

Returns on Illiquid Assets: Are They Fair Games?

John Krainer
Federal Reserve Bank of San Francisco

Stephen F. LeRoy
University of California, Santa Barbara

sleroy@econ.ucsb.edu
(805) 893-2438

September 11, 1997 (first draft April 25, 1997)

*The simplest tests of capital market efficiency are tests of the fair game model: conditional expected returns less the interest rate are equal to zero.*¹ *The fair game model is thought to obtain only when markets are perfectly liquid. We show that this conjecture is false. In a model of the housing market where heterogeneous agents must search for partners in order to trade, excess returns on housing wealth are fair games if, as is appropriate, returns are defined to include shadow prices measuring illiquidity (JEL classifications G12, D40, D83).*

1 Introduction and Summary

Under simplifying assumptions, excess asset returns 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, be-

ing risk neutral, would continue to trade until fair game asset prices were reestablished.

However, in the case of illiquid assets—defined here as assets for which the optimal sale or purchase strategy entails time-consuming search—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.²

We present a model of equilibrium valuation of illiquid assets in this paper. For the present purpose, illiquidity has three components. First, the asset in question is het-

¹The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

We are indebted to John McCall for helpful comments.

²Wheaton [10] and Williams [11] have models of real estate illiquidity that are related to ours.

erogeneous. Heterogeneity by itself 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.³ The model to be presented has all three features.

The term “liquidity” is often used with connotations different from those listed in the preceding discussion or incorporated in the model to be presented. For example, illiquidity is often held to imply that assets can be sold immediately only at fire sale prices, if at all. The model to be presented (under minor modification) has the property that assets can be sold immediately at (wholesale) prices that are lower than the (retail) prices obtainable if the assets are marketed. However, these are fair prices, not fire sale prices. It would seem that the phenomenon of fire sales can only occur when credit markets fail, a specification that we do not consider (see Stein [8]).

In the market microstructure literature in finance the term “liquidity” is used in a still different way. There illiquidity is manifested in a bid-ask spread, which is how the specialist protects himself when trading with agents who may have superior information (Glosten and Milgrom [3]).

As in the model to be presented, the analysis of illiquidity in the market microstructure literature is predicated on asymmetric information. However, in the market microstructure literature the asset in question—shares of stock—is homogeneous, and there is no element of irreversibility.

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 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 in its absence 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. 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 goods decreases by 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 to buy 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, so 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

³The link between liquidity and flexibility is emphasized by Jones and Ostroy [4]

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 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.

Note that market clearing is not involved in the notion of equilibrium relevant for the analysis of illiquid asset valuation. This implies that an argument frequently involved in analyzing valuation of liquid assets—prices are as they are because otherwise markets could not clear—is not available in analyzing valuation of illiquid assets. 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, but it remains true that if assets have the wrong prices some or all market participants are acting suboptimally in their responses to each other, which is inconsistent with Nash equilibrium.

The optimal decision rules are easily characterized. 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 with a good fit.

The model just described captures the essential features of illiquidity as characterized above. It is described more fully 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 illiquidity. Strictly, this discussion is a digression from the central concern of this paper, equilibrium return distributions on illiquid assets. The discussion of measures of liquidity is presented because it is of interest in its own right, and also because it brings out some general aspects of valuation of illiquid assets.

One unambiguous measure of illiquidity is the average time to sale under optimal pricing. This measure is relevant to both buyers and sellers.

Another measure is the difference between the price at which houses sell and the value to its owner of a house for sale prior to the arrival of the successful buyer. We refer to the former as the retail price of a house and the latter as its wholesale price. This terminology is appropriate because agents would be willing to transfer ownership of houses in arbitrary quantities at the wholesale price if we had specified the model so as to permit them to do so.

It is easiest to suppress the wholesale market in housing, along with the markets for other liquid financial assets, since our assumptions imply that in equilibrium agents would be indifferent about trading on this market if permitted to do so. Despite the fact that we suppress the wholesale market, it is convenient to refer to the value of a house being offered for sale as its wholesale price. The wholesale price, of course, is predicated on the presumption that the owner continues to offer it for sale at the retail price, as is optimal.

The retail-wholesale price spread is a measure of illiquidity because it equals the price sacrifice a seller would have to accept in order to achieve immediate sale (in the version of the model permitting wholesale trading of houses). This measure of illiquidity is relevant only for sellers.

A third possible measure of illiquidity is the value of the search option. If, as suggested in the introduction, illiquidity involves irreversible commitment, then the cost of that commitment should be a measure of illiquidity; here the cost of committing to the purchase of a house is the value of the search option. We will see in Section 3 that the value of the search option behaves differently from the other two measures of illiquidity, raising questions about this interpretation.

In Section 4 we address the question of whether equilibrium excess returns are fair games. There are three distinct assets here: houses currently being offered for sale, owner-occupied houses and the search option. The equilibrium conditions directly imply that excess returns on the former are fair games: the excess capital gain if the house sells just offsets the foregone interest income

if it does not sell (each multiplied by the relevant probability).

With regard to the second, in our model the return on an owner-occupied house consists of its housing services plus its next-period value, adjusted for whatever loss in value occurs if the match is broken next period, all divided by the current value of the house. Excess returns on owner-occupied housing are fair games.

Finally, we show that the excess return on the search option is also a fair game.

2 The Model

As noted in the introduction, an agent without a match evaluates 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. There is 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 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.

If the agent buys the house, he receives housing services at rate ϵ beginning at the end of that 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 at the end of the 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 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 fit, and will begin searching for a house.

As noted, we seek a 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. It is assumed that buyers and sellers are anonymous; since they have no repeated interaction their strategy sets consist of the decision rules at a single date.

For any agent the state at any date consists of three variables: ψ , ϵ and h . Here ψ is a 0-1 variable specifying whether the agent has a match ($\psi = 1$) or not ($\psi = 0$) at the beginning of the period. As noted above, ϵ is the fit for an agent with a match (if $\psi = 0$, ϵ does not affect any subsequent behavior). The variable h is the number of unoccupied houses an agent owns, and therefore offers for sale.

An agent who is without a match and who does not own any houses still has positive wealth (defined as the expected present value of future consumption). This wealth equals the expected discounted value of future consumer's surpluses that occur whenever an agent buys a house.

Call the capitalized value of this surplus s . At any date the wealth of an agent equals

$$hq + \psi v(\epsilon) + (1 - \psi)s, \quad (1)$$

where s and q , the wholesale value of a house, are determined endogenously. Here $v(\epsilon)$, the expected utility of the housing services generated by a house for its owner-occupant if the fit is ϵ , is given by

$$v(\epsilon) \equiv \beta\epsilon + \beta\pi v(\epsilon) + \beta(1 - \pi)(q + s). \quad (2)$$

The house generates housing services ϵ at the end of the period, leading to the term $\beta\epsilon$. If the match persists, which occurs with probability π IID, the house is worth $v(\epsilon)$ again next period. If not, the value of the house to its owner is its wholesale value q . However, if the match fails the owner also recovers the opportunity to search, which contributes s to the current value of the house.

The optimal decision rule ⁴ for the buyer is:

$$\text{buy if } \eta \geq \epsilon^*$$

$$\text{do not buy if } \eta < \epsilon^*, \quad (3)$$

where η is the fit of the house the prospective buyer is currently evaluating. Here ϵ^* , the reservation fit, satisfies

$$v(\epsilon^*) = \bar{p} + \beta s. \quad (4)$$

Eq. (4) states that the marginal buyer finds that the value $v(\epsilon^*)$ of the house he is evaluating just equals the price \bar{p} asked by the seller plus the opportunity cost s of giving up the option to search again next period.

⁴These decision rules are derived in the appendix. Decision rules of the form (3) are frequently encountered in search models; see, for example, Lippman and McCall [6].

A bar over a variable indicates that that variable is determined by the behavior of another agent. Of course, the adopted equilibrium concept implies that the value of \bar{p} chosen by the seller whose house the agent is evaluating will equal the price p that the agent will charge for his own house, or houses, that he is selling. Nevertheless, the two variables must be distinguished in deriving the equilibrium.

The value of the option to search s satisfies

$$s = \mu \left(v \left(\frac{\epsilon^* + 1}{2} \right) - \bar{p} \right) + \beta(1 - \mu)s. \quad (5)$$

Here μ equals the probability that an agent searching for a house will buy the house he is currently evaluating. With probability μ the agent will buy the house. Conditional upon purchase, the expected fit is $(\epsilon^* + 1)/2$ —halfway between the reservation fit ϵ^* and the maximal fit of unity. Since v is linear, the capitalized expected consumer surplus of a house conditional upon purchase is therefore $v((\epsilon^* + 1)/2) - \bar{p}$. With probability $1 - \mu$ the buyer chooses not to buy, in which case he retains the search option next period.

When the match is broken the agent immediately puts his house up for sale. He will keep it on the market until it sells whether or not he succeeds in finding a new house, since either way the old house yields no housing services to the seller. The retail and wholesale prices are related by

$$q = \mu p + \beta(1 - \mu)q. \quad (6)$$

The optimal decision rule for the seller is to set the retail price p to maximize q :

$$p = \operatorname{argmax}_y (\mu(y)y + \beta(1 - \mu(y))q). \quad (7)$$

Here $\mu(y)$ is the probability that the house sells as a function of its sale price y . This function, which turns out to be linear, is taken as given by the seller. It is derived from the buyer's rule, eq. (3), where y is the price the buyer faces.

The first-order condition for this maximization is

$$(\pi - \beta^{-1})(p - \beta q) + \mu = 0, \quad (8)$$

where $\pi - \beta^{-1}$ equals the derivative of μ with respect to y (Appendix).

The model is closed by the equation

$$\mu = 1 - \epsilon^*, \quad (9)$$

which follows from the assumption that the fit is uniformly distributed on the interval $[0, 1]$.

We have five equations—(4), (5), (6), (8) and (9)—in the five unknowns q , p , μ , ϵ^* and s . A solution to this

system of equations is a symmetric Nash equilibrium. These equations, although nonlinear, are easily solved numerically.

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 of a sale contract is good news for both buyer and seller (and their agents).

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 they occur with low probability if the number of agents and houses is large. The problem can be avoided altogether if it is assumed that the number of agents is infinite, but that would entail analytical complications.⁵

3 Measures of Liquidity

The model just presented makes possible a general discussion of the meaning of liquidity and of several possible measures of liquidity. Optimal pricing implies that a house sells with probability less than one each period. This suggests that a natural measure of liquidity is the probability μ of sale during the current period or, equivalently, the expected time to sale, $(1 - \mu)/\mu$. This measure is appropriate for both buyer and seller.

Another measure of liquidity, appropriate for the seller, is the spread between the retail price of a house p and its wholesale price q ; this spread measures the capital gain a seller experiences when a house sells. Because this variable always equals zero for liquid assets (the value of a liquid asset to its owner just prior to sale equals its value when sold), it may also be a suitable measure of liquidity. One expects that the higher the expected time to sale, or the higher the retail-wholesale price spread, the greater the illiquidity.

⁵This 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. See Werner [9].

A third candidate measure of liquidity, appropriate for the buyer, is s , the value of the search option. The rationale for this measure is that, as noted in the introduction, illiquidity involves an element of irreversible cost when the asset is purchased. Here the irreversible cost arises from the fact that an agent who buys a house foregoes the opportunity to search again until the new match is broken. The value of this foregone search option, s , appears to be a natural measure of liquidity, with high values of s implying a high degree of illiquidity.

To investigate whether the interpretation of these variables is correct, we conducted a comparative statics experiment designed to vary liquidity. In our model houses are illiquid because buyers can evaluate only one house per period. The easiest way to vary liquidity is therefore to alter parameter values (the discount rate, for example) 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 proposed measures indicate high liquidity.

We first computed a benchmark equilibrium based on $\pi = 0.8$ and $\beta = .95$ (corresponding to an average occupancy duration of four 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 endogenous variables p , q and s do not require rescaling, but the expected time to sale was redefined to equal $(1 - \mu)/n\mu$, so as to measure in years rather than periods.

The accompanying graphs show the equilibrium values of p , q , μ and $(1 - \mu)/n\mu$ as a function of n , for selected values of n . For the most part the interpretation is as expected. 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.

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). Thus the higher n , the lower the vacancy rate.

When n is high, both p and q are high, but the spread between them is small (this is, of course, a direct consequence of eq. (6) with β replaced by β_n , since β_n converges to unity as n rises). Thus the first two measures of liquidity proposed above imply that when n is high, markets are liquid, as expected.

The behavior of s for different values of n is more interesting. First, note that s rises with n . However, the relevant measure of the value of search is surely s expressed as a proportion of the value of housing (p or q), and p and q also rise with n , as just observed.

As it happens, s/q equals one-half regardless of the value of n , so that measured liquidity does not depend on n . To understand why s/q equals one-half, note that the fact that eqs. (5) and (6) have the same form implies that

$$\frac{s}{q} = \frac{v((\epsilon^* + 1)/2) - p}{p}. \quad (10)$$

Thus showing that s/q equals one-half is equivalent to showing that the right-hand side of (10) equals one-half. But the fact that under linear utility the expected consumer surplus equals one-half the purchase price is well known.⁶

There are two possible interpretations of the fact that s/q does not rise with n . The first is that s/q is not a good measure of liquidity. However, rejecting s , however normalized, as a measure of liquidity raises questions about the characterization of liquidity offered in the introduction. There we suggested that illiquidity involves an element of irreversible cost; in our model s is the obvious measure of this cost. Thus rejecting s as a measure of illiquidity appears to imply rejecting the association between illiquidity and irreversible cost.

The other interpretation is that the idea of liquidity does not lend itself to a single unambiguous measure. It is true that in economies with high n buyers without a match will buy a house more quickly on average than in economies with low n , suggesting that liquidity is higher in the former case. However, it is also true that when n is high the number of searches that a buyer foregoes upon purchase of a house is higher. As observed above, it turns out that the value of these foregone searches rises

⁶In a wide variety of models incorporating risk neutrality, the expected surplus to the buyer of an optimally-priced commodity equals one-half the price paid. For example, consider a single-date model in which a seller owns an asset the value of which to a buyer is uniformly distributed on $[0, 1]$. This is the static counterpart of the dynamic model considered here. The seller knows that the buyer will buy if the value exceeds p , which will occur with probability $1 - p$ for any p . Therefore the optimal price is that which maximizes expected revenue $p(1 - p)$, or one-half. But then the expectation of the value of the object to the buyer conditional on purchase is three-quarters, implying that the expected surplus equals one-half the purchase price.

in proportional to the value of the house, suggesting that liquidity does not change with n .

Along these lines one would accept the fact that different aspects of liquidity are appropriately measured using different variables, and these variables cannot be expected to behave in the same way in different settings. One would give up the idea that a single comparative statics experiment, such as changing the length of the period, necessarily corresponds unambiguously to increasing or decreasing liquidity. Rather, both the variable used to measure liquidity and the experiment chosen to vary liquidity will depend on what aspect of liquidity one has in mind.

4 Are Returns Fair Games?

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 (recall eq. (1)). 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 return on a house offered for sale is

$$r = \begin{cases} p/\beta q & \text{with probability } \mu \\ 1 & \text{with probability } 1 - \mu \end{cases} . \quad (11)$$

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. (11). The conditional expected return is given by

$$E(r) = \mu p/\beta q + (1 - \mu). \quad (12)$$

Using eq. (6), eq. (12) simplifies to

$$E(r) = \beta^{-1}. \quad (13)$$

Thus the return equals investors' time preference.

Second, the return on an owner-occupied house is

$$r = \begin{cases} \epsilon/v(\epsilon) + 1 & \text{with probability } \pi \\ (\epsilon + q + s)/v(\epsilon) & \text{with probability } 1 - \pi \end{cases} . \quad (14)$$

Eq. (14) 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. (2), it follows that the expected rate of return on an owner-occupied house is also given by eq. (13).

Third, an agent without a match owns the search option, the current value of which is s . The return on the search option depends on η , the outcome of the search (which is not known at the beginning of the period). If $\eta \geq \epsilon^*$, so that the agent buys the house he is about to evaluate, the payoff consists of housing services η plus the expected next-period value of the house, $\pi v(\eta) + (1 - \pi)(q + s)$ less the next-period value of the purchase price, equal to $\beta^{-1}p$. If $\eta < \epsilon^*$ the agent does not buy the house, so the payoff is just s . We have

$$r = \begin{cases} (\eta + \pi v(\eta) + (1 - \pi)(q + s) - \beta^{-1}p)/s, & \eta \geq \epsilon^* \\ 1, & \eta < \epsilon^* \end{cases} . \quad (15)$$

Using eq. (2), eq. (15) can be simplified to

$$r = \begin{cases} (\beta^{-1}(v(\eta) - p))/s, & \eta \geq \epsilon^* \\ 1, & \eta < \epsilon^* \end{cases} . \quad (16)$$

Let \mathcal{F} be a partition of $[0, 1]$ consisting of the two events $[0, \epsilon^*)$, $[\epsilon^*, 1]$. Then the conditional expectation of r is given by

$$E(r|\mathcal{F}) = \begin{cases} (\beta^{-1}(v((\epsilon^* + 1)/2) - p))/s & \text{w. p. } \mu \\ 1 & \text{w. p. } 1 - \mu \end{cases} , \quad (17)$$

using the facts that $E(\epsilon|\eta \geq \epsilon^*) = (\epsilon^* + 1)/2$ and that v is linear. Taking the expectation of r and using eq. (2), eq. (13) results.

The excess returns $(r - \beta^{-1})$ just characterized are fair games: the expected excess returns conditional on the values of any or all of an agent's state variables are zero.

The argument just presented prepares the way for the analysis of returns in a more general version of the model, one in which there exist economy-wide shocks. Specifically, suppose that we modify the model so that housing services ϵ are replaced by $\epsilon + x$, where x is the economy-wide component of housing services. The current value of x is common to, and known by, all agents, but changes randomly over time. In equilibrium all the endogenous variables depend on x ; nevertheless, it turns out that the conditional expected return on owner-occupied houses, non-owner-occupied houses and the search option conditional on x is equal to β^{-1} , just as here⁷.

⁷See Krainer [5], where a model with aggregate risk is solved

It is well known that for liquid assets universal risk neutrality implies that conditional expected returns equal investors' rate of time preference. The result of this paper is that exactly the same thing is true for illiquid assets. Illiquidity affects asset prices, but it also affects the distribution of the equilibrium service flow on houses (since the equilibrium value of ϵ^* is endogenous) in such a way that expected returns are unaffected.

A large quantity of empirical evidence (for example, Case and Shiller [1], Meese and Wallace [7]) 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 it will not be easy to test directly the fair game proposition.

A more promising research strategy is to test the model by determining its predictions for return measures 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 return changes, if these are identified with price changes as in the empirical literature. However, as noted above 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].

5 Conclusion

The major conclusion of this paper is that excess returns are a fair game in illiquid markets, just as in liquid markets.

This conclusion, of course, depends critically on the assumption of risk neutrality. The present result suggests a conjecture about economies that incorporate both illiquidity and risk aversion. The conjecture is that the general equation of consumption-based asset pricing—that the expected excess return on any asset is proportional to the covariance of its payoff with the marginal utility of consumption—will apply even in the presence of illiquidity, as long as asset prices and payoffs are defined so as to reflect option values correctly.

explicitly assuming that x is a two-state Markov chain; the model is used to analyze such problems as how the various measures of liquidity depend on the aggregate state.

If this conjecture is correct, the implication is that incorporating illiquid assets in equilibrium models implies a much less radical departure from established analytical procedures than one might expect.

Appendix

We want to derive the policy functions associated with symmetric Nash equilibrium. The buyer's policy rule is

$$\theta = \operatorname{argmax}_z (zv(\eta) + (1-z)(\bar{p} + \beta s)), \quad (18)$$

where η is the buyer's fit with the house he is evaluating. Here $z \in \{0, 1\}$ is the buyer's decision variable, and θ is the expected-utility-maximizing value of z . The seller's policy rule is to offer for sale all houses that he owns but does not live in. The sale price is p , given by

$$p = \operatorname{argmax}_y (\mu(y)y + \beta(1 - \mu(y))q). \quad (19)$$

The endogenous variables p , q , ϵ^* , μ and s are determined endogenously by equations (4), (5), (6), (8) and (9), reproduced for convenience below:

$$v(\epsilon^*) = \bar{p} + \beta s, \quad (20)$$

$$s = \mu \left(v \left(\frac{\epsilon^* + 1}{2} \right) - \bar{p} \right) + \beta(1 - \mu)s, \quad (21)$$

$$q = \mu p + \beta(1 - \mu)q, \quad (22)$$

$$(\pi - \beta^{-1})(p - \beta q) + \mu = 0, \quad (23)$$

$$\mu = 1 - \epsilon^*. \quad (24)$$

Justification for the proposed policy rules (18) and (19) begins with the conjecture that the wealth function

$$\psi v(\epsilon) + (1 - \psi)s + hq \quad (25)$$

((1) above) is a value function in the dynamic programming sense. Then it is necessary (1) to show that the proposed policy rules implement the maximization in the Bellman equation, and (2) to verify that the conjectured form of the value function satisfies the required recursion.

The derivation to follow turns out to be more digestible if the cases $\psi = 1$ and $\psi = 0$ are handled separately. Thus we write the conjectured form of the value function as

$$J(\epsilon, h) = v(\epsilon) + hq \text{ if the agent has a match} \quad (26)$$

$$J(0, h) = s + hq \text{ if the agent has no match,} \quad (27)$$

where we substitute 0 for ϵ if there is no match.

Consider first the case where there is a match. The Bellman equation is

$$\begin{aligned} J(\epsilon, h) = & \max_y [h\mu(y)y + \beta\epsilon \\ & + \beta\pi J(\epsilon, (1 - \mu(y))h) \\ & + \beta(1 - \pi)J(0, (1 - \mu(y))h + 1)]. \end{aligned} \quad (28)$$

The first line is the agent's expected consumption from sales of housing, $h\mu(y)y$, and of housing services, $\beta\epsilon$. The second line is the probability π that the match persists multiplied by the discounted expected value function in that case. Here the second argument of J represents the expected next-period housing portfolio as one minus the expected rate of sale multiplied by the original portfolio. The third line is the product of the probability that the match fails and the discounted value function in that case; note that the expected housing portfolio is augmented by 1, reflecting the fact that the agent will want to sell the house.

Now insert the value functions (26) and (27) into (28). Collecting terms, there results

$$\begin{aligned} J(\epsilon, h) = & \max_y [h\mu(y)y + \beta(1 - \mu(y))hq] \\ & + \beta\epsilon + \beta\pi v(\epsilon) + \beta(1 - \pi)(q + s). \end{aligned} \quad (29)$$

The right-hand side of the first line equals hq , from (19) and (22), and the second line equals $v(\epsilon)$, from the definition

$$v(\epsilon) \equiv \beta\epsilon + \beta\pi v(\epsilon) + \beta(1 - \pi)(q + s) \quad (30)$$

of v . Therefore we have (26).

Now turn to the case where there is no match. The Bellman equation is

$$\begin{aligned} J(0, h) = & E_\eta [\max_{y,z} \{h\mu(y)y + z(\beta\eta - \bar{p}) \\ & + \beta z \{ \pi J(\eta, (1 - \mu(y))h) + (1 - \pi)J(0, (1 - \mu(y))h + 1) \} \\ & + \beta(1 - z)J(0, (1 - \mu(y))h) \}]. \end{aligned} \quad (31)$$

As before, the first term on the right-hand side of (31) indicates that the agent allocates the proceeds of housing sales to nonhousing consumption. If he chooses to buy the house he is evaluating ($z=1$) he consumes in addition its housing services η , but also decreases nonhousing consumption by an amount equal to the purchase price of the house.

The second line of (31) gives the expected next-period value functions if the agent decides to buy the house. Because of the possibility that the match just undertaken may fail going into the next period, there are two

possibilities: $J(\eta, 1 - \mu(y)h)$ if the match persists and $J(0, (1 - \mu(y))h + 1)$ if it fails. The third line gives the next-period expected value function if the agent decides not to buy the house ($z = 0$).

Inserting expressions (26) and (27) for the next-period value functions in the right-hand side of (31) and collecting terms, we get

$$\begin{aligned} J(0, h) = & E_\eta [\max_{y,z} \{h\mu(y)y + z(\beta\eta - \bar{p}) + z\beta\pi v(\eta) \\ & + (1 - \mu(y))hq(z\beta\pi + z\beta(1 - \pi) + (1 - z)\beta) \\ & + z\beta(1 - \pi)(q + s) + \beta(1 - z)s \}]. \end{aligned} \quad (32)$$

Observe that the second line above simplifies to $\beta(1 - \mu(y))hq$, so we have

$$\begin{aligned} J(0, h) = & E_\eta [\max_{y,z} \{h\mu(y)y + \beta(1 - \mu(y))hq \\ & + z(\beta\eta - \bar{p}) + z\beta\pi v(\eta) + z\beta(1 - \pi)(q + s) + \beta(1 - z)s \}]. \end{aligned} \quad (33)$$

Again, (19) and (22) may be used to replace $\max_y (\mu(y)y + \beta(1 - \mu(y))q)$ by q , which may be passed out of the expectation. Therefore we have

$$\begin{aligned} J(0, h) = & hq + E_\eta [\max_z \{ \\ & z(\beta\eta + \beta\pi v(\eta) + \beta(1 - \pi)(q + s)) \\ & - z\bar{p} + \beta(1 - z)s \}]. \end{aligned} \quad (34)$$

From (30), the second line is recognized as $zv(\eta)$, so (34) becomes

$$\begin{aligned} J(0, h) = & hq + E_\eta [\max_z \{z(v(\eta) - \bar{p}) + (1 - z)\beta s \}]. \quad (35) \\ = & hq + E_\eta [\max_z \{zv(\eta) + (1 - z)(\bar{p} + \beta s) - \bar{p} \}] \quad (36) \end{aligned}$$

after adding and subtracting \bar{p} .

The policy rule (18) implements the maximization, so (36) becomes

$$\begin{aligned} J(0, h) = & hq + E_\eta [\theta v(\eta) + (1 - \theta)(\bar{p} + \beta s) - \bar{p}] \quad (37) \\ = & hq + E_\eta [\theta(v(\eta) - \bar{p}) + (1 - \theta)\beta s]. \end{aligned} \quad (38)$$

The decision variable θ , now treated as a random variable, takes on value 1 with probability μ and value 0 with probability $1 - \mu$. Further, the expectation of η conditional on $\theta = 1$ is $(\epsilon^* + 1)/2$. Therefore we have

$$\begin{aligned} J(0, h) = & hq + \mu \left(v \left(\frac{\epsilon^* + 1}{2} \right) - \bar{p} \right) + \beta(1 - \mu)s \quad (39) \\ = & hq + s, \end{aligned} \quad (40)$$

from (21).

The last step is to evaluate $d\mu/dy$ in the first-order condition

$$\frac{d\mu}{dy}(y - \beta q) + \mu = 0 \quad (41)$$

for maximization of q . The seller knows that if he charges price y , the buyer with fit ϵ will make the purchase if $\epsilon \geq \epsilon^*(y)$, where $\epsilon^*(y)$ satisfies

$$v(\epsilon^*(y)) \equiv \frac{\beta\epsilon^*(y) + \beta(1 - \pi)(q + s)}{1 - \beta\pi} = y + \beta s. \quad (42)$$

Eq. (42) implies

$$\frac{d\epsilon^*(y)}{dy} = \beta^{-1} - \pi. \quad (43)$$

Finally, (24) implies that

$$\frac{d\mu}{dy} = -\frac{d\epsilon^*}{dy}, \quad (44)$$

so there results

$$\frac{d\mu}{dy} = \pi - \beta^{-1}, \quad (45)$$

leading to (23).

References

- [1] B. 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] Larry Glosten and Paul Milgrom. Bid, ask and transaction prices in a specialist model with heterogeneously informed traders. *Journal of Financial Economics*, 14, 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] 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.
- [8] Jeremy C. Stein. Prices and trading volume in the housing market: A model with down-payment effects. *Quarterly Journal of Economics*, 110:379–406, 1995.
- [9] Jan Werner. Diversification and equilibrium in securities markets. forthcoming, *Journal of Economic Theory*, 1994.
- [10] William Wheaton. Vacancy, search and prices in a housing market matching model. *Journal of Political Economy*, 61:1270–1292, 1990.
- [11] Joseph T. Williams. Pricing real assets with costly search. *Review of Financial Studies*, 8:55–90, 1995.

Figure 1

Prices as function of buyer arrival rate




Figure 2

