JOURNAL OF FINANCIAL INTERMEDIATION **6,** 64–86 (1997)
ARTICLE NO. JF970209

Thin Markets, Asymmetric Information, and Mortgage-Backed Securities*

EDWARD L. GLAESER

Harvard University, Hoover Institution, and NBER

AND

HÉDI D. KALLAL

New York University and Salomon Brothers

Received July 21, 1993

This paper tries to explain why the issuers of an asset would restrict what information is available about their asset. In a world where knowledge is valued, market forces should induce disclosure, but we often see markets (such as the market for mortgage-backed securities) where assets' issuers refuse to release valuable information. We present a model of market liquidity and find that market liquidity can both rise and fall with the quantity of released information. More information may increase asymmetries of information and "lemons" style breakdowns. We find that asset bundling is more advantageous when private information is more accurate, which may be the case in the mortgage-backed securities market. *Journal of Economic Literature* Classification Numbers: G14, G32.

I. INTRODUCTION

Consider an asset that pays a fixed quantity in a year if and only if the Dow Jones industrial average lay between 5000 and 5200. Imagine a second asset that pays a fixed quantity at the same date if and only if the Dow

* We acknowledge financial support from the University of Chicago Department of Economics, the Caisse Autonome de Refinancement, Paris, and NSF Grant SES 8420930. Glaeser also thanks the Javits Fellowship Program, the Bradley Foundation, and NSF Grant SBR 9309808. Kallal thanks the Frank H. Knight Fellowship Fund. Helpful comments were provided by Joseph Beaulieu, In-Koo Cho, Bruno Hallak, Andrei Shleifer, Jean-Luc Vila, Laurence Weiss, and participants in the University of Chicago Theory Workshop. José Scheinkman provided us with the topic and guidance throughout the research. We are also grateful to a large number of financial professionals for the time and insights. Guillermo Babatz provided superb research assistance.

1042-9573/97 $25.00

Jones Industrial average lay between 5200 and 5400. Surely, it makes sense for an issuer to sell these assets separately. An investor who wants both assets can always buy both, and investors who only want one of the assets would be willing to pay to have the assets available separately.

Consider the same two assets and imagine that the issuer sells one of the assets and does not tell the purchaser which one of the assets he is selling. The purchaser only finds out what asset he owns at the end of the year. This failure to reveal information creates unnecessary risk that would be undesirable to a risk-averse investor; furthermore, it limits an investor's ability to use the assets to diversify a portfolio. Surely the issue would command a higher price if the issuer reveals what asset is being sold.

These arguments suggest that assets' issuers should neither bundle assets nor restrict available information. Yet in the mortgage-backed security (hereafter MBS) market, mortgages are sold on the secondary market in large bundles. MBS issuers often refuse to disclose even the most rudimentary information about the ingredients of a particular bundle. This fact presents a puzzle.

To approach this question, this paper joins the small literature on asset bundling (Gorton and Pennachi, 1989; Boot and Thakor, 1993; Subrahamanyam, 1992; Tashjian and Weissman, 1995)[1] and the somewhat larger literature on information disclosure (Diamond, 1985; Diamond and Verrechia, 1991; Kim and Verrechia, 1994; Fishman and Hagerty, 1995; Yosha, 1995). We explain asset bundling and limited information disclosure as methods of increasing liquidity in the mortgage-backed security market.[2] While many of our results on information disclosure are close to the existing literature (particularly to Diamond and Verrechia, 1991), our model and our findings are quite different from most previous work.

The central difference in our model and our results comes from a fundamental difference in our modeling approach. In most liquidity models (such as Subrahamanyam, 1992), market structure follows Kyle (1985) and there are uninformed liquidity traders (i.e., traders who need to trade) and informed market participants. These traders may choose not to trade for fear of suffering losses due to their relative ignorance. Limiting both asset variance and private information in that model allows the uniformed liquidity traders to trade without fear of informed traders and thus increases the liquidity of the market. Limiting private information, however, can also lower liquidity. For example, in Diamond and Verrechia (1991), limited

[1] Subrahamanyam (1992) and Gorton and Pennachi (1989) are both information-based models of bundling and are close to this paper. The other two papers offer quite different models. In Boot and Thakor (1992) bundling increases competition rather than working through the distribution of asset values. In Tashjian and Weissman (1995) bundling allows for a form of price discrimination.

[2] The connection between information and liquidity generally comes from Akerlof (1970).

private information decreases the returns to market makers, causing them to exit and thus decreasing liquidity.

In our model, the costs and benefits to limiting information asymmetries will be quite different from standard models. Our framework has three actors: an initial issuer, informed financial intermediaries who buy the asset directly from the issuer and an uninformed market. Following the security design literature (Allen and Gale, 1988; Chiesa, 1992; Boot and Thakor, 1993; DeMarzo and Duffie, 1995), asset attributes such as bundling and information disclosure will be chosen to maximize the price that the issuer can charge the informed financial intermediaries. The initial issuer is a well established entity that sells all of its assets (e.g., Fannie Mae) and we assume that there are no informational problems during the stage where the issuer sells to the intermediary. The informational problems occur when the intermediary attempts to resell the asset to the market.

The equilibrium level of liquidity in this secondary market is found to induce "accurate" pricing of the asset, where accurate means that the intermediary's asking price reflects the expected value of the asset to the market, based on the intermediary's private information. Informed intermediaries sell the asset for its expected value because an overpriced asset will be less likely to sell, or because they will only be able to sell the overpriced asset after a delay.

Initial issue price of an asset will be a function of that asset's expected liquidity, so the issuer will choose bundling and optimal information disclosure to maximize liquidity. Both of these choices determine liquidity, because the asymmetries of information between intermediaries and final buyers that determine illiquidity are a function of the variance of the intermediaries' valuation of the asset. Both asset bundling and limited information disclosure by the issuer will act to reduce the variance of the intermediaries valuation. Bundling smoothes out extreme values; limited information forces the sellers' estimates to be closer to his priors, and these estimates will have a lower variance.

The connection between asset variance and liquidity comes from two features of the liquidity function, i.e., the function that maps price into the probability of a quick sale (illiquidity is alternatively modeled as a lower probability of making a sale or as a delay before sale at the asked price) (see Fig. 1). First, the liquidity function is a decreasing function of the support of the sellers' valuation of the asset. When the support contracts, liquidity rises. This feature produces the result in our model that less information, and bundling, can increase liquidity if less information contracts the extreme values of the support of the asset. Our result, that less information can increase liquidity, closely parallels the results of the information disclosure literature (Diamond and Verrechia, 1991) and the bundling literature (Subrahamanyam, 1992).

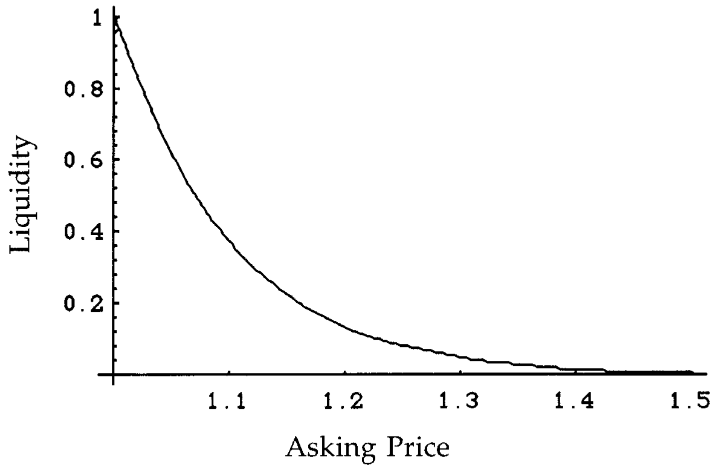

FIG. 1. The liquidity function generated when the possible value of the asset to the seller ranges from 1 upward and the cost to the seller of not selling the asset is 0.1.

Convexity of the liquidity function is the second feature of our model. This convexity comes from the fact that liquidity is found to induce accurate pricing. Accurate pricing requires that a marginal overpricing of the asset leads to a reduction in the probability of sale that is just large enough to dissuade the seller from mispricing the asset. As the asset value rises and liquidity falls, the change in liquidity needed to stop mispricing also falls. This convexity implies that a mean preserving contraction in the distribution of asset valuations (coming from either bundling or limited information disclosure) will actually decrease liquidity, if the support of the asset valuations is fixed and thus the liquidity function is constant. As opposed to previous models (such as Diamond and Verrechia, 1991), information acquisition (that holds the support of asset valuations constant) will increase liquidity without any effect on the number of market makers.

Thus, there are two competing effects of bundling or information disclosure on liquidity. One effect occurs because these actions contract the support of the asset's valuations to the intermediary, and this effect causes liquidity to rise. A second effect occurs because these actions create a mean preserving contraction in the distribution, and this effect causes liquidity to fall. These tensions exist in two different versions of our model. A general tradeoff seems to exist.

Bundling will increase liquidity, especially when there is already a great deal of private information and when bundling actually changes the support of the asset. These conditions seem roughly to be met in the mortgage-backed security market. There are many other markets, however, where bundling and limited information may reduce liquidity. In particular, older mortgage pools are in a sense less bundled than newer mortgages; they are traded individually not in the To Be Announced (TBA) market. Our model predicts this since more information is available about these older mortgages.

II. THE PUZZLE OF MORTGAGE-BACKED SECURITIES

Mortgage-backed securities are ''pass through'' securities where the initial lender still services the mortgages; the lender collects fees and then passes through interest and principal to the owners of the bundled mortgages. The Government National Mortgage Association (GNMA or ''Ginnie Mae''), the Federal Home Loan Mortgage Corporation (FHLMC or ''Freddie Mac''), and the Federal National Mortgage Association (FNMA or ''Fannie Mae'') are the three main issuers of mortgage-backed securities. These agencies purchase mortgages from initial lenders (banks, savings, and loans). The initial lenders have their own asymmetric information problem with borrowers, primarily concerning default. However, these problems have been modeled (Dunn and Spatt, 1988; Chari and Jagannathan, 1989) and are not all that relevant to the secondary mortgage market, where mortgages are generally insured against default by agencies such as the Federal Housing Administration, the Veterans' Administration, and other agencies including private insurers.

The issuing agencies (FNMA, GNMA, FHLMC) guarantee mortgages and then bundle mortgages from different regions and risk groups. Pools can be quite small; the minimum pool size for GNMA II mortgage bundles is $250,000, but pools can also be much larger (often over $10 million). The demand for bundling is such that GNMA has a program allowing for smaller pools to be combined into larger bundles. Many aspects of the mortgages are hidden and bundles are made up independently of specifics, such as location, size of mortgage (up to certain maximum loan sizes), and borrower characteristics. However, all mortgages within a bundle share a common duration (i.e., 30 year mortgages are not bundled with 15 year mortgages).

Mortgage bundles are first sold to primary dealers such as Salomon Brothers and Bear Sterns. These primary dealers often hold some quantity of MBSs for their own portfolios, but they primarily resell these bonds to final investors, such as pension funds or mutual funds. The final demand for these securities lies within institutions that have substantial long term liabilities and who desire to have long term assets to balance these liabilities. The tendency of the primary dealers to hold on to these securities is often explained as an attempt to use their superior understanding of these instruments to earn capital. These primary dealers often repackage the mortgages into instruments such as CMOs (Collateralized Mortgage Obligations).

CMOs are an attempt to finetune MBSs so that they provide a very particular payment stream.

The important, idiosyncratic risk of mortgage-backed securities comes from the possibility of prepayment.[3] At any time, mortgages can be repaid at par. Prepayment comes when homeowners (i) default, (ii) move, or (iii) gain financially by repaying at par. As the borrower only gains financially by prepaying when the MBS is trading for more than par, the borrower's financial gain is always the bondholder's financial loss. When prepayment occurs because of the sale of the house or default, then prepayment might hurt the bondholder (if the mortgage is trading for more than par). Prepayment risk is substantial and the prepayment probability differs across mortgages. Different geographic areas have different prepayment probability. Mortgages with different maturities and different coupons also have different probabilities of prepayment.

Since it is possible to use the information about mortgages to predict prepayment, firms diligently try to predict prepayment likelihoods. Indeed, all traders desire more information about securities and occasionally are willing to go to absurd lengths to get that information. Michael Lewis, in *Liar's Poker*, describes Salomon Brothers sending young analysts out to suburban malls to try to gauge prepayment probabilities. As of 1989, Bear Sterns seems to have acquired a reputation for knowing about particular mortgage pools. Right now most of this knowledge centers on (1) the location of the issuers of the mortgages in a particular pool, (2) the past history of that pool, and (3) regional prepayment characteristics. While intermediaries occasionally are successful at predicting prepayment, anecdotally the predictive power is said to be fairly weak.

Indeed, Freddie Mac, Fannie Mae, and Ginnie Mae give out only the most limited information about the mortgages in a pool. Freddie and Fannie reveal the location of the mortgage's issuer. Since location of issuer does not necessarily reveal the location of the mortgages (i.e., Citibank's official state of residence is Michigan), this information may not be particularly relevant. Agencies could certainly reveal more than they do about the location, maturity, and coupon of mortgages in pools.

The mortgage-backed security market is vast. Bartlett (1989) reports that over $708 billion worth of mortgages have been securitized—this represents more than one-third of the more than $2 trillion home mortgage market. The size of this market makes it obvious that transaction costs cannot explain much of the bundling or the limited information. MBSs are traded over the counter, and it would be cheap (relative to the size of the markets)

[3] Like most other securities, the value of mortgage-backed securities depends on the overall interest rate environment.

to have separate markets for pools containing only New York mortgages and pools containing only Houston mortgages.

There exist two distinct types of markets for MBS securities. Newer mortgage backed securities are traded in the TBA market. Older mortgage-backed securities are almost all traded in the market for specific pools not in the TBA market. This occurs, perhaps, because there is more information available on these older MBS pools. In that market, traders exchange obligations to provide mortgage-backed securities at a later date. Traders accept that they will be provided with the cheapest possible mortgage-backed security at that later date. In a sense, this step involves a second bundling, because in the TBA market, traders buy and sell an undisclosed bundle of many possible bundles. This suggests a further desire of the market to bundle these securities. Later, we will discuss why this market structure may be optimal in the context of our model.

The mortgage-backed security market has certain features that models should consider. First, there are two steps of sale. The first step involves a sale from a long-term actor who can essentially commit itself because of its reputation and because of its highly regulated status not to take advantage of its private information (indeed, no one has ever suggested a lemons problem in the initial sale of mortgages to the primary dealer). The second step involves a sale from the primary dealers to the final purchasers. There are large gains to these sales going through quickly, but adverse selection creates limits to liquidity. The primary dealers who have spent more time analyzing prepayment risk can probably value these assets more accurately than the portfolio managers who will eventually hold them. Second, there exists an uneven distribution of information. Analyzing prepayment risk is difficult and very mortgage specific. Third, there seem to be several points where players in the market try to further bundle these securities. Finally, the issuers choose not to issue all the information. The following model will try to explain the preference for bundling and the limits of information acquisition as an outcome of the potential and existing levels of asymmetric information.

III. THE BASIC FRAMEWORK

This model examines asset resale in a world of endogenous information acquisition. At time zero, an asset is sold by an issuer (e.g., Fannie Mae, Freddie Mac) who posts a price for the asset and is able to extract the asset's expected value. There is assumed to be no adverse selection problems in this initial sale.[4] We denote this initial issue price, which will be equal to

[4] This assumption can be justified if the issuer sells all his assets or if there is an underlying reputation model.

the assets' expected return, as P_0. The issuer will be able to change the features of the asset. We will try to understand why particular features of the asset (such as bundling or restricted information) might raise P_0 and therefore be desirable to the issuer.

After the asset is offered, one unit of the asset is bought by an intermediary (e.g., an investment bank). The intermediary holds the asset and may receive a noisy signal of the asset's true value, between time zero and time one. At time one, the intermediary offers to sell the asset to the general market for a price $\tilde{P}$. If the intermediary holds the asset in its own portfolio, the ultimate value to the intermediary will be a random variable V. If the intermediary sells the asset to final consumers (such as mutual funds), the ultimate value to those final consumers will be $V + \varepsilon$, where ε is a nonstochastic gain from trade.[5] The intermediaries' knowledge at time one is denoted by I, and the expected value of $V + \varepsilon$, based on I, will be denoted P.

The market will always have less information about the asset than the intermediary. This assumption seems quite realistic in the MBS market where intermediaries have developed a large quantity of specialized information about the prepayment risk of mortgages. The assumption is particularly close to Kim and Verrechia (1994).

To recapitulate, the model has five phases: (1) at time zero the issuer sells the asset to the intermediary for a price P_0, (2) between time zero and time one, with a certain probability, the intermediary acquires a possibly noisy signal of the bond's value (in the first model, this probability will be the result of search by the intermediary; in the second model, this probability will be fixed), (3) at time one, the intermediary offers the asset for sale to the market at a price $\tilde{P}$, (4) after time one, the market may purchase some quantity of the asset, and (5) eventually, the value of the asset is revealed to everyone, and the intermediary can sell off the remainder of the asset.

While there are multiple equilibria in our models, we will restrict ourselves to those equilibria where there is truth telling—i.e., where the value charged by the intermediary equals the intermediary's expectation of the value of the asset as of time one. In equilibrium, this behavior will be privately optimal.[6] Truth-telling implies that $\tilde{P} = P$.

We now proceed with two different models within this framework. In

[5] Some benefit of selling the asset is necessary to avoid no trade results. This cost might be stochastic and represent a liquidity shock to the intermediary. It might be a result of greater diversification-based demand for this asset among particular investors.

[6] We believe that truth-telling has a variety of beneficial features, in particular being stable (in a Kohlberg and Mertens (1986) sense—this game has a standard signaling structure and the intuition of the stability follows arguments in Cho and Kreps (1987)). A fully separating equilibrium where agents state the value of their assets is also more intuitive and appealing than any one of the possible pooling equilibria. Alternatively, we can appeal to the Revelation Principle; see, for example, Myerson (1979).

the first model, there is a fixed market ready to purchase the good, and illiquidity will mean that only a share of the goods offered will be purchased by the market. In the second model, there will be a flow of potential buyers and illiquidity will mean that it takes longer to sell the asset. The first model is simpler and is easier to analyze. We provide the second model to show that these results go through in a more realistic depiction of the mortgage-backed securities market.

IV. NOISY SIGNALS AND BINARY VARIABLES

In the first model, the underlying value of the asset, V, is either $\beta + \alpha$ or $\beta - \alpha$, each with unconditional probability 0.05. Between time zero and time one, a proportion of the population receives a signal that noisily reveals the asset's true worth. As in Subrahamanyam (1992), this proportion of informed intermediaries will be endogenous. The costs of information acquisition are individual specific and denoted by c, which has a cumulative distribution function $G(c, I)$ and corresponding density $g(c, I)$. The term I reflects the amount of information that the seller chooses to divulge, so that $\partial G(c, I)/\partial I > 0$. Intermediaries purchase the asset before they observe their level of c. We will denote marginal information acquirer with c^*, and thus the number of informed intermediaries will always be $G(c^*, I)$. The distribution of c is assumed to be such that in equilibrium there are always some intermediaries who acquire information and some who choose not to acquire that information.

The signal is correct σ percent of the time and 1-σ percent of the time the signal gives no information (thus, when the signal is valueless the expected value of the asset is β), and the intermediary does not know when the signal is valueless. At time one, an intermediary may fall into three classes: (1) having no signal, so the expected value of the asset is still β, (2) having a positive noisy signal, so the expected value of the asset is $\beta + \sigma\alpha$, and (3) having a negative noisy signal, so the asset is expected to be worth $\beta - \sigma\alpha$.

The intermediary offers the asset at a price to the market, and some portion of the market will purchase the security from the intermediary. The amount of the asset bought by the market at the asked price will be denoted $\lambda(\tilde{P})$; this quantity will be interpreted as market liquidity. The function, $\lambda(\tilde{P})$, which is the ability to sell the good at its value quickly to the public, corresponds with intuitive notions of liquidity such as Keynes' (1930) idea that an asset is liquid if "it is certainly realizable at short notice without loss."

In equilibrium, market liquidity will ensure that the intermediary offers the asset at the expected value to the market. To induce this "truth-telling" behavior, liquidity at high prices must be just low enough so that intermediaries are not tempted to offer the asset at a price above the asset's expected value. Alternatively, liquidity at high prices must be high enough so that intermediaries with a high value asset will not choose to underprice the asset. As a general rule, there must always be perfect liquidity of assets sold at the lowest expected value that the asset can possibly have, denoted $\underline{P}$. Since it would be impossible for an asset to be worth less than $\beta - \sigma\alpha + \varepsilon$ to the market (which is $\underline{P}$ in this case), the market must offer complete liquidity to assets sold for $\beta - \sigma\alpha + \varepsilon$.[7] For intermediaries with low signals to prefer truth telling instead of charging β it must be true that

$$\beta - \sigma\alpha + \varepsilon \geq \lambda(\beta + \varepsilon)(\beta + \varepsilon) + (1 - \lambda(\beta + \varepsilon))(\beta - \sigma\alpha). \tag{1}$$

The benefit of mispricing is the higher price if the asset sells; the cost of mispricing is the lost gains from trade (ε) if the asset fails to sell. Stability arguments imply that (1) must hold with equality (as in Cho and Kreps, 1987, the Pareto optimal separating equilibrium is the stable equilibrium), which implies $\lambda(\beta + \varepsilon) = \varepsilon/(\sigma\alpha + \varepsilon)$.

For agents with no signal to prefer truth-telling to charging $\beta + \sigma\alpha + \varepsilon$ it must be true that

$$\begin{aligned} &\lambda(\beta + \varepsilon)(\beta + \varepsilon) + (1 - \lambda(\beta + \varepsilon))\beta \\ &\quad \geq \lambda(\beta + \sigma\alpha + \varepsilon)(\beta + \sigma\alpha + \varepsilon) \\ &\quad + (1 - \lambda(\beta + \sigma\alpha + \varepsilon))\beta. \end{aligned} \tag{2}$$

Again stability arguments tell us that (2) should hold with equality, which implies that $\lambda(\beta + \sigma\alpha + \varepsilon) = \varepsilon^2/(\sigma\alpha + \varepsilon)^2$.[8] With this we have solved for the values of $\lambda(\tilde{P})$ for all the relevant prices; therefore, we can solve for the equilibrium level of information acquisition, which is determined so that the marginal intermediary (with costs c^*) is exactly indifferent between acquiring information and not acquiring information,

$$\beta + 0.5(\varepsilon + \varepsilon\lambda(\beta + \sigma\alpha + \varepsilon)) - c^* = \beta + \varepsilon\lambda(\beta + \sigma\alpha + \varepsilon) \tag{3}$$

or $c^* = \alpha\sigma\varepsilon(\alpha\sigma + 2\varepsilon)/2(\alpha\sigma + \varepsilon)^2$. The number of informed agents (i.e.,

[7] If it didn't, the intermediary would offer to sell the asset for $\beta - \sigma\alpha + \varepsilon - \xi$, for some arbitrarily small ξ. Since the market would be earning positive profits at no risk at such a sale price, they would, of course, buy all of such an offer. The intermediary would then deviate and charge $\beta - \sigma\alpha + \varepsilon - \xi$ if the liquidity of assets sold at $\beta - \sigma\alpha + \varepsilon$ was strictly less than one.

[8] If (1) and (2) hold with equality, it follows that agents with high signals will not choose to imitate any other agents, and agents with no signals will not choose to imitate agents with low signals. Agents with low signals will not choose to imitate agents with high signals.

c^*) rises with α, σ, and ε. The equilibrium offering price of the asset at period zero is set by the issuer so that the profits of the intermediary are zero:

$$P_0 = \beta + \varepsilon \left[\frac{\varepsilon^2 + \sigma\alpha\varepsilon + 0.5\,\sigma^2\alpha^2 G(c^*, I)}{(\varepsilon + \sigma\alpha)^2} \right] - \int_{c \leq c^*} cg(c, I)\, dc. \quad (4)$$

The asset price equals the expected value of the asset to the intermediary (β), plus the expected gains from selling the asset to the final investors (ε) times the probability of resale, or the amount of liquidity, i.e.,

$$\left(\frac{\varepsilon^2 + \sigma\alpha\varepsilon + 0.5\,\sigma^2\alpha^2 G(c^*, I)}{(\varepsilon + \sigma\alpha)^2} \right),$$

minus the expected search costs search

$$\left(\int_{c \leq c^*} cg(c, I)\, dc \right).$$

One natural representation of bundling is to imagine that this single asset is now replaced by an average of two identical but independently distributed assets. In the first case, we assume that when an agent searches, the agent learns the value of both assets in the bundle. When the assets are sold together and the intermediary acquires information there are three possible cases: (1) the intermediary learns that one asset is worth $\beta + \sigma\alpha$ and one asset is worth $\beta - \sigma\alpha$ so the combined asset is worth β, (2) the intermediary learns that both assets are worth $\beta + \sigma\alpha$, and (3) the intermediary learns that both assets are worth $\beta - \sigma\alpha$.

If information acquisition yields signals on each asset, then this bundling reduces the probability that the asset will take on each of its extreme values from one-half to one-quarter, and one-half of the time the asset will be worth its average value.[9] We will capture this type of bundling by denoting the probability that the asset is worth each of its extreme values as θ, and its value of being worth its average value as $1 - 2\theta$ (thus when there is one security in the bundle $\theta = 0.5$ and when there are two securities in the bundle $\theta = 0.25$). We can then investigate comparative statics on θ to examine the costs and benefits of bundling. If we include the parameter θ, then (4) becomes

[9] Of course, this will not be true for bundling more than two securities together. More generally, bundling of this sort will represent a mean preserving contraction in the distribution of the asset's values that leaves the extreme points of the value distribution untouched.

$$P_0 = \beta + \varepsilon \frac{\varepsilon^2 + \sigma\alpha\varepsilon + \theta\sigma^2\alpha^2 G(c^*, I)}{(\varepsilon + \sigma\alpha)^2} - \int_{c \le c^*} cg(c, I)\, dc. \qquad (4')$$

Alternatively, we can assume that when the intermediary acquired information about the bundle of two assets, it could only learn about one of the component assets in the bundle. This assumption represents an alternative view of bundling. Then the asset's value to the intermediary (conditional upon search) will equal $\beta + \sigma\alpha/2$ or $\beta - \sigma\alpha/2$, each with probability of one-half. Since bundling in this case is equivalent to a reduction in α, we can examine the value of bundling by examining comparative statics of α.

Comparative Statics

Differentiation tells us that both liquidity and the issue price rise with I (the availability of information). This effect occurs because search increases the probability that the asset will be worth its extreme values, relative to its average values. The average liquidity of the extreme values is higher than the liquidity of the average value, so liquidity rises with search. Therefore, actions that will make it easier for individuals to search will be desired by the initial issuer in this model. This comparative static is completely opposite to the findings of Subrahamanyam (1992) and many other authors in this literature. Obviously, the liquidity problems would vanish if there was no asymmetry of information, but as long as there is some possibility for private information, liquidity is increased with the level of information gathering. There is, thus, a nonmonotonic relationship between information and value; when there is some asymmetric information, it pays to make more available.

All other comparative statics generally have two effects. The first effect is the direct effect of changes in the parameter on price and liquidity, holding c^* constant. The second or indirect effect works through the changes in c^* caused by changes in the parameter. Unlike in Subrahamanyam, primarily because information gathering is liquidity enhancing in this model, the direct effect of parameters on liquidity will generally go against the indirect effect working through less information gathering. A low enough value of $g(c^*, I)$ generally guarantees that the direct effect on liquidity is more than offset by the indirect effect. This occurs because a low value of $g(c^*, I)$ implies that changes in the incentives to search will have little effect on the overall amount of searching. It is worth stressing that the indirect effects of parameter changes through information gathering will effect liquidity but not price, because lower liquidity through less search will always be exactly offset by lower search costs.

The asset price falls with α and σ. Liquidity also generally falls with these variables if $g(c^*, I)$ is small. These comparative statics imply that if intermediaries only learn the value of one piece of a bundle, then bundling will increase asset value and increase liquidity. Because increases in these parameters lead to reductions in c^*, it is possible for liquidity to rise as α and σ rise. The overall price, however, must fall. Since these parameters determine the extreme values of the asset to the intermediary, they will determine the level of liquidity. Moreover, when the extreme values become more extreme, liquidity falls. It is also easy to show that, as long as $g(c^*, I)$ is small, reductions in α are more advantageous when σ is large. Thus when signals are accurate and there is a large amount of asymmetric information, bundling is more advantageous.

Increases in θ must lead to an increase in price and will lead to an increase in liquidity again as long as $g(c^*, I)$ is small. When bundling does not affect the extreme values of the asset, but rather only makes it less likely that the asset will take on its extreme values, then bundling reduces both liquidity (generally) and price.

Overall, this model suggests something of a puzzle. When the asset is more likely to be worth its extreme values (i.e., when I or θ rise), then average liquidity rises, since $\lambda(\tilde{P})$ is convex. This comparative static is also found by Boot and Thakor (1993), but in their work, increased liquidity is connected to the number of agents who choose to become informed traders. In this paper, liquidity rises because it becomes easier to enforce truth-telling.

However, when the extreme values of the asset become more extreme (i.e., when either σ or α increase) then liquidity and initial price fall. Thus bundling enhances value and liquidity if it reduces the extreme values of the asset, but bundling is harmful if it reduces the likelihood that the asset will be worth its extreme values. This point is, to our knowledge, a new and somewhat counterintuitive finding. The convexity of the liquidity function lies behind this result. The convexity itself results because it is more difficult to dissuade intermediaries from small amounts of misrepresentation than from large amounts of misrepresentation in our model.

The amount of liquidity rises with ε (the holding cost of the asset). The intuition, again, is that truth-telling becomes relatively more appealing as ε rises because the costs of having one's offer rejected become higher. However, the initial price of the asset falls with ε despite the higher liquidity as higher values of ε lead to higher losses from whatever illiquidity exists.

Understanding the Convexity of the Liquidity Function

The two competing effects of bundling are a result of the convexity of $\lambda(\ )$. This convexity is fairly general within the class of models related to the previous model (it shows up in the subsequent model), and it seems to be a fairly pervasive feature. For example, consider the general situation in which P can take on many values (a continuum). Failure to sell generates a loss $g(P)$, which is ε in the previous model. The basic maximization problem for the intermediary will be to choose sale price, $\tilde{P}$, to maximize $\lambda(\tilde{P})\tilde{P} + (1 - \lambda(\tilde{P})(P - g(P))$. Assuming that in equilibrium the liquidity function will be twice differentiable, the first-order condition is $\lambda(\tilde{P}) + \lambda'(\tilde{P})(\tilde{P} - P + g(P)) = 0$, which in a truth-telling equilibrium implies that $\lambda(\tilde{P}) = -\lambda'(\tilde{P})g(P)$. Solving this differential equation, using K to denote the constant of integration, the liquidity function is $\lambda(P) = Ke^{-\int dP/g(P)}$, which is convex as long as $g'(P) > -1$. This is a fairly weak condition (we suspect that the costs from failing to trade should be constant with price, as in the above model, or perhaps increasing with P). In this framework, when the liquidity function is well behaved (i.e., twice differentiable), the liquidity function will also be convex under most reasonable specifications of the loss function.

The convexity can be best understood from examining the first-order condition $\lambda(\tilde{P}) = -\lambda'(\tilde{P})g(P)$. This condition implies that the gain from increasing the price (which is the level of liquidity at P) must be equated to the cost of higher prices (which is the reduction in liquidity) times the loss from sale. At high levels of liquidity, or low prices, the reduction in liquidity must be high. At low levels of liquidity, or high prices, the reduction in liquidity will be much lower. As the needed decline in liquidity falls with price, the function will be convex (unless $g(P)$ is falling so sharply with P to offset this effect).

V. ENDOGENOUS LIQUIDITY

In this form of the model, we assume, again, that the asset's true value is distributed over the real line. Now all sellers receive some signal about the asset's value. Intermediaries all receive a common set of information I and $P = E(V|I)$. We denote $\underline{P}(I)$ as the lowest possible expected value of P that can be realized given a fixed quantity of information. Since an increase in information should increase the variance of the signal P, we naturally assume that the lower bound of the distribution of P falls with I, e.g., $\underline{P}'(I) \leq 0$.

The structure of this model changes somewhat from the previous section. Again, at time one, the seller posts a price, $\tilde{P}$, for his asset. After this point in time, a continuous flow of buyers q appears ready to purchase the good. This measure q of buyers is one type of liquidity, and while it will be treated initially as an exogenous parameter, this measure will eventually be endogenized. The seller discounts the future at a continuous discount rate of r and prefers selling earlier to selling later; this discounting creates the costs of failing to sell.

Over time there is also a probability that the sellers' information will be revealed publicly. If the information becomes common knowledge, the price will adjust immediately to this new price P. The probability that this signal is revealed follows a standard Poisson distribution, where the marginal probability of information revelation is $\phi e^{-\phi t}$. If no new information occurs, the price is assumed to stay fixed at $\tilde{P}$; the seller charges a fixed price.

Again, we are interested in the truth-telling equilibrium. The waiting time before purchase is the mechanism which induces truth-telling. High posted prices bring about long delays; low posted prices create short delays. In particular, there will be an equilibrium delay function $D(P)$ which indicates the length of time before buyers start purchasing the asset. In the case that the information is publicly revealed, purchasing starts immediately. It must also be true that $D(\underline{P}(I)) = 0$, because an agent who charges the lowest possible price cannot possibly be lying; a price marginally lower would yield positive profits and the buyers would immediately purchase.

Once purchasing starts, the agents sell q of their good per period. The first quantity worth determining is the profits once purchasing begins. This quantity is found by considering both the sales at the asking price $\tilde{P}$ and the probability that the information will be revealed. Given these conditions, the asset will be sold at price P (the true price). In particular, the revenues are

$$\int_{t=0}^{1/q} e^{-rt}[e^{-\phi t}\overline{P} + (1 - e^{-\phi t})P]q\, dt = q(\overline{P} - P)\frac{1 - e^{-r+\phi/q}}{r + \phi} + qP\frac{1 - e^{-r/q}}{r}. \quad (5)$$

The integral represents the discounted sum of the sales profits. At any point of time in the future, there is a probability that information has been revealed $(1 - e^{-\phi t})$ and a probability that information has not been revealed $(e^{-\phi t})$. The price depends on whether information has been revealed or not. In fact, under truth-telling these two prices will be equal. Initially, however, we must determine the profits conditional on any price offer. To determine the actual profits of the seller, we must consider discounting to this point. We must also consider the possibility that information is distributed before this time period. If that information is revealed, the total profits will then be $qP(1 - e^{-r/q})/r$.

Total profits, denoted Π, conditional on an offer price $\tilde{P}$ are thus

$$\Pi \equiv e^{-(r+\phi)D(\tilde{P})}\left[q(\tilde{P} - P)\frac{1 - e^{-r+\phi/q}}{r + \phi} + qP\frac{1 - e^{-r/q}}{r}\right] + \int_{t=0}^{D(\tilde{P})} \phi e^{-(r+\phi)t} qP\frac{1 - e^{-r/q}}{r}\, dt. \quad (6)$$

This expression gives profits as a function of the true value of the asset and as a function of its stated price. Maximizing this expression with respect to $\tilde{P}$ gives the marginal product of changing the offer price:

$$\frac{d\Pi}{d\tilde{P}} = qe^{-(r+\phi)D(\tilde{P})}\left[\frac{1-e^{(r+\phi)/q}}{r+\phi}\right] + D'(\tilde{P})qe^{-(r+\phi)D(\tilde{P})}$$
$$\left[-\tilde{P}(1-e^{-(r+\phi)/q}) + P(e^{-r/q}-e^{-(r+\phi)/q})\right]. \tag{7}$$

For truth-telling to be optimal, this term must equal zero when $P = \tilde{P}$ (second-order conditions hold). This first-order condition implies

$$\frac{1-e^{-(r+\phi)/q}}{(r+\phi)(1-e^{-r/q})P} = D'(P). \tag{8}$$

Integration, and using the fact that $D(\underline{P}(I))$ must equal zero, tells us that the value of the $D(P)$ function is

$$D(P) = \frac{1-e^{-(r+\phi)/q}}{(r+\phi)(1-e^{-r/q})}\operatorname{Log}\left(\frac{P}{P}\right). \tag{9}$$

Since $\underline{P}'(I) < 0$, it must be true that the delay for any given price rises with the overall quantity of information. It is also true that delay is concave with respect to price, or inversely if we consider liquidity to be the opposite of delay, liquidity is again convex with respect to price. The average delay time (which determines the initial asking price) will be

$$\int_{P=\underline{P}}^{\infty} \frac{1-e^{-(r+\phi)/q}}{(r+\phi)(1-e^{-r/q})}\operatorname{Log}\left(\frac{P}{\underline{P}}\right) f(P,I)\,dP. \tag{10}$$

The derivative of (10) with respect to I is

$$\int_{P=\underline{P}}^{\infty} \frac{1-e^{-(r+\phi)q}}{(r+\phi)(1-e^{-r/q})}\operatorname{Log}\left(\frac{P}{\underline{P}}\right)\frac{\partial f(P,I)}{\partial I}\,dP$$
$$-\int_{P=\underline{P}}^{\infty} \frac{1-e^{-(r+\phi)/q}}{(r+\phi)(1-e^{-r/q})}\frac{\underline{P}'(I)}{\underline{P}} f(P,I)\,dP. \tag{11}$$

The first term is negative. Because delay is concave in price, average delay will fall with information by Jensen's inequality. The second term is positive since $\underline{P}'(I) \leq 0$. The overall value of (11) is ambiguous, just as in the previous model. More information raises the level of delay at any given asking price, but it also shifts the distribution of asking prices in a way that will act to lower the average delay. Likewise, if bundling contracts the extreme values (i.e., raises $\underline{P}(I)$), bundling will increase liquidity. However, if bundling contracts the distribution of P (holding $\underline{P}$ fixed), then bundling will reduce liquidity and asset value.

In general, when the first term is small (perhaps because ϕ is large), then the second term will dominate. When there is little intrinsic liquidity, adding more information hurts the situation. When there is much liquidity, adding information helps the situation. Therefore, perhaps the tendency to bundle and obscure information in the MBS market is optimal particularly because there are natural tendencies toward illiquidity in that market.

Endogenizing Liquidity or q

The previous section treated the flow of buyers as exogenous. In fact, the number of buyers as well as the length of time delay both effect market liquidity. The time delay is the amount of apprehension that the market has before leaping in and buying a highly priced offer, but q determines how quickly the market can act once it has decided to act or how quickly the market can purchase a good once the information about that good is common knowledge.

We assume a supply of financial intermediaries who face a loss function if they enter this market but do not trade. This assumes that buyers focus on assets where they are likely to make a purchase. If buyers are unlikely to purchase an asset, conditional upon investigation, then these buyers will avoid the asset. We assume that potential buyers do not observe price, nor time since posting, before they actually appear ready to buy the asset.

The expected time period during which an asset will be available to be bought is

$$e^{-\phi D(P)}D(P) + \int_{t=0}^{D(P)} \phi t e^{-\phi t}\, dt + \frac{1}{q} = \frac{1}{q} + \frac{1 - e^{-\phi D(P)}}{\phi}. \tag{12}$$

During this period, only $1/q$ of the time will a buyer who investigated this available asset end up purchasing this asset. Therefore the fraction of the time that an investigation will end up in a purchase is

$$\int_{P=\underline{P}}^{\infty} \frac{\phi}{\phi + q(1 - e^{-\phi D(P)})} f(P, I)\, dP$$
$$= \int_{P=P}^{\infty} \frac{1}{1 + q/\phi\, [1 - (P/\underline{P})^{-\phi((1-e(r+\phi)/q)/(r+\phi)(1-e-r/q))}]} f(P, I)\, dP. \tag{13}$$

The supply of buyers is an increasing function $S(\quad)$ of this quantity (i.e., more suppliers show up if they are more likely to make a purchase), thus equilibria will be fixed points of the function

$$q = S\left(\int_{P=\underline{P}}^{\infty} \frac{\phi}{\phi + q(1 - e^{-\phi D(P)})} f(P, I)\, dP \right). \tag{14}$$

As long as the function inside the integral is declining in q, then this function will behave nicely and have no more than one fixed point. We are interested in comparative statics around a stable equilibrium (i.e., an equilibrium where $S(\quad)$ strikes the 45° line from above). Monotonic comparative statics occur when

$$\frac{q}{\phi}\left[1 - \left(\frac{P}{\underline{P}}\right)^{-\phi((1-e^{(r+\phi)/q})/(r+\phi)(1-e^{-r/q}))}\right] \tag{15}$$

moves monotonically with a given parameter.[10] The number of investors, q, will increase when this term declines.

This term is easily seen to be falling in $\underline{P}$[11] and therefore q is rising with $\underline{P}$. Liquidity rises with $\underline{P}$ because as the range of possible values shrinks, the necessary time delay shrinks, and more buyers are attracted to the security. Overall, this result suggests that market demand and attention to the asset will rise as the amount of information declines and as the variance of the asset shrinks.

[10] We are interested in fixed points which are "stable" in the sense that $S(\quad)$ strikes the 45° line from above, which we assume (for the purposes of this discussion) exist. It is fairly easy to assume sufficient structure on $S(\quad)$ to guarantee us at least one fixed point that satisfies this condition (i.e., we can assume that $S(0) > 0$ and $S(1) < 1$ and that $S(\quad)$ is continuous). It is easy to show that when all the P's are sufficiently close to $\underline{P}$, this condition will be met.

[11] To infer comparative statics, we are holding $f(P, I)$ fixed. This experiment assumes that there is negligible weight on the lower reaches of the distribution of P's and this distribution, with negligible weight, has raised its lower bound.

However, it is also true that

$$\frac{\phi}{\phi + q(1 - e^{-\phi D(P)})},$$

the probability of there being a sale, is convex in P. An increase in the level of information which creates a mean preserving spread in P (holding $\underline{P}$ constant) should, therefore, cause

$$\int_{P=\underline{P}}^{\infty} \frac{\phi}{\phi + q(1 - e^{-\phi D(P)})} f(P, I)\, dP$$

to increase, and q then must also increase by Jensen's inequality. The intuition is again that increased information leads to greater variance in P. This, in turn, means holding the liquidity function constant, which will lead to a greater likelihood of making a sale and a greater value of q (since supply is higher when, in equilibrium, sales are more frequent).

Thus, increased information has two familiar effects on the supply of liquidity. With more information, the given amount of delay for a given P increases. Also, with more information, when we hold delay for a given P constant, the prices that are charged will be more likely to be unusually high or low. This larger spread leads to, on average, more sales. The overall effect is ambiguous, but again if P values are generally close to $\underline{P}$, then the first effect will dominate. There will, therefore, be an increase in q as the amount of information falls.

Additional comparative states are that q is rising in r. As r rises, it becomes relatively easier to enforce truth-telling by the intermediary so the delay time shrinks and buyers find the security more appealing. Again this is the intuition that sellers' who find delay really costly are very trustworthy and less illiquidity is needed to guarantee truth-telling with these intermediaries. As long as P is relatively close to $\underline{P}$, the overall quantity of liquidity must be rising with ϕ. As the amount of information available to the flow of buyers rises, the flow of buyers will also rise, and the market will become more liquid.

VI. IMPLICATIONS FOR ASSET DESIGN AND CONNECTION TO THE PREVIOUS LITERATURE

This section uses the model to present some implications for an asset's issuer who is looking to maximize initial price.

Summary of Comparative Statics

Our primary focus has been on the disclosure of information and bundling. We find that both of these actions work in the same way, by altering the variance of the sellers' valuations of the asset. However, decreased variance of sellers' valuations can either increase or decrease the liquidity of the asset. Generally, when the asset is intrinsically and relatively illiquid and there is a tendency toward asymmetries of information, then bundling and reduced information disclosure are more likely to be optimal. Moreover, if we know that reduced information or bundling significantly contracts the support of the sellers' valuations, then bundling and reduced valuation will improve liquidity. However, when there is no such contraction of the support, then bundling or limited information disclosure will reduce liquidity.

In this paper, there are no gains from spanning. In fact, there may be real reasons a market in Houston mortgages is highly desirable. If there were real gains from having markets in unbundled assets, then the advantages of bundling should unsurprisingly fall. Likewise, information does not really play any positive role here (except when it increases liquidity). If more information allows the asset to be properly placed in the right portfolio, then the value of reducing information should decline. Therefore, we should expect to see bundling more often when the gains from spanning or the advantages of knowing exactly the correct market are small.

We also find that liquidity is strictly increasing in holding costs. It is still not optimal, however, to design assets with larger degrees of holding costs. The gains in liquidity are more than offset by the losses in increased costs of illiquidity. Still, the implication that liquidity rises with holding costs should hold across markets or across intermediaries.

How These Results Differ from the Previous Literature

One contribution of this paper, over previous asymmetric information models of bundling, such as Subrahamanyam, is that we have presented a model of asymmetric information and bundling that is radically different in its structure from standard models. We believe our structure conforms reasonably well to the actual attributes of the MBS and other markets. We believe it useful to know that the basic intuition of asymmetric information, bundling, and market liquidity is not model specific. Asymmetric information models of bundling, such as ours and Subrahamanyam's, are connected to the Boot and Thakor (1993) model of bundling. Ultimately, since their model focuses on competition among informed intermediaries and our paper has no competition of that kind, the explanations are quite different. Tashjian and Fishman (1995) suggest that bundling exists because it enables price discrimination. In principle, it might be possible to empirically distinguish the three explanations.

Second, we connect bundling of information and limited information disclosure. We argue that these two actions have a common effect limiting the variance of asset valuations for some financial actors. Understanding that these two actions both achieve a very similar end and should be seen in similar circumstances (such as the MBS market) is also a new insight. This insight, to some extent, means that the information disclosure literature can be used, in many cases, to understand bundling.

Third, we argue (and to our knowledge this contribution is new) that contracting the variance of asset valuations has two effects: (1) increasing the minimum value that the asset could have and (2) creating a mean preserving contraction in the distribution of asset valuations. Increasing the minimum value which the asset could have will increase liquidity, since we have eliminated some of the worst cases in the lemons problem. However, it turns out that the further apart asset valuations are (holding the minimum value constant), the easier it becomes to enforce truth-telling. Subsequently, liquidity will rise.

MBSs and the TBA Market

Our model follows Boot and Thakor (1993) and suggests that withholding information can often increase liquidity. These liquidity motives can explain why MBS issuers (such as GNMA or FHLMC) withhold information about the members of mortgage pools. However, in our model, bundling is only optimal if it works to contract the range of possible values for the mortgage pool. While we have no hard evidence on this claim, we believe that this is likely to be an accurate representation of the reality of the mortgage-backed securities. In areas where bundling just acts as a mean-preserving contraction of asset values and does not lead to a reduced range of asset values, we should not expect to see such bundling.

The value of bundling can also explain why there exists a market in undifferentiated bundles of mortgages. Each time an agent sells a mortgage bundle in that market, he gains more by withholding information than he would by explaining what exactly is in the bundle being sold. We think that this is one explanation why the TBA market exists.

Another possible view of the TBA market is that it is a market in the "worst" possible mortgage bundle. Since agents are selling an obligation to provide any bundle, the only bundle they are required to provide is the worst bundle. Naturally, this security will be highly liquid, since there will be only one value which the market as a whole assigns to the worst security. In both models, the liquidity of the worst possible asset was always one. This worst bundle also has the attributes Gorton and Pennachi (1990) expect in a liquid security. It will be free from idiosyncratic risk elements but still allows trading as a systematic risk object.

VII. CONCLUSION

We have argued that restriction of information may play a natural role in optimal asset design. Better information can act to increase the amount of informational asymmetries and decrease the overall quantity of liquidity. In both examples, the market liquidity of an asset sold at a price P is weakly decreasing in the quantity of overall information. This generation of liquidity can explain why asset issuers withhold certain types of information.

Of course, even when buyers are completely uninformed, more information can still act to increase overall market liquidity. The liquidity functions that we examined were convex, and since more information leads to a greater variance of ask prices, this greater variance will translate into more liquidity. In the first model, we found that this result suggests that issuers will try to guarantee that all intermediaries have a high probability of being informed, although the negative effects of information described above suggests that issuers want these intermediaries to be as badly informed as possible. In the second model, we found that this positive effect of information is likely to dominate when initial liquidity is high. When initial liquidity is low, the negative effects of information are likely to be more important.

These results seem to explain some of the paradoxes of mortgage-backed securities. Since mortgage-backed securities have high degrees of variance, they seem like assets where more information might reduce liquidity. Furthermore, since liquidity is already weak in the specific pool market, the conditions under which more information seems to be undesirable are satisfied.

REFERENCES

Akerlof, G. (1970). The market for ‘lemons’: Quality uncertainty and the market mechanism, *Quart. J. Econ.* **84,** 488–500.

Allen, F., and Gale, D. (1988). Optimal security design, *Rev. Finan. Stud.* **1,** 229–263.

Bartlett, R. (1989). “Mortgage-Backed Securities,” New York Financial Institute, New York.

Boot, A., and Thakor, A. (1993). Security design, *J. Finance* **48,** 1349–1378.

Chari, V., and Jagannathan, R. (1989). Adverse selection in a model of real estate lending, *J. Finance* **44,** 499–508.

Chiesa, G. (1992). Debt and warrants: Agency problems and mechanism design, *J. Financ. Intermediation* **2,** 237–254.

Cho, I., and Kreps, D. (1987). Signaling games and stable equilibria, *Quart. J. Econ.* **102,** 179–221.

DeMarzo, P., and Duffie, D. (1995). A liquidity based model of security design, *J. Econ. Theory* **68,** 185–206.

Diamond, D. (1985). Optimal release of information of firms, *J. Finance* **40,** 1071–1094.

Diamond, D., and Verrechia, R. (1991). Disclosure, liquidity and the cost of capital,'' *J. Finance* **46,** 1325–1359.

Dunn, K., and Spatt, C. (1988). Private information and incentives: Implications for mortgage contract terms and pricing, *J. Real Estate Finan. Econ.* **1,** 33–56.

Fishman, M., and Hagerty, K. (1990). The optimal amount of discretion to allow in disclosure, *Quart. J. Econ.* **105,** 427–444.

Fishman, M., and Hagerty, K. (1995). The incentive to sell financial market information, *J. Financ. Intermediation* **4,** 95–115.

Gorton, G., and Pennachi, G. (1990). Financial intermediaries and liquidity creation, *J. Finance* **45,** 45–71.

Gorton, G., and Pennachi, G. (1989). Security baskets and index-linked securities, mimeographed.

Grossman, S., and Miller, M. (1988). Liquidity and market structure, *J. Finance* **43,** 617–637.

Keynes, J. (1930). ''Treatise on Money,'' Macmillan, London.

Kim, O., and Verrechia, R. (1994). Market liquidity and volume around earnings announcements, *J. Accounting Econ.* **17,** 41–67.

Kohlberg, E., and Mertens, J. (1986). On the strategic stability of equilibria, *Econometrica* **54,** 1003–1037.

Kyle, A. (1985). Continuous auctions and insider trading, *Econometrica* **53,** 1313–1335.

Lewis, M. (1989). ''Liar's Poker,'' Norton, New York.

Myerson, R. (1979). Incentive compatibility and the bargaining problem, *Econometrica* **54,** 1003–1037.

Subrahamanyam, A. (1992). A theory of trading in stock index futures, *Rev. Finan. Stud.* **4,** 17–52.

Tashjian, E., and Weissman, M. (1995). Advantages to competing with yourself: Why an exchange might design futures contracts with correlated payoffs, *J. Financ. Intermediation* **4,** 133–157.

Yosha, O. (1995). Information disclosure costs and the choice of financing source, *J. Finan. Intermediation* **4,** 3–20.