Equilibrium Valuation of Illiquid Assets

John Krainer
Federal Reserve Bank of San Francisco

Stephen F. LeRoy
University of California, Santa Barbara

March 9, 2001

Abstract

We develop an equilibrium model of illiquid asset valuation based on search and matching. We propose several measures of illiquidity and show how these measures behave. We also show that the equilibrium amount of search may be less than, equal to or greater than the amount of search that is socially optimal. Finally, we show that excess returns on illiquid assets are fair games if returns are defined to include the appropriate shadow prices (JEL classifications G12, D40, D83).

We are indebted to Tom Cooley, John McCall and Cliff Smith for helpful comments. We have also benefited from comments by a referee and an associate editor of this journal. The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

1 Introduction

Illiquid markets are characterized in informal discussion as markets in which transactions can be completed only with a delay. By this it is meant that optimal behavior by buyers and sellers is inconsistent with immediate completion of transactions; immediate completion of transactions in illiquid markets either is impossible or is attainable only on disadvantageous terms. Some markets—those for Treasury bills—are highly liquid. Others—retail markets for real estate or used cars—are fairly illiquid. Still others—collectibles—are very illiquid.

In the preceding paragraph we characterized markets as illiquid, not assets. Many assets are traded both on illiquid and liquid markets. For example, real estate assets are illiquid when traded retail, liquid when traded as shares of real-estate investment trusts. The assets of Ford Motor Company, consisting of auto factories, are very illiquid, but Ford stock is liquid. The principal asset of Microsoft is Bill Gates' marketing ability, which cannot be directly traded at all due to the constitutional prohibition of slavery but, again, Microsoft stock is very liquid. Mortgages are of intermediate liquidity when traded directly, but are much more liquid when combined into mortgage pools. The reason government agencies insure these pools is to increase their liquidity. Thus assets themselves cannot be characterized as to liquidity since they can be traded either directly on illiquid markets or indirectly on liquid markets as securities, or both.

In this paper we present a model of equilibrium valuation of assets traded in both illiquid and liquid markets. For the present purpose, illiquidity has four components. First, the asset in question is heterogeneous. Heterogeneity by itself, however, does not imply illiquidity: Ricardian land is heterogeneous, but not illiquid. Second, asset quality can be determined only via costly search, resulting in noncompetitive markets. Third, illiquidity implies an element of irreversibility: acquisition of an illiquid asset involves a cost that cannot be recouped completely if the asset is subsequently sold.[1] Fourth, the assets traded on illiquid markets are indivisible: one can buy a small house, but not half a house. The model to be presented has all four components.

The term "liquidity" is often used with connotations different from those listed in the preceding discussion or incorporated in the model to be presented. In the market microstructure literature in finance the term "illiquidity" refers to the bid-ask spread that a market-maker imposes in dealing with buyers and sellers. Imposition of bid-ask spreads is how security specialists protect themselves when trading with agents some of whom have superior information (Glosten and Milgrom [3]). The presence of bid-ask spreads brings home the point that elements of illiquidity remain even when assets are traded as securities. Since the emphasis in this paper is on the heterogeneity of the assets being traded, the model here does not apply, directly at least, to the analysis of liquidity in securities markets.

In our model, agents consume two goods: housing services and a background good. They are risk neutral in both goods. Agents have an infinite horizon, and have a common rate of time preference β. Consumption of the background good can be either positive or negative. Agents' endowments of the background good are zero, so an agent's consumption of the background good at any date equals the negative of his net expenditure on housing at that date. Under this specification there is no need to incorporate markets in financial claims on the background good in the model: agents have no incentive either to shift consumption over time or to transfer risk among themselves. Including financial markets in the model would be possible—in fact, easy—precisely because doing so would not materially alter the equilibrium.

Agents can consume housing services only by buying a house. They can own more than one house, but can consume housing services only from one house at a time. An agent who lives in a house is said to have a "match", and the quantity of

[1] The link between liquidity and flexibility was emphasized by Jones and Ostroy [4]. Wheaton [11] and Williams [12] have models of real estate illiquidity that are related to ours.

housing services provided per period, ϵ, is called the "fit". An agent with a match does not search for new housing; he consumes housing services from his current home until the match fails, an event that occurs with probability $1-\pi$ at each date.

The assumption that agents must forego the opportunity to search upon buying a house is our (admittedly ad hoc) way of capturing the element of irreversibility that, on our definition, is inherent in the idea of illiquidity. We choose this specification instead of other possible specifications because under the alternatives the analysis would be considerably more difficult.

The interpretation of the match failing is that the agent now needs a house with different characteristics—location, size or amenities, for example. In the model, when the match fails the house no longer furnishes any housing services. Therefore the agent begins searching for a new house. Agents without a match visit exactly one house that is for sale per period. Having inspected the house, the prospective buyer knows the fit. After comparing the fit with the sale price, the buyer decides whether or not to buy the house.

The fit is not observed by the seller and cannot be credibly communicated to him. The seller posts a take-it-or-leave-it price for the house, with no subsequent bargaining. If the prospective buyer buys the house his consumption of the background good equals the negative of the purchase price of the house, and he consumes housing services until the new match fails. At that time he offers the house for sale and again begins a search for housing. If he declines the house, he consumes no housing services in that period and continues the search for a house in the next period.

As soon as a match fails, the house in question becomes a financial asset to be disposed of optimally. There is no rental market, so the agent will immediately offer the house for sale, and will keep the house on the market until it is sold. It is assumed that the number of agents equals the number of houses, and that each house that is for sale is visited by exactly one prospective buyer per period. It is possible for an owner to have no houses, one house or several houses on the market, depending on his luck at finding buyers and at maintaining his own match.

The agent's problem as a buyer consists in formulating a decision rule that governs whether he buys the house he inspected. As a seller he must decide how much to charge for a house (or houses) that he is selling. These rules, of course, apply only when the agent does not have a match in the first case, and only when the agent has a positive inventory of houses in the second case.

It turns out that linear utility has the agreeable implication that these problems are decoupled: the optimal buy rule is independent of how many houses the agent is selling, and the optimal sale price does not depend on whether the agent has a match or on the number of houses that he has for sale.

We seek a stationary symmetric Nash equilibrium: an equilibrium in which each agent's decision rules are best responses to the same decision rules when adopted by other agents, and in which equilibrium variables are constant over time.

Note that market clearing is not involved in the notion of equilibrium relevant for the valuation of assets traded on illiquid markets: at the end of a typical period many would-be buyers have not bought, and many would-be sellers have not sold. This implies that an argument that is available in analyzing valuation on liquid markets—prices are as they are because otherwise markets could not clear—is not available in analyzing asset valuation on illiquid markets. If the seller overprices the house he is selling, prospective buyers who at a lower price might have bought the house will pass on it. Therefore the seller will wait too long before selling the house, on average. There is no sense in which markets fail to clear here. If assets have the wrong prices, the interpretation is that some or all market participants are acting suboptimally in their responses to each other.

The optimal decision rules are easy to characterize informally. With regard to the buy rule, an agent can compute the value of owning a house by capitalizing the expected housing services the house provides. In this calculation the agent makes appropriate allowance for the possibility that the match will fail, implying that the house will then be offered for sale. Under the optimal buy rule the agent buys the house under consideration only if the estimated value exceeds the price by an amount which equals the discounted value of the opportunity to continue to search for housing.

With regard to the sell rule, the seller weighs the benefit of a high price—higher revenue if the house sells—against a lower probability of the house selling. If the house does not sell the seller must hold it without receiving revenue until the next period, which is costly because of the time value of money. The optimal price is high enough to afford an adequate capital gain, but not so high as to reduce prohibitively the probability of sale. A "motivated seller", in realtors' parlance, could sell a house quickly by setting a low sale price, but optimization entails setting a higher price and waiting for a buyer who is willing to pay it.

Note that in illiquid markets, the sale of a house is a positive net-present-value event for both the buyer and the seller, in contrast to the case in liquid markets. The buyer has a wealth increase equal to the capitalized value of the consumer surplus. Similarly, the seller receives a capital gain upon sale: precisely because of the possibility that the house will not sell immediately, its value unsold is strictly less than the sale price. These features of our model correspond to real-world housing markets, where signing a sale contract is good news for both buyer and seller (and their agents).

The model just described captures the essential features of illiquidity as characterized above. It is described more formally in the next section.

In Section 3 we go on to present a general discussion of liquidity in the context of our model. Specifically, we consider the suitability of several possible measures of liquidity.

In Section 4 the welfare implications of our model are considered. Since participants in illiquid markets are not price takers, there is no reason to expect that equilibrium will be Pareto optimal, and (generically) it is not. It turns out that there may be either too much or too little search, depending on parameters.

In Section 5 we address the question of whether equilibrium excess returns on assets traded on illiquid markets are fair games. This depends on how prices, payoffs and returns are defined. We believe that the most useful definitions are those that incorporate certain shadow prices that reflect the illiquidity of the underlying assets. Under these definitions the equilibrium conditions directly imply that excess returns are fair games.

2 The Model

Figure 1A displays the timing conventions governing the agent as buyer; Figure 1B does the same for the agent in his role as seller. Definitions of variables refer to the boxed entries in these figures (for example, to understand the determination of s, refer to the box in the lower left corner of Figure 1A).

As noted in the introduction, an agent without a match evaluates for possible purchase one and only one house at each date. The fit ϵ of any house for any prospective buyer is a random variable distributed uniformly on [0,1], IID. This distribution is common knowledge. Upon evaluating the house the buyer learns the fit, but the seller does not. As noted above, there is assumed to be no credible way for the buyer to communicate the fit, nor can the seller induce or compel him to reveal it. Thus the seller must calculate the probability that the house will sell from the distribution of ϵ and the equilibrium buy rule, whereas the buyer makes his decision based on the realization of ϵ. The seller will set the sale price accordingly.

Since the seller does not know the prospective buyer's fit, he must ask the same price regardless of the fit. Therefore the prospective buyer who decides to buy will realize a consumer's surplus the magnitude of which depends on the fit, but is always nonnegative.

INSERT FIGURE 1A HERE

If the agent buys the house, he receives housing services at rate ϵ beginning in the next period and continuing until the match is broken. By convention the housing services on a newly bought house, like those on a house bought at some time in the past, occur in the next period (so that housing is priced ex-housing services, corresponding to the convention usually adopted in finance that stocks and bonds are priced ex-dividend and ex-coupon). If the buyer elects not to buy the house he consumes no housing services, and will continue the search next period.

At the end of the period agents who entered the period with a match and those who entered the period unmatched but bought a house during the period draw random variables which determine whether their matches continue into the next period or are broken. If an agent's match persists he continues to consume housing services at the rate ϵ; if the match is broken the agent will go into the next period without a match, and will then search for a house. As noted, we seek a stationary symmetric Nash equilibrium: each agent's decision rules are a best response to other agents' behavior when other agents act according to the same decision rules. Further, the equilibrium values of decision variables (and all variables that depend on them, such as prices) are constant over time. It is assumed that buyers and sellers are anonymous, so they have no repeated interaction.

INSERT FIGURE 1B HERE.

We consider the problem of an agent without a match who has one house for sale. This specification involves no material loss of generality: when the agent has a match, or when he has no houses in inventory, then he has no role as a buyer or seller; when he has more than one house for sale he asks the same price for each house as he does when he has one house for sale, as is easily verified.

It is also readily verified that the assumptions made on preferences imply that the agent's buying problem is decoupled from his selling problem. Therefore we may consider the two separately. Consider first the agent in his role as buyer. The buyer's strategy set is assumed to consist of a linear function that expresses his reservation fit $\bar{\epsilon}$ as a function of current price:

$$\bar{\epsilon} - \bar{\epsilon}^* = \delta(p - p^*), \tag{1}$$

so the decision variables are δ and p^*. Here we express reservation fit and price as deviations from their respective equilibrium values; (throughout we will use * to denote the equilibrium values of variables); this is an arbitrary normalization. The parameter δ, then, measures the effects of deviations from the equilibrium price on the reservation fit that is optimal for the buyer.

Being unmatched, the agent owns an asset that consists of the right to search for a house. Define the value of this right as s. Then s is given by

$$s = \mu\left(v\left(\frac{\bar{\epsilon}+1}{2}\right) - p^*\right) + \beta(1-\mu)s^*. \tag{2}$$

Here μ is the probability of sale, β is the agent's discount factor and $v(\epsilon)$ is the value of a house with fit ϵ. In turn, $v(\epsilon)$ is given by

$$v(\epsilon) \equiv \beta\epsilon + \beta\pi v(\epsilon) + \beta(1-\pi)(q+s). \tag{3}$$

In (3) π is the probability of preserving the match, and q is the value of a house to the owner after he has lost his match. The argument of v in (2) equals the expectation of ϵ conditional on $\epsilon \geq \bar{\epsilon}$, so $v\left(\frac{\bar{\epsilon}+1}{2}\right) - p^*$ equals the expectation of the buyer's surplus conditional on having a fit that exceeds the reservation fit $\bar{\epsilon}$. Solving for $v(\epsilon)$, (3) becomes

$$v(\epsilon) \equiv \frac{\beta\epsilon + \beta(1-\pi)(q+s)}{1-\beta\pi}. \tag{4}$$

The buyer's decision problem is to find the value of $\bar{\epsilon}$ that maximizes s in (2), for any p. Under symmetric Nash equilibrium the buyer takes the value of s on the right-hand side of (2) as given. This specification reflects the assumption that the buyer will set future values of $\bar{\epsilon}$ at the equilibrium level in deciding whether to buy now, and also that the future values of p that the buyer will face will equal the equilibrium value. The buyer evaluates μ from

$$\mu = 1 - \bar{\epsilon}, \tag{5}$$

which follows from the fact that ϵ is uniformly distributed on the unit interval.

Substituting (5) in (2) and using (4), the first-order condition for a maximum of s with respect to $\bar{\epsilon}$ is

$$\frac{\beta(1-\bar{\epsilon})}{2(1-\beta\pi)} - v\left(\frac{\bar{\epsilon}+1}{2}\right) + p^* + \beta s^* = 0. \tag{6}$$

This condition can be put in a form that is more readily interpreted. Using

$$v\left(\frac{\bar{\epsilon}+1}{2}\right) = v(\bar{\epsilon}) + \frac{\beta(1-\bar{\epsilon})}{2(1-\beta\pi)}, \tag{7}$$

which in turn follows from (4), (6) simplifies to

$$v(\bar{\epsilon}) = p^* + \beta s^*. \tag{8}$$

This equation states that at the reservation fit, the expected utility of owning the house equals its price plus the discounted value of search, reflecting the fact that the buyer gives up the right to search if he elects to buy the house. The value of $\bar{\epsilon}$ that solves (8) is the equilibrium value $\bar{\epsilon}^*$:

$$v(\bar{\epsilon}^*) = p^* + \beta s^*. \tag{9}$$

To derive the value of δ in (1) it is necessary to relax the assumption that the current value of p in (9) equals its equilibrium value. The optimal value of $\bar{\epsilon}$ for arbitrary p satisfies

$$v(\bar{\epsilon}) = p + \beta s^*. \tag{10}$$

Subtracting (9) from (10) and using (4) to solve for $\bar{\epsilon} - \bar{\epsilon}^*$, there results

$$\bar{\epsilon} - \bar{\epsilon}^* = (\beta^{-1} - \pi)(p - p^*), \tag{11}$$

so we have

$$\delta^* = (\beta^{-1} - \pi). \tag{12}$$

We now turn to the seller's problem. The seller owns an asset, an unsold house, with value q. His problem is to choose the optimal price p. The wholesale price q and retail price p satisfy

$$q = \mu p + \beta(1-\mu)q^* \tag{13}$$

or, using (5),

$$q = (1-\bar{\epsilon})p + \beta\bar{\epsilon}q^*. \tag{14}$$

The first-order condition associated with maximizing q in (14) with respect to p is

$$(1-\bar{\epsilon}) + \frac{d\bar{\epsilon}}{dp}(\beta q^* - p) = \mu + \frac{d\bar{\epsilon}}{dp}(\beta q^* - p) = 0; \tag{15}$$

here the term that includes $d\bar{\epsilon}/dp$ reflects the seller's recognition that his choice of p affects the buyer's reservation fit. From (11) we have

$$\frac{d\bar{\epsilon}}{dp} = \beta^{-1} - \pi, \tag{16}$$

so the first-order condition becomes

$$\mu + (\beta^{-1} - \pi)(\beta q^* - p) = 0. \tag{17}$$

The values of μ and p that solve the first-order condition are equilibrium values:

$$\mu^* + (\beta^{-1} - \pi)(\beta q^* - p^*) = 0. \tag{18}$$

The model has five equations. The first three,

$$s^* = \mu^*\left(v\left(\frac{\bar{\epsilon}^* + 1}{2}\right) - p^*\right) + \beta(1-\mu^*)s^*, \tag{19}$$

$$q^* = \mu^* p^* + \beta(1-\mu^*)q^*, \tag{20}$$

and

$$\mu^* = 1 - \bar{\epsilon}^*, \tag{21}$$

are the equilibrium counterparts of (2), (13) and (5), respectively. The others are the equilibrium versions of the first-order conditions (9) and (18). There are five unknowns: q^*, p^*, μ^*, $\bar{\epsilon}^*$ and s^*. A solution to this system of equations is a stationary symmetric Nash equilibrium. These equations, although nonlinear, are easily solved numerically.

The model has a minor loose end. We have not specified the number of agents. If there exists a finite number of agents, then a single agent could conceivably own all the houses in the economy at some date. In that case there arises the question of what house he inspects if his match fails. We ignore such events since over any finite time interval they occur with low probability if the number of agents and houses is large. The problem can be avoided altogether if it assumed that the number of agents is infinite, but that would entail analytical complications.[2]

To determine existence of a solution to the model, begin by using (21) to eliminate $\overline{\epsilon}^*$ in (9) and (19). If we fix μ^*, the remaining equations are affine (and, as is easily checked, linearly independent). Therefore p^*, q^* and s^* are uniquely determined as functions of μ^*. Thus the equations of the model define a map—call it Ψ—from the unit interval to itself. Since Ψ is continuous, the Schauder fixed-point theorem (Stokey and Lucas, with Prescott [10]) implies existence of a solution (see also Krainer [5]).

There is a parallel between this argument and the derivation of equilibrium in Lucas's [7] study of asset prices. Recall that Lucas used two contraction arguments: in the first, equilibrium policy functions and value functions were derived via a contraction, taking asset prices as given. Second, a map taking the space of asset prices into itself was found that was also a contraction. Here the first stage of the derivation—solution for $\overline{\epsilon}^*$, p and μ as functions of q and s—is similar to Lucas's derivation of equilibrium policy rules and value functions; in the present case the contraction argument can be dispensed with since the relevant functions are linear. Determining equilibrium values of the state variables q and s is similar to Lucas's second-stage contraction. Here the second stage involves continuation values of state variables rather than asset prices; this difference reflects the fact that the solution concept here is symmetric Nash equilibrium rather than competitive equilibrium as in Lucas [7].

3 Measures of Liquidity

The model just presented suggests several possible measures of liquidity. One is the expected time to sale, $(1-\mu^*)/\mu^*$, with low values of this measure corresponding to high liquidity. This measure is appropriate for both buyer and seller.

A second measure of liquidity, appropriate for the seller, is the ratio of the retail price of a house p^* to its wholesale price q^*; the difference between p^* and q^* measures the capital gain a seller experiences when a house sells. This variable equals 1 for liquid assets (the value of a liquid asset to its owner just prior to sale equals its value when sold). The lower the value of p^*/q^*, the higher the level of liquidity.

To investigate whether the interpretation of these variables is correct, we conducted a comparative statics experiment designed to vary liquidity. In our model

[2] This difficulty occurs frequently in economics and finance. For example, in discussing the arbitrage pricing theory it is customary to discuss diversified portfolios in a setting where only finite portfolios, which cannot be completely diversified, are explicitly modeled. This practice is acceptable because it is known that if infinite portfolios are specified, then diversified portfolios can be explicitly modeled, and omitting doing so does not distort the results.

houses are illiquid because buyers can evaluate only one house per period. The easiest way to vary liquidity is therefore to alter parameter values so as to change the effective length of the period. The expectation is that when the period is short, so that buyers search frequently, the housing market behaves much like a liquid market: the average fit is high, the average time to sale is short and the wholesale price of a house is almost equal to its retail price.

We first computed a benchmark equilibrium based on $\pi = 0.9$ and $\beta = .95$ (corresponding to an average occupancy duration of nine years and a real interest rate of five per cent per year). Then we assumed that there are n periods per year, for various values of n. For each run we defined $\beta_n \equiv \beta^{1/n}$ and $\pi_n \equiv \pi^{1/n}$. Also, we assumed that housing services are distributed uniformly on $[0, 1/n]$ instead of $[0, 1]$, so as to preserve the scale of housing prices. The ratio p^*/q^* does not require rescaling, but the expected time to sale is redefined to equal $(1 - \mu^*)/n\mu^*$ so as to measure in years rather than periods.

INSERT FIGURE 2 HERE

Figure 2 shows the equilibrium values of p^*, q^* and s^* as functions of n, for selected values of n. Figure 3 shows the measures of liquidity μ^*, $(1 - \mu^*)/n\mu^*$ and p^*/q^* as functions of n. When n is high the probability of sale during any period, μ^*, is low since the prospective buyer will buy the house only if the fit is very high. The buyer is willing to pass on the house currently being evaluated unless the fit is very high, since he does without housing services for only a short interval before searching again. Correspondingly, when n is high the seller charges a high price for the house since he knows that if the current prospective buyer does not buy, another prospective buyer will be along shortly, and the cost of holding the house vacant for a short time is low. The wholesale price of a house also rises with n since for high n rapid sale is very likely. The value of search s^* also rises with n.

INSERT FIGURE 3 HERE

Figure 3 shows that the measures of liquidity behave as expected. Even though the probability of sale during any period is low when n is high, the expected time to sale is low (since $(1 - \mu^*)/\mu^*$ increases more slowly than n). When n is high, both p^* and q^* are high, but the spread between them is small. Therefore p^*/q^* is only slightly higher than 1.

4 Welfare

As with most models involving search and matching, equilibrium here is not Pareto optimal. This is to be expected: the choice the seller faces between high price and high probability of sale is formally identical to the choice the monopolist faces between high price and high quantity sold. With non-price-taking behavior, the first welfare theorem does not apply. However, in contrast to the case of the static monopolist, here the equilibrium reservation fit $\overline{\epsilon}^*$ can be either higher or lower (or, in a borderline case, equal to) the optimal reservation fit, depending on β and π.

The planning problem that corresponds to the equilibrium analyzed here consists of determining a reservation match $\hat{\epsilon}$ such that agents stop searching and move into a given house when $\epsilon \geq \hat{\epsilon}$. The planner chooses $\hat{\epsilon}$ to maximize aggregate expected utility.

Aggregate expected utility is a weighted average of discounted average housing consumption of matched and unmatched agents, where the weights are the respective fractions of each in the population. Formally, the planner maximizes

$$\begin{aligned} W &= \sum_t \beta^t [\Pr(\text{matched})_t (\text{avg. housing consumption})_t \\ &\quad + \Pr(\text{unmatched})_t (\text{avg. housing consumption})_t .] \end{aligned} \tag{22}$$

At each t, average housing consumption by matched agents is $(\hat{\epsilon}+1)/2$, where $\hat{\epsilon}$ is the reservation match.[3] Housing consumption by unmatched agents equals zero. In a steady state, the probability of being matched is constant over time, so maximizing W is equivalent to maximizing

$$\Pr(\text{matched}) \left(\frac{\hat{\epsilon}+1}{2} \right). \tag{23}$$

The (invariant) probability that an agent is matched for a given reservation fit $\hat{\epsilon}$ is easily calculated from the conditional probabilities that an agent remains in his current state (see Figure 1A). The transition matrix is

$$T = \begin{bmatrix} \pi & 1-\pi \\ \pi(1-\hat{\epsilon}) & \hat{\epsilon} + (1-\pi)(1-\hat{\epsilon}) \end{bmatrix}. \tag{24}$$

Here the element T_{22} is the probability that an unmatched agent in the beginning of the period will be unmatched at the beginning of the next period. This event can occur in two ways (see Figure 1A, where the states at each date are enclosed in boxes). First, the search fails with probability $\hat{\epsilon}$. Second, with probability $(1-\pi)(1-\hat{\epsilon})$ the agent buys a house but loses the match in the next period. The other elements of T are self-explanatory.

It follows that the unconditional probability that an agent has a match is $(\pi(1-\hat{\epsilon}))/(1-\pi\hat{\epsilon})$. The optimal value of $\hat{\epsilon}$ is therefore that which maximizes

$$\left(\frac{\pi(1-\hat{\epsilon})}{1-\pi\hat{\epsilon}} \right) \left(\frac{\hat{\epsilon}+1}{2} \right). \tag{25}$$

Denote this $\tilde{\epsilon}$. We obtain

$$\tilde{\epsilon} = \arg\max_{\hat{\epsilon}} \left(\frac{\pi(1-\hat{\epsilon})}{1-\pi\hat{\epsilon}} \right) \left(\frac{\hat{\epsilon}+1}{2} \right) = \frac{1-\sqrt{1-\pi^2}}{\pi}. \tag{26}$$

[3] Consumption of the background good can be deleted from eq. (22) because the negative consumption of the background good by buyers of houses cancels the positive consumption of sellers.

Note that $\tilde{\epsilon}$, in contrast to $\overline{\epsilon}^*$, does not depend on β.

INSERT FIGURE 4 HERE

INSERT FIGURE 5 HERE

INSERT FIGURE 6 HERE

Figure 4 displays $\tilde{\epsilon}$ and $\overline{\epsilon}^*$ against π for $\beta = 0.95$, and Figure 5 displays $\tilde{\epsilon}$ and $\overline{\epsilon}^*$against β for $\pi = 0.9$. Figure 6 shows the values of π and β for which the optimal reservation fit is greater than (equal to, less than) the equilibrium reservation fit.

Figure 4 shows that $\tilde{\epsilon} < \overline{\epsilon}^*$ for low values of π. This makes sense: as $\pi \to 0$ the expected duration of stay in a house converges to one period. A planner charged with maximizing expected consumption of the housing good will raise the homeownership rate in this environment because the returns to searching for a good match are low. This leads to $\tilde{\epsilon} \to 0$. In the equilibrium model, however, the seller will maximize expected revenue by asking a price such that the reservation fit is 0.5.

In contrast, we have $\tilde{\epsilon} > \overline{\epsilon}^*$ for low values of β (when π is sufficiently high). To understand this, recall that β does not figure in the planner's choice of b. However, β does affect $\overline{\epsilon}^*$: when $\beta \ll 1$ the seller will set a low price so as to avoid the long vacancy period that will occur if he fails to sell. In response to this low price, and also to avoid a long period of homelessness, the buyer will set a low reservation fit $\overline{\epsilon}^*$.

5 Are Returns Fair Games?

Under simplifying assumptions—principally risk neutrality—excess returns on liquid assets are fair games: the conditional expected return on any asset less the interest rate is zero. The fair-game model plays a central role in settings where one is willing to assume stationarity and to abstract away from the effects of risk aversion on asset prices. For example, the market efficiency tests reported in Fama [2] are, for the most part, tests of the fair game model.

It is generally supposed that the fair game model describes returns only in markets that are perfectly liquid. The basis for this presumption is that the simplest justification for the fair game model does in fact require market liquidity. This justification consists of the observation that if there existed some asset with an expected return that differed from the interest rate, then a single (well-financed, risk-neutral, price-taking) investor could generate an expected utility gain by borrowing and buying the mispriced asset, or the reverse. This investor, being risk neutral, would continue to trade until fair game asset prices were reestablished.

However, in the case of illiquid assets, transaction costs generally prevent the investor from bidding away the return differentials. Therefore, the argument concludes, one would not expect to end up with a fair game. It would seem that autocorrelated returns to real estate, for example, could coexist with a constant interest rate because the illiquid nature of real estate prevents any investor from conducting the trades that in liquid markets would restore fair games.

This argument is unsatisfactory. It confuses necessity and sufficiency. It is correct that if markets are liquid, then one can justify the fair game model by appealing to the behavior of a single risk-neutral investor. It is also correct that this argument fails if markets are illiquid. It does not follow from these facts that perfect liquidity is necessary for the fair game model (as, in fact, Fama was careful to point out). Asset returns in liquid markets behave as they do, not because otherwise a single agent could conduct profitable trades, but because otherwise the optimal trading rules of agents collectively are mutually incompatible.

So far, it appears, we have no argument either way about whether returns on illiquid assets are fair games. The question has not been investigated, no doubt due to the fact that we have little experience building models of equilibrium valuation of illiquid assets.

In this section the properties of the equilibrium distributions of returns are analyzed. Here the (gross) return on an asset has the usual definition as the value of its payoff (dividend or service flow plus next-period asset value) divided by current asset value.

In the model of this paper there are three sources of wealth. First, any agent, matched or not, may own one or more houses that he no longer lives in. All unoccupied houses are always offered for sale at price p^*. Prior to sale they have value q^* per house. Second, a matched agent with fit ϵ owns an asset with value $v(\epsilon)$. Third, an unmatched agent owns the search option, which has value s^*. We consider the returns on each asset in turn.

First, the equilibrium distribution of the return on a house offered for sale is

$$r^* = \begin{cases} p^*/\beta q^* \text{ with probability } \mu^* \\ 1 \text{ with probability } 1-\mu^* \end{cases}. \tag{27}$$

To see this, observe that if the house sells its payoff is p^*. However, under our convention on notation the proceeds of the sale are paid to the seller in the current period, not the next period. The next-period value of the payoff if the house sells is therefore $\beta^{-1}p^*$. If the house does not sell, its next-period value is q^*. Since the current value of the house is q^*, the return distribution is as shown in eq. (27). The expected return is given by

$$E(r^*) = \mu^* p^*/\beta q^* + (1-\mu^*). \tag{28}$$

Using eq. (20), eq. (28) simplifies to

$$E(r^*) = \beta^{-1}. \tag{29}$$

Thus the expected return equals investors' time preference.

Second, the return distribution on an owner-occupied house is

$$r^* = \begin{cases} \epsilon/v(\epsilon) + 1 \text{ with probability } \pi \\ (\epsilon + q^* + s^*)/v(\epsilon) \text{ with probability 1-}\pi \end{cases}. \tag{30}$$

Eq. (30) is based on the fact that the value of an owner-occupied house to its owner is $v(\epsilon)$, not p^* or $p^* + \beta s^*$. The next-period payoff on the house is $\epsilon + v(\epsilon)$ if the match is not broken, and $\epsilon + q^* + s^*$ if the match is broken. Taking the expectation and using eq. (4), it follows that the expected rate of return on an owner-occupied house is also given by eq. (29).

Third, an agent without a match owns the search option, the current value of which is s^*. The expected return on the search option conditional on buying or not buying is

$$r^* = \begin{cases} \left(\frac{\bar{\epsilon}^*+1}{2} + \pi v\left(\frac{\bar{\epsilon}^*+1}{2}\right) + (1-\pi)(q^*+s^*) - \frac{p^*}{\beta} \right) / s^* \text{ with probability } \mu^* \\ 1 \text{ with probability } 1-\mu^* \end{cases} . \quad (31)$$

Note that p^* is multiplied by the interest rate because returns are defined as next period payoffs divided by current period value. From eq. (3) and (19), eq. (29) results.

In all three cases the excess returns $r - \beta^{-1}$ just characterized are seen to be fair games: the expected excess returns conditional on the values of any or all of an agent's state variables are zero.

A large quantity of empirical evidence (for example, Case and Shiller [1], Meese and Wallace [8]) supports the conclusion that returns on housing are positively autocorrelated. However, in empirical work returns are defined as price changes, whereas we have seen that the appropriate definition implies not only that implicit rent should be included in the payoff of housing, but also that the capitalized consumer surplus and the value of the search option should be included in the value of housing. All these variables are unobservable, and it is not easy to think of proxies. Thus testing the fair game proposition as it applies to illiquid assets is not straightforward.

A more promising research strategy is to test the model by determining its predictions for return and price variables that one can measure, rather than by trying to construct a proxy for the theoretically correct return measure. The present version of the model is not well suited to this task, since it predicts that there are no price changes. However, the model can be modified to include aggregate shocks to housing services. If this is done then the empirical association between returns and the various liquidity measures can be investigated. Preliminary results along these lines are reported in Krainer [5].

6 Conclusion

Asset illiquidity in this paper is generated by asymmetric information between buyers and sellers and by a restriction that agents search sequentially for trading partners. As the time between potential transactions shrinks, expected time to sale decreases. In this model, as in many models of search, equilibrium is not efficient: two matched agents could exchange houses to mutual advantage. More important, however, equilibrium asset liquidity will most likely not be socially optimal.

The analysis of this paper depends critically on the assumption that agents are risk neutral. Excess returns on assets traded in illiquid markets, like those traded in liquid markets, will not be fair games (with respect to the natural probabilities) if agents are risk averse.

The analysis of the preceding section makes clear that the principles of valuing illiquid assets are essentially the same as those that apply to liquid assets, and also that these principles will carry over from the case of risk neutrality analyzed in this paper to the general case of risk aversion. Specifically, risk premia on assets traded on illiquid markets will be governed by the covariance of their payoffs with the marginal utility of consumption, just as with liquid assets. Thus the theory of consumption-based asset pricing applies just as much to illiquid assets as to liquid assets. The problem posed by illiquid assets is not that there is any ambiguity about the relevant theory, but that the definitions of the relevant returns include shadow prices (in our model, q^* and s^*) that are difficult to measure.

There is a presumption that assets traded in illiquid markets have positive risk premia. This is so because, as seen in the present paper, the sale of an asset traded in illiquid markets creates wealth for the seller, since the sale price exceeds the value of the asset being offered for sale that has not yet been sold. This effect induces a negative covariance between the payoff of an asset being offered for sale and the marginal utility of consumption. Of course, this negative covariance can be offset by other factors, depending on other aspects of the model. This qualification aside, the conclusion is that illiquidity by itself gives rise to positive risk premia.

We pointed out above that agents are indifferent between selling houses as illiquid assets on the retail market at price p^* (but probably with a delay) or immediately as liquid assets on the wholesale market at price q^*. Thus there is no room in the present model for fire sale prices. Distress sales reflect capital market imperfections, which we ruled out. A logical next step in the analysis of illiquid asset valuation would be to incorporate capital market imperfections in the analysis of illiquidity (see Stein [9]).

References

[1] Karl E. Case and Robert J. Shiller. Forecasting prices and excess returns in the housing market. AREUEA Journal, 18(3): 253-273, 1990.

[2] Eugene F. Fama. Efficient capital markets: A review of theory and empirical work. Journal of Finance, 25: 283-417, 1970.

[3] Lawrence R. Glosten and Paul R. Milgrom. Bid, ask, and transaction prices in a specialist model with heterogeneously informed traders. Journal of Financial Economics, 14: 71-100, 1985.

[4] Robert A. Jones and Joseph M. Ostroy. Flexibility and uncertainty. Review of Economic Studies, LI: 13-32, 1984.

[5] John Krainer. Pricing illiquid assets with a matching model. reproduced, University of Minnesota, 1997.

[6] Steven A. Lippman and John J. McCall. An operational measure of liquidity. American Economic Review, 76: 43-55, 1986.

[7] Robert E. Lucas. Asset prices in an exchange economy. Econometrica, 46: 1429-1445, 1978.

[8] Richard Meese and Nancy Wallace. Testing the present value relation for housing prices: should I leave my house in San Francisco? Journal of Urban Economics, 35: 245-266, 1994.

[9] Jeremy C. Stein. Prices and trading volume in the housing market: a model with downpayment effects. Quarterly Journal of Economics, 110: 379-406, 1995.

[10] Nancy Stokey, Robert E. Lucas, with Edward C. Prescott. Recursive Methods in Economic Dynamics. Harvard University Press, Cambridge, MA, 1989.

[11] William C. Wheaton. Vacancy, search, and prices in a housing market matching model. Journal of Political Economy, 61: 1270-1292, 1990.

[12] Joseph T. Williams. Pricing real assets with costly search. Review of Financial Studies, 8: 55-90, 1995.