 0$. If the market-making sector holds zero inventory, it would be willing to accept the first share at the risk-neutral valuation today. Thus, the first seller would avoid the liquidation risk tomorrow without any price penalty today.

II.B. A CARA-Normal Example

In order to to solve algebraically for equilibrium prices, we now assume that (i) the stock payoff $\tilde{Z}$ is normally distributed with mean μ and variance σ^2, and (ii) the market-making sector has negative exponential utility $u(w) = -e^{-\gamma \cdot w}$, where γ is the coefficient of absolute risk aversion.

If the market-maker has initial wealth W_0 and ignores future trading opportunities, i.e., behaves myopically, his random wealth at date 2 is $\tilde{W}_2 = W_0 + \alpha \cdot (\tilde{Z} - p_0)$. The share price $p_0(\alpha)$ makes the market-maker indifferent between buying α shares at date 0 and maintaining zero inventory of shares:

$$E[-e^{-\gamma \cdot \tilde{W}_2}] = E[e^{-\gamma \cdot W_0}],$$

(3) $\quad \Rightarrow E[W_0 + \alpha \cdot (\tilde{Z} - p_0)] - \gamma \cdot \text{var}\,[W_0 + \alpha \cdot (\tilde{Z} - p_0)]/2 = W_0,$

$$\Rightarrow p_0(\alpha) = \mu - \gamma \cdot \sigma^2 \cdot \alpha/2.$$

We begin by assuming that liquidity shocks are perfectly correlated, in which case all investors want to sell the remaining ($1 - \alpha$) shares in period 2 with probability s. The market-maker already holds α shares, and the price $p_1((1 - \alpha); \alpha)$ makes the market-maker indifferent between buying ($1 - \alpha$) new shares at date 1 and maintaining an inventory of α shares:

(4) $\quad E[\tilde{W}_2 + (1 - \alpha) \cdot (\tilde{Z} - p_1)] - \gamma$

$\qquad \cdot \text{var}\,[\tilde{W}_2 + (1 - \alpha) \cdot (\tilde{Z} - p_1)]/2 = E[\tilde{W}_2] - \gamma \cdot \text{var}\,[\tilde{W}_2]/2,$

$$\Rightarrow p_1((1 - \alpha); \alpha) = \mu - [2 \cdot \alpha + (1 - \alpha)] \cdot \gamma \cdot \sigma^2/2.$$

THEOREM 1. If liquidity shocks are perfectly correlated, the unique symmetric Nash equilibrium is

(5) $\qquad \alpha^* = \begin{cases} (s/1 - s) & \text{for } s \leq \tfrac{1}{2}. \\ 1 & \text{for } s > \tfrac{1}{2}. \end{cases}$

Proof of Theorem 1. Substitute the pricing functions (3) and (4) into equation (2). Note that $F(0) > 0$ for all $s > 0$ and F_{α^*}, the derivative of F with respect to α^*, is negative. Thus, there are two possibilities. If $F(1) \geq 0$, then there is a unique pure strategy equilibrium $\alpha^* = 1$; and if $F(1) < 0$, there is a unique mixed strategy α^*, where $F(\alpha^*) = 0$. For $s > \frac{1}{2}$, $F(1) > 0$; thus, $\alpha^* = 1$. For $s \leq \frac{1}{2}$, solving for α^* yields the result.   QED

The equilibrium market-maker inventory increases in the liquidity shock probability s. Even though the efficient outcome would be for market-makers to hold zero inventory at date 0, the desire of investors to preempt other investors forces the risk-averse market-making sector to inefficiently hold shares. This inefficient allocation of risk is reflected in a lower equilibrium price for the stock.

The market-maker inventory is convex in the liquidation probability, which implies an "accelerator" effect: fear of other investors liquidating has an immediate influence on each investor's own decision to liquidate. For very small values of s, i.e., very little chance of future liquidity shocks, an investor sees other investors waiting and thus does not mind waiting herself. The market-making sector needs to hold almost no shares today (α^* close to zero), and the outcome is close to the Pareto optimum. With increasing s, the fraction of tendering investors rises ever more quickly. Similarly, the resulting volatility of stock returns $\tilde{R}_{0,1} \equiv (p_1 - p_0)/p_0$ is increasing and convex in s.[5] Thus, small changes in s can significantly change both market-maker inventory and market volatility. Finally, in the extreme, if there is "only" a 50-50 chance of investors facing a future liquidity shock, and even if the market-making sector is extremely risk-averse ($\gamma \to \infty$), risk-neutral investors find themselves unwilling to hold *any* stock today.

Although these are not distinct equilibria, there is a flavor of two distinct scenarios here: a good scenario, in which the probability of liquidation is low, and the market-making sector is not holding much inventory; and a bad (or run) scenario, in which the probability of individual liquidation is average, and the risk-averse market-making sector has to absorb *all* shares in the economy.

---

5. For $s \leq \frac{1}{2}$, stock return volatility is $\sigma(\tilde{R}_{0,1}) = (\gamma \cdot \sigma^2 \cdot s)/[2 \cdot (1 - s) \cdot \mu - \gamma \cdot \sigma^2 \cdot s]$. Volatility increases in s, γ, and σ^2 and decreases in μ.

The theorem readily generalizes to other correlations of investor liquidity shocks.[6] For example, if liquidity shocks are uncorrelated, the *Law of Large Numbers* implies that $s \cdot (1 - \alpha)$ shares will be sold for sure at date 1.[7] Replacing $(1 - \alpha)$ with $s \cdot (1 - \alpha)$ in equation (4) yields a date 0 equilibrium inventory of $\alpha^* = s^2/(1 - s)^2$.

Interestingly, with CARA utility, the risk-absorption capacity of the market-making sector (γ) and the riskiness of the stock (σ) play no role in the equilibrium market-maker inventory outcome (α^*). Expanding the market-making sector in both good and bad times would not solve the allocation problem created by the fear of facing a liquidity shock. The reason is that there are two countervailing forces when the market-making sector is deep: on the one hand, the average in-run price is higher because the market-making sector is close to risk-neutral. On the other hand, the marginal price obtained after the run is also higher. In other words, higher risk-capacity for the market-making sector not only allows investors to unload shares at an attractive price in a run, but it also allows them to enjoy a better price after a run. With CARA preferences, these two effects exactly offset each other in the investors' selling decision, because the market-making price is linear in inventory.[8] Although risk aversion and payoff variance affect the slope of the linear demand curve, they do not affect the relation between average and marginal prices. Thus, the trade-off between tendering today and waiting is independent of these parameters. The prime ingredient in this version of our model is investors' *fear* of future liquidation, s.

II.C. A CRRA-Binomial Example

Now assume that the market-maker has CRRA utility with risk-aversion parameter γ and the stock payoff is either U (with probability π) or D. Although we cannot solve this model algebraically, the numerical solutions illustrate richer comparative statics. We already know from the CARA case that it is not the

---

6. For arbitrary correlations ρ among investor liquidity shocks, replace $(1 - \alpha)$ with $(1 - \alpha) \cdot s(\rho)$ in (4). $s(\rho)$ is decreasing in ρ ($\rho = 0 \Rightarrow s(\rho) = s$, $\rho = 1 \Rightarrow s(\rho) = 1$). Then, $p_1((1 - \alpha); \alpha) = \mu - [2 \cdot \alpha + (1 - \alpha) \cdot s(\rho)] \cdot \gamma \cdot \sigma^2/2$, and $\alpha^* = [s \cdot s(\rho)]/\{1 + [s(\rho) - 2] \cdot s\}$.

7. Here, it is an undesirable model artifact that the *Law of Large Numbers* eliminates aggregate uncertainty. It would not be difficult to add other sources of uncertainty to eliminate the consequent arbitrage. Dow and Gorton [1994], Allen, Morris, and Shin [2002], Shleifer and Vishny [1997], and Liu and Longstaff [2000] deal with similar concerns.

8. Of course, when the market-making sector is deep, prices are close to risk-neutral even if no risk-neutral investor is willing to hold shares and thus the welfare loss is small.

steepness of the demand curve itself (i.e., the "depth") that matters to market-making inventory. But in the CRRA case, the other parameters (such as wealth, risk aversion, and riskiness) matter for the relative share allocations to the extent that they affect the curvature of the market-making demand function.

Figure II graphs the market-making sector's equilibrium holdings (α^*) as a function of exogenous parameters for the case of independent liquidity shocks across investors. (The numerical results are qualitatively similar when liquidity shocks are perfectly correlated across investors.) Typically, we find that the market-making sector holds more inventory (α^*)

- when the market-making sector has greater wealth;
- when the market-making sector has greater risk-absorption capacity (risk-aversion coefficient γ is lower);
- when the asset is less risky (when $U - D$ is smaller holding the mean payoff $\pi \cdot U + (1 - \pi) \cdot D$ constant);
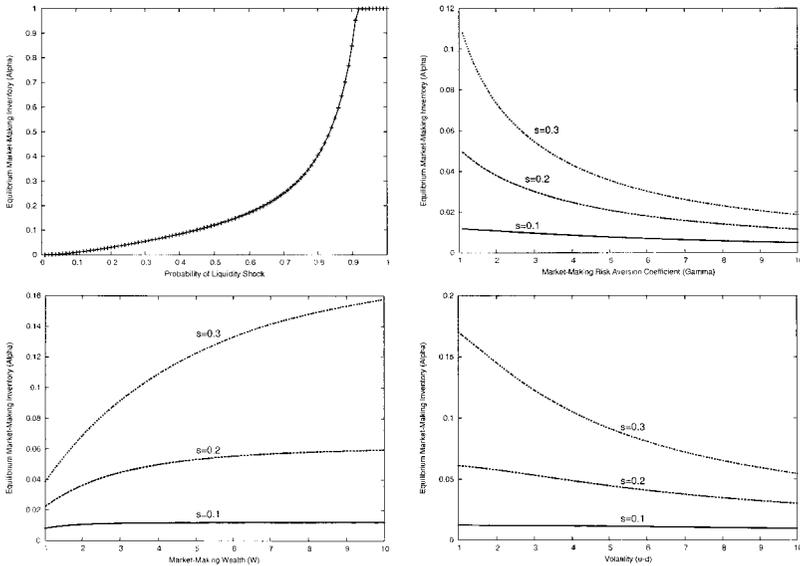- when the probability of a liquidity shock (s) is higher.

FIGURE II

Comparative Statics under Market-Making CRRA Utility

Comparative statics when investors face independent liquidity shocks. Our base parameters are a down-stock-value of $D = 10$ and an up-stock-value $U = 20$ with equal probability $\pi = 0.5$, a risk-aversion coefficient of $\gamma = 3$, and market-making wealth of $W = 1.5$ (i.e., roughly ⅒ of the value of the financial market).

III. Equilibrium with Margin Constraints

Margin calls, which force investors to sell more shares if the share price declines, are important during financial market crashes (see, e.g., Chowdhry and Nanda [1998]). Margin calls can endogenize liquidity constraints and can produce the very high-frequency "phase transitions" (as well as multiple equilibria) that characterize stock market crashes.

We now sketch a simple model of margin constraints. As before, the stock payoff is normally distributed with mean μ and variance σ^2 and the market-maker has CARA utility with risk-aversion parameter γ. Suppose that every individual investor has an external source of income at date 0 of $\tilde{\mathscr{W}}$, which is uniformly distributed over the interval $[0,B]$ and independent of any stock price movements, and she has financed her purchase of the stock with margin. At date 1, if her wealth (including the stock) has fallen by too much, a margin call will force liquidation. Otherwise, holding on will be optimal. Suppose that she had purchased shares at price p (given outside the model), and let $m \in [0,1]$ be the proportion of the investment financed with margin. Thus, if the price falls from p to p_0, she would need to come up with cash of $m \cdot (p - p_0)$ in order to hold onto the shares until the final period. Margin constraints are thus triggered by a *decline* in price from the purchase price p to p_0. Therefore, her endogenous probability of liquidation is

$$(6) \qquad s(\alpha) = \text{prob}\,[\tilde{\mathscr{W}} < m \cdot (p - p_0)] = \frac{m \cdot (p - p_0)}{B}$$

$$= \frac{m \cdot (p - \mu)}{B} + \left(\frac{m \cdot \kappa}{B}\right) \cdot \alpha \equiv c_0 + c_1 \cdot \alpha,$$

where $\kappa \equiv \gamma \cdot \sigma^2/2$, $c_0 \equiv m \cdot (p - \mu)/B$, and $c_1 \equiv m \cdot \kappa/B$ (>0). We consider only the case where the income shocks are perfectly correlated; i.e., *all* or *no* investors face a liquidity shock at date 1 depending on the realization of $\tilde{\mathscr{W}}$.[9]

THEOREM 2. If (i) $B > p > \mu > \kappa$, where $\kappa \equiv \gamma \cdot \sigma^2/2$; (ii) B is sufficiently large; and (iii) income shocks are perfectly corre-

---

9. The independent liquidity shocks case solution involves cubic equations. It is difficult to find the parameter restrictions ensuring a unique equilibrium so that we can do the comparative statics. However, there is no reason why the intuition of the equilibrium discussed in this section would not carry over to the independent liquidity shock scenario.

lated, then there is a unique tendering equilibrium, $\alpha^* \in$ (0,1), in which α^* increases in m and ($p - \mu$) and decreases in B.[10]

Proof of Theorem 2. In the perfectly correlated income shocks case, we have $p_0 = \mu - \kappa \cdot \alpha$ if α proportion tender at date 0, and $p_1 = \mu - \kappa \cdot (1 + \alpha)$ if all remaining investors are forced to liquidate at date 1. Substituting these two price functions and the liquidation probability (6) into (2) yields

$$(7) \qquad F(\alpha) \equiv \kappa \cdot [(c_0 + c_1 \cdot \alpha) \cdot (1 + \alpha) - \alpha].$$

First, note that $F(0) = \kappa \cdot c_0 > 0$, so $\alpha = 0$ is not an equilibrium. Second, note that $F(1) = \kappa \cdot [2 \cdot (c_0 + c_1) - 1] < 0$ by assumptions (i) and (ii), so $\alpha = 1$ is not an equilibrium. Finally, $F_\alpha = \kappa \cdot [c_0 + c_1 + 2 \cdot c_0 \cdot \alpha - 1] < 0$ by assumptions (i) and (ii), so there is a unique $\alpha^* \in$ (0,1) such that $F(\alpha^*) = 0$.

The comparative statics results follow from the facts that (i) $F_{c_0} = \kappa \cdot (1 + \alpha) > 0$; (ii) $F_{c_1} = \kappa \cdot \alpha \cdot (1 + \alpha) > 0$; (iii) c_0 is increasing in m and p and decreasing in B and μ; and (iv) c_1 is increasing in m and decreasing in B.   QED

The intuition for the comparative statics are straightforward.

Margin Constraint (m): the more investors can borrow, the greater will be the tendering proportion at date 0, because it is less likely that investors will be able to meet margin calls at date 1.

Original Purchase Price (p): the more investors paid relative to the current mean ($p - \mu$), the greater the tendering proportion at date 0, because it is less likely investors will be able to meet the margin call at date 1. One way to interpret $p - \mu$ is the innovation in beliefs from purchase of the stock to now. This states that if there is a big negative shock to beliefs, then margin calls exacerbate the price move through early liquidation: this is overreaction to bad news. However,

---

10. The restrictions (i) $B > p > \mu > \kappa$ and (ii) B sufficiently large, are made for the following reasons. The assumption $\mu > \kappa$ ensures that $p_0 > 0$ so we get a meaningful margin requirement. The assumption $p > \mu$ ensures that we are considering cases where the price of the stock has fallen since purchase, so that we can consider the effect of margin calls. Finally, the assumption B sufficiently large ensures an interior solution for α^*. If B is not large enough, the probability of not meeting the margin call at date 1 is very high, and the only equilibrium we get at date 0 is everyone tendering (which has no interesting comparative statics).

there is no overreaction to good news, because there is not a margin call in that case!

Expected Income (B): the higher external expected income is likely to be (to meet future margin calls), the smaller is the tendering proportion at date 0.

Ultimately, the only important aspect of our model is that a lower price can further increase the probability of future liquidation needs. The liquidity run phenomenon then interacts with and rationally amplifies feedback trading [Shleifer 2000]. Our margin assumptions have produced the particular linear mapping of price declines into liquidation probabilities in (6). Risk management systems, principal-agent problems, limited horizons, or empirical liquidation estimates could produce other mappings. If $c_0 = 0$ and $c_1 = 1$, there are three equilibria: a stable one in which no investor tenders and therefore no investor is afraid of liquidation; a stable one in which every investor tenders *because* every investor tenders; and an instable one in which there is an interior tendering equilibrium. Because these are simultaneously feasible, sudden equilibrium switches could potentially occur.

IV. DISCUSSION, EXTENSIONS, AND WELFARE

IV.A. *Preventing Runs*

What mechanisms could prevent the need for the market-making sector to absorb run inventory from the public?

The first answer lies in the enforcement of perfect sequentiality. With sequential execution the last investors (who now know they are the last investors!) would be better off just waiting it out instead of being the last in-the-run investors. In response to the 1987 crash, the NYSE massively expanded its communication infrastructure, a mechanism to prevent the conversion of the sequential market into a random-execution market in times of declines. In contrast, a widespread belief that front-running others is possible can encourage run equilibria, because successful front-running increases the expected payoff to tendering early.[11]

---

11. Brunnermeier and Pedersen [2002] demonstrate that a strategic trader optimally front-runs and sells when he knows another trader must sell and then buys back when the share price bottoms out. Within the context of the informational cascades literature [Bikhchandani, Hirshleifer, and Welch 1992; Welch 1992]; Chen 1995 has modeled such informational interactions in a banking run context.

Naturally, in an equilibrium with homogeneous agents, no one can expect to front-run anyone else. However, in the real world, some heterogeneous investors may rationally or irrationally believe in their ability to front-run. Indeed, portfolio insurance attempts to precommit to withdraw funds in the case of large moves, which will thus worsen the liquidity effects described in our own paper. Leland and Rubinstein [1988, p. 46f] describe some possible front-running in 1987: "With the sudden fall in the market during the last half hour of trading on October 16, many insurers found themselves with an overhang of unfilled sell orders going into Monday. In addition, several smart institutional traders knew about this overhang and tried to exit the market early Monday before the insurers could complete their trades."

The second answer lies in providing liquidity during runs. For example, suppose that there is limited market-making entry at date 1, so that it is deeper at date 1 than at date 0 (i.e., $\gamma_{t=1} \leq \gamma_{t=0}$). Again assume CARA utility, normally distributed payoffs, and perfectly correlated shocks. The prices now reflect the different market-making depth at each date; thus, $p_0(\alpha) = \mu - \gamma_0 \cdot \sigma^2 \cdot \alpha/2$, and $p_1(1 - \alpha; \alpha) = \mu - (1 + \alpha) \cdot \gamma_1 \cdot \sigma^2/2$. Substituting $p_0(\alpha)$ and $p_1(1 - \alpha; \alpha)$ into $F(\alpha)$ yields

THEOREM 3. If the market-making sector is deeper at date 1 (i.e., $\gamma_1 < \gamma_0$) and liquidity shocks are perfectly correlated, there is a unique symmetric Nash equilibrium with

$$(8) \qquad \alpha^* = \begin{cases} \left(\dfrac{s}{\gamma_0/\gamma_1 - s}\right) & \text{if } s \leq \gamma_0/(2 \cdot \gamma_1) \\ 1 & \text{if } s > \gamma_0/(2 \cdot \gamma_1). \end{cases}$$

Market-making inventory α^* decreases in γ_0 and increases in γ_1. If the market-making sector is more shallow at date 0 than at date 1, investors would be less eager to tender to market-makers and more inclined to take their chances. Thus, a shallow market-making sector in ordinary markets (high γ_0) can be as important as intervention in bad markets (low γ_1)! Conversely, if "standby liquidity" for a financial crisis is low (i.e., the market-making risk aversion γ_1 is unusually high), then the postrun price will be lower, which prompts investors to sell more at date 0.

Unfortunately, the ordinary private market-making sector could sometimes even behave as if it is *more* risk-averse during runs. Gammill and Marsh [1988] document that exchange market-makers, other exchange members, and option market-makers

indeed provided liquidity from October 15 to October 19, 1987, but then actually *sold* $606 million on October 20, 1987. (In fact, the market-makers' advantageous location could have allowed some of them to front-run external investors.) The biggest providers of liquidity on October 20 were pension funds, trading-oriented investors, firms repurchasing their own shares, and individual investors. Furthermore, many large institutions, other potential providers of liquidity, often run portfolio-insurance schemes which tend to sell more into a crash rather than against a crash. From October 15 to October 20, 1987, these portfolio insurers withdrew $853 million, $2,419 million, $5,223 million, and $2,148 million.

Greenwald and Stein [1988, p. 19] suggest one alternative private mechanism, in which large financial insurers would agree to cover some of the losses of market-makers if the market drops significantly. Still, we find it unlikely that the private sector could provide liquidity *only* in bad scenarios, but not in good scenarios. Therefore, government intervention which commits to provide market depth *in "bad" but not in "good" times* might usefully mitigate run inefficiencies: if correctly done, standby liquidity could help prevent many financial runs in the first place.[12]

IV.B. A Multiperiod View

Our model is multiperiod robust, so a liquidity run can occur even if the liquidity shock is far away. The two important conditions are only that investors must experience the liquidity shocks simultaneously and that the market-making sector must face the risk of being stuck with inventory that is subject to fundamental price risk. For example, suppose that investors have two opportunities to sell, date 0 and date 1, prior to the occurrence of a liquidity shock. An equilibrium is now a pair (α_0, α_1), for which—given that α_0 proportion sell at date 0—it is optimal for an α_1 proportion to sell at date 1, and vice versa. One condition for optimality is that someone who sells at one date does not have the incentive to deviate and sell at the other date. But there is only one case for which this is true: α proportion sell at date 0 and no one sells at date 1! In this case, the date 0 price exceeds the date 1 price so no one has an incentive to deviate and sell at date 1. Moreover, the possibility of a liquidity shock in the future makes

12. Indeed, this is the equivalent of the national petroleum reserves, which are rarely released, but whose presence may in itself prevent runs.

an investor indifferent between selling at date 0 and waiting if she conjectures that α^* (as in our earlier model) proportion of investors will sell at date 0.

The probability of a future liquidity shock may constantly fluctuate, even though the liquidity shock itself can be off on the horizon. Consequently, an empiricist could observe dramatic price movements and market-making inventory changes without observing any actual liquidity shocks. And, for the rare empiricist able to measure the *fear* of liquidity shocks (s), depending on its value, seemingly small changes can cause large sudden changes in the desire of investors to unload shares onto the market-making sector.[13]

Thus, time-varying probability assessments of *future* liquidity shocks could lead to active trading and time-varying market-making inventory adjustments, even in the absence of any *current* liquidity shocks. This model can potentially explain relatively high trading volume and price fluctuations in the presence of only mild news.

IV.C. The Social Cost of Investor Fear

In our model, there is no asymmetric information or trading costs—and yet the market outcome can be significantly worse than the Pareto-optimal allocation. In the CARA-normal case, we can compute the social cost of investor fear. Our benchmark is *not* a price of μ, but a requirement that risk-neutral investors must not sell at date 0 (similar to the analysis in Diamond and Dybvig [1983]). In this Pareto-optimal outcome, the risk-neutral investors hold all the shares at date 0 and sell to the market-making sector at date 1 *only if* they are actually hit by a liquidity shock. Assume that liquidity shocks are perfectly correlated,[14] so every investor would sell shares with probability s at a price $p_1 = \mu - \gamma \cdot \sigma^2/2$ (assuming that the market-maker sector executes these sell orders at a price that yields no utility gain for them) and would retain shares with probability $1 - s$ (with expected value μ). Thus, investors' utility would be

---

13. Although our model has emphasized purely rational behavior, where the fear of liquidity shocks is rationally assessed or derived from margin constraints, our equilibrium could also be embedded in a world of "nonrational behavioral economics," if the fear of a liquidity shock (the need to terminate an investment early during a market run) were itself nonrational.

14. When liquidity shocks are independent, the expressions and intuition are similar.

(9) $\mu - (\gamma \cdot s \cdot \sigma^2)/2.$

In contrast, in a financial run, risk-neutral investors sell with probability α^* at date 0 at the average price $p_0 = \mu - (\gamma \cdot \alpha^* \cdot \sigma^2)/2$, liquidate with probability $(1 - \alpha^*) \cdot s$ at date 1 at the average price $p_1 = \mu - (1 + \alpha^*) \cdot \gamma \cdot \sigma^2/2$, and retain shares with probability $(1 - s) \cdot (1 - \alpha^*)$ at expected value μ. Thus, investors' utility is

(10)
$$
\begin{cases}
\mu - \dfrac{\gamma \cdot \sigma^2}{2} \cdot \dfrac{s}{1 - s} & \text{if } s \leq \dfrac{1}{2} \\[2ex]
\mu - \dfrac{\gamma \cdot \sigma^2}{2} & \text{if } s > \dfrac{1}{2}.
\end{cases}
$$

By assumption, the market-making sector has zero expected utility gain, so a total welfare comparison only requires a comparison of the investors' utility. The equilibrium welfare (expected selling price) is below the Pareto-optimal level of welfare by the amount

(11)
$$
\begin{cases}
\dfrac{\gamma \cdot \sigma^2}{2} \cdot \dfrac{s^2}{1 - s} & \text{if } s \leq \dfrac{1}{2} \\[2ex]
\dfrac{\gamma \cdot \sigma^2}{2} \cdot (1 - s) & \text{if } s > \dfrac{1}{2}.
\end{cases}
$$

The welfare loss is increasing in the market-maker's risk aversion and the payoff variance σ^2, because inefficient risk-sharing is exacerbated. The welfare loss is greatest when $s = \frac{1}{2}$ because the market-making sector must absorb *all* shares, not just those of the liquidity-shocked individuals. Because α^* increases at a faster rate as s approaches $\frac{1}{2}$, the welfare loss increases in s for $s \in [0,0.5)$. However, because (i) $\alpha^* = 1$ for all $s \geq \frac{1}{2}$ and (ii) as s increases the market-makers would hold an increasing proportion of shares in the Pareto-optimal benchmark outcome, the welfare loss decreases in s for $s \in (0.5,1]$.

IV.D. Contagion

Contagion of liquidity *fear* across investors falls naturally out of the model. In the bad scenario, there are spillovers in the decisions of investors to sell their shares. For example, Schnabel and Shin [2002] argued that the forced sale of commodities by the de Neufville banking house culminated in the financial crisis of 1763. Limited liquidity in the commodity markets caused a decline in commodity prices which in turn put a strain on other

investors who were then also forced to sell to meet liquidity needs. Similar arguments have been made in regard to the collapse of the Long Term Capital Management Hedge Fund in 1998 which also led to the demise of other hedge funds with similar investment strategies [Brunnermeier and Pedersen 2002].

Contagion across markets and institutions, as modeled, e.g., in Allen and Gale [2000], also fall naturally out of our model. For example, if investors hold positions in many markets, liquidity needs in one market can lead to early liquidation of assets in other markets. This suggests that the nature of cross-liquidity constraints could be more important than previously thought. If financial markets investors were unaffected in a liquidity crisis, they could bolster failing financial institutions, and vice versa. It is the universality of liquidity fears across the economy that is important.

V. RELATED LITERATURE

We adapted the Diamond and Dybvig [1983] model into a financial markets setting in which prices can fluctuate, there is no sequential service constraint, no productive inefficiency if investors liquidate, and investors do not suffer complete losses if they fail to join the run.[15] Indeed, working out the endogenous pricing and market-making sectors' inventory is a major focus of our paper. Our models share the fear of liquidity shocks, the assumption of payoff externalities, and the conclusion that a "lender of last resort" can earn money or prevent the run or both.

This strengthens the argument in De Long, Shleifer, Summers, and Waldmann [1990] and Shleifer and Vishny [1997], where potential arbitrageurs may not find it in their interests to trade *against* noise traders. Their arbitrageurs fear future adverse noise trader risk (liquidity shocks), because they may have to liquidate if mispricing further increases. In our model, potential arbitrageurs may even find it in their interest to trade *with* noise traders. The risk of experiencing future liquidity shocks themselves can actually induce potential arbitrageurs to sell

---

15. Allen and Gale [2003] and similar financial contagion models, though quite different, also adopt the Diamond and Dybvig [1983] framework, as does the liquidity crisis and international runs on currency reserves literature, e.g., Caballero and Krishnamurthy [2001]. Instead of their assumptions, our model had to assume a division of the economy into a sector willing to absorb shares in crashes and a breakdown of perfectly sequential execution.

(rather than purchase) when and because other possibly irrational investors are also selling. Closely related to this idea, Brunnermeier and Pedersen [2002] present a model in which traders strategically sell ahead of other large traders who need to reduce their positions to meet margin requirements. This leads to price overshooting and, as in our model, the market becomes illiquid just when liquidity is most needed.

There is also related theoretical literature on stock market crashes. Grossman and Miller [1988] present a two-period model in which the limited risk-bearing capacity of market-makers causes prices to fall excessively in response to an order imbalance in the first trading period but then return to fundamental values in the second trading period when individual traders can come in and supply liquidity. Greenwald and Stein [1991] extend the Grossman and Miller [1988] analysis so that traders can only submit market orders in the second trading period. This introduces transactional risk (uncertainty about the price at which their trades will execute) which reduces the willingness of traders to absorb the market-makers inventory in the second trading period. Knowing this, market-makers demand an even larger risk premium than in the Grossman-Miller analysis to absorb any temporary order imbalances in the first trading period.

Another large and related literature examines the impact of portfolio insurance (e.g., Grossman [1988], Brennan and Schwartz [1989], Genotte and Leland [1990], Jacklin, Kleidon, and Pfleiderer [1992], Donaldson and Uhlig [1993], Grossman and Zhou [1994], and Basak [1995]). Portfolio insurers are usually modeled as agents who display positive feedback trading (of an accelerating kind) for exogenous (often assumed) reasons. This literature's primary goal is to show that portfolio insurers can exacerbate crashes.

Other papers have also presented ingenious mechanisms that can elicit large price changes. In Madrigal and Scheinkman [1997] an informed strategic market-maker attempting to control both the order flow she receives and the information revealed to the market by the prices she sets may choose an equilibrium price schedule that is discontinuous in order flow, thus prompting large changes in price for arbitrarily small changes in market conditions. In Romer [1993] uncertainty about the quality of others' information is revealed by trading, and large price movements, such as the October 1987 crash, may be caused not by news about fundamentals but rather by the trading process itself. In San-

droni [1998], market crashes can be a self-fulfilling prophecy when agents have different discount rates and different beliefs about the likelihood of rare events (even if these beliefs converge in the limit). Barlevy and Veronesi [2001] present a model in which uninformed traders with rational expectations have locally upward sloping demand curves which can generate an equilibrium price function discontinuous in fundamentals.

Finally, we are not the first to employ margin constraints to generate (multiple) equilibria. In Chowdhry and Nanda [1998], perhaps the paper most similar to our own endogenous liquidity constraint section, some investors engage in margin borrowing to obtain their desired investment portfolio. Because shares can be used as collateral, there is a link between the price of the stock and the capacity to invest in it which introduces the possibility of multiple equilibria.[16]

VI. CONCLUSION

Our paper has extended the Diamond and Dybvig [1983] model of bank runs to the "other half" of the financial system. Fearing possible liquidation in the future, investors prefer to sell their shares today to avoid coming in last. Their fear causes a financial market run in which prices can fall precipitously, even when the changes in the likelihood of liquidation are small. In such settings, financial markets do not provide low-cost insurance against liquidity shocks. Our model's natural policy implications are for markets to expand communication and capacity in order to maintain trade sequentiality whenever possible, to avoid market closures during a financial run, to disallow front-running in crashes by market-makers and other close-to-the-market investors (who are designated to provide liquidity in ordinary situations), and to assure investors that the Fed will continue to play the same role as a liquidity-provider-of-last-resort that it played impromptu during the 1987 crash.

UNIVERSITY OF CALIFORNIA, LOS ANGELES
YALE UNIVERSITY

---

16. For a more recent example, Yuan [2000] demonstrates that margin constraints can be beneficial because they may apply to informed investors and thus reduce the adverse selection problem with uninformed investors.

REFERENCES

Allen, Franklin, and Douglas Gale, "Financial Contagion," *Journal of Political Economy,* CVIII (2000), 1–33.

Allen, Franklin, and Douglas Gale, "Financial Fragility, Liquidity and Asset Prices," manuscript, University of Pennsylvania, 2003.

Allen, Franklin, Stephen Morris, and Hyun Song Shin, "Beauty Contests, Bubbles and Iterated Expectations in Asset Markets," manuscript, University of Pennsylvania, Yale University, and London School of Economics, 2002.

Amihud, Yakov, Haim Mendelson, and Robert A. Wood, "Liquidity and the 1987 Stock Market Crash," *Journal of Portfolio Management,* XVI (1990), 65–69.

Barlevy, Gadi, and Petro Veronesi, "Rational Panics and Stock Market Crashes," manuscript, University of Chicago, 2001.

Basak, Suleyman, "A General Equilibrium Model of Portfolio Insurance," *Review of Financial Studies,* VIII (1995), 1059–1090.

Baumol, William J., John C. Panzar, and Robert D. Willig, *Contestable Markets and the Theory of Industry Structure* (San Diego, CA: Harcourt Brace Jovanovic, 1982).

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy,* C (1992), 992–1026.

Brennan, Michael J., and Eduardo Schwartz, "Portfolio Insurance and Financial Market Equilibrium," *Journal of Business,* LXII (1989), 455–472.

Brunnermeier, Markus K., and Lasse Heje Pedersen, "Predatory Trading," manuscript, Princeton University and New York University, 2002.

Caballero, Ricardo J., and Arvind Krishnamurthy, "International Liquidity Illusion: On the Risks of Sterilization," manuscript, Massachusetts Institute of Technology and Northwestern University, 2001.

Chen, Yehning, "Banking Panics: The Role of the First-Come, First-Served Rule and Informational Externalities," manuscript, University of California, Los Angeles, 1995.

Chowdhry, Bhagwan, and Vikram Nanda, "Leverage and Market Stability: The Role of Margin Rules and Price Limits," *Journal of Business,* LXXI (1998), 179–210.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, "Noise Trader Risk in Financial Markets," *Journal of Political Economy,* XCVIII (1990), 703–738.

Dennis, Patrick J., and Deon Strickland, "Who Blinks in Volatile Markets, Individuals or Institutions?" *Journal of Finance,* LVII (2002), 1923–1949.

Diamond, Douglas, and Philip Dybvig, "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy,* XCI (1983), 401–419.

Diamond, Douglas W., and Robert E. Verrecchia, "Disclosure, Liquidity, and the Cost of Capital," *The Journal of Finance,* XLVI (1991), 1325–1355.

Donaldson, R. Glenn, and Harald Uhlig, "The Impact of Large Portfolio Insurers on Asset Prices," *The Journal of Finance,* XLVIII (1993), 1943–1953.

Dow, James, and Gary Gorton, "Arbitrage Chains," *The Journal of Finance,* XLIX (1994), 819–849.

Fair, Ray, "Events That Shook the Market," *Journal of Business,* LXXV (2002), 713–732.

Gammill, James F., and Terry Marsh, "Trading Activity and Price Behavior in the Stock and Stock Index Futures Markets in October 1987," *Journal of Economic Perspectives,* II (1988), 25–44.

Garman, Mark, "Market Microstructure," *Journal of Financial Economics,* III (1976), 257–275.

Genotte, Gerard, and Hayne Leland, "Market Liquidity, Hedging, and Crashes," *American Economic Review,* LXXX (1990), 999–1021.

Goldstein, Michael A., and Kenneth A. Kavajecz, "Trading Strategies during Circuit Breakers and Extreme Market Movements," manuscript, Babson College and University of Pennsylvania, 2003.

Greenwald, Bruce, and Jeremy Stein, "The Task Force Report: The Reasoning Behind the Recommendations," *Journal of Economic Perspectives,* II (1988), 3–23.

Greenwald, Bruce, and Jeremy Stein, "Transactional Risk, Market Crashes, and the Role of Circuit Breakers," *Journal of Business,* LXIV (1991), 443–462.
Grossman, Sanford J., "An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Trading Strategies," *Journal of Business,* LXI (1988), 275–298.
Grossman, Sanford, and Merton Miller, "Liquidity and Market Structure," *The Journal of Finance,* XLIII (1988), 617–633.
Grossman, Sanford, and Zhogquan Zhou, "Equilibrium Analysis of Portfolio Insurance," *The Journal of Finance,* LI (1994), 1379–1403.
Hasbrouck, Joel, and George Sofianos, "The Trades of Market Makers: An Empirical Analysis of NYSE Specialists," *The Journal of Finance,* XLVIII (1993), 1565–1593.
Ho, Thomas, and Hans Stoll, "Optimal Dealer Pricing under Transactions and Return Uncertainty," *Journal of Financial Economics,* IX (1981), 47–73.
Jacklin, Charles J., Allan W. Kleidon, and Paul Pfleiderer, "Underestimation of Portfolio Insurance and the Crash of October 1987," *Review of Financial Studies,* V (1992), 35–63.
Leland, Hayne, and Mark Rubinstein, "Comments on the Market Crash: Six Months After," *Journal of Economic Perspectives,* II (1988), 45–50.
Liu, Jun, and Francis A. Longstaff, "Losing Money on Arbitrages: Optimal Dynamic Portfolio Choice in Markets with Arbitrage Opportunities," manuscript, University of California, Los Angeles, 2000.
Madhavan, Ananth, and Seymour Smidt, "An Analysis of Changes in Specialist Inventories and Quotations," *The Journal of Finance,* XLVIII (1993), 1595–1628.
Madrigal, Vicente, and José A. Scheinkman, "Price Crashes, Information Aggregation, and Market-Making," *Journal of Economic Theory,* LXXV (1997), 16–63.
Romer, David, "Rational Asset-Price Movements without News," *American Economic Review,* LXXXIII (1993), 1112–1130.
Sandroni, Alvaro, "Learning, Rare Events, and Recurrent Market Crashes in Frictionless Economies without Intrinsic Uncertainty," *Journal of Economic Theory,* LXXXII (1998), 1–18.
Schnabel, Isabel, and Hyun Song Shin, "Foreshadowing LTCM: The Crisis of 1763," manuscript, London School of Economics, 2002.
Schwert, G. William, "Stock Market Volatility: Ten Years After the Crash," *Brookings-Wharton Papers on Financial Services,* I (1998), 65–114.
Shleifer, Andrei, *Inefficient Markets—An Introduction to Behavioral Finance* (Oxford, UK: Oxford University Press, 2000).
Shleifer, Andrei, and Robert W. Vishny, "The Limits of Arbitrage," *The Journal of Finance,* LII (1997), 35–55.
Welch, Ivo, "Sequential Sales, Learning, and Cascades," *The Journal of Finance,* XLVII (1992), 695–732.
Yuan, Kathy, "Asymmetric Price Movements and Borrowing Constraints," manuscript, Massachusetts Institute of Technology, 2000.