

Valuing Thinly Traded Assets

Francis A. Longstaff^{a,b}

^aUCLA Anderson School of Management, University of California, Los Angeles, Los Angeles, California 90095; ^bNational Bureau of Economic Research, Cambridge, Massachusetts 02138

Contact: francis.longstaff@anderson.ucla.edu (FAL)

Received: October 12, 2015

Accepted: November 10, 2016

Published Online in Articles in Advance:
May 24, 2017

<https://doi.org/10.1287/mnsc.2016.2718>

Copyright: © 2017 INFORMS

Abstract. We model illiquidity as a restriction on the stopping rules investors can follow in selling assets, and apply this framework to the valuation of thinly traded investments. We find that discounts for illiquidity can be surprisingly large, approaching 30%–50% in some cases. Immediacy plays a unique role and is valued much more than ongoing liquidity. We show that investors in illiquid enterprises have strong incentives to increase dividends and other cash payouts, thereby introducing potential agency conflicts. We also find that illiquidity and volatility are fundamentally entangled in their effects on asset prices. This aspect may help explain why some assets are viewed as inherently more liquid than others and why liquidity concerns are heightened during financial crises.

History: Accepted by Gustavo Manso, finance.

Keywords: asset pricing • liquidity • finance • investment

1. Introduction

Thinly traded assets are often defined as investments for which there is no liquid market available. Thus, investors holding illiquid or thinly traded assets may not be able to sell their positions for extended periods—if ever. At best, investors may only be able to sell in infrequent privately negotiated transactions. The economics of these private transactions, however, are complicated since prospective buyers realize that they will inherit the same problem when they later want to resell the assets. Not surprisingly, sales of thinly traded assets typically occur at prices far lower than would be the case if there were a liquid public market.

The valuation of thinly traded assets is one of the most important unresolved issues in asset pricing and has many fundamental implications for individuals, firms, markets, and policymakers. One reason for this is that thinly traded assets collectively represent a large fraction of the aggregate wealth in the economy. Key examples where investors may face long delays before being able to liquidate holdings include the following:

- Sole proprietorships.
- Partnerships, limited partnerships.
- Private equity and venture capital.
- Life insurance and annuities.
- Pensions and retirement assets.
- Residential and commercial real estate.
- Private placements of debt and equity.
- Distressed assets and fire sales.
- Compensation in the form of restricted options and shares.
- Investments in education and human capital.

Other examples include transactions that take public firms private such as leveraged buyouts (LBOs) that

result in residual equity holders having much less liquid positions. Many hedge funds have lockup provisions that prohibit investors from withdrawing their capital for months or even years. Investors in initial public offerings (IPOs) are often allocated shares with restrictions on reselling or “flipping” the shares.

Many insightful approaches have been used in the asset pricing literature to study the effects of illiquidity on security prices, and these are briefly described in the literature review section below. This paper approaches the challenge of valuing illiquid assets from a new perspective. We view illiquidity as a restriction on the set of stopping rules that an investor is allowed to follow in selling the asset. This approach allows us to use an option-theoretic framework to place realistic bounds on the values of securities that cannot be traded continuously.

To illustrate the intuition behind our approach, let Ω denote the set of all stopping rules available to an investor when the asset is fully liquid. Let $A \subset \Omega$ denote the restricted subset of stopping rules available because of the illiquidity of the asset. Finally, let $B \subset A$ be a subset of A that contains only one element—the worst-case stopping rule in A . Our approach consists of finding an upper bound on the amount that would be required to fully compensate an investor for the welfare loss from restricting his stopping rules to subset B . Using a simple dominance argument, however, it is clear that this amount also represents an upper bound on what would be required to fully compensate an investor who was limited to the less-restricted subset A . Thus, by finding a bound for the worst-case scenario, our results also apply for less-severe forms of illiquidity encountered in actual markets.

This paper contributes to an extensive literature that uses dominance arguments to provide asset pricing bounds. Key examples include the no-arbitrage option-pricing bounds derived by Merton (1973) and the bounds on the moments of the stochastic discount factor obtained by Hansen and Jagannathan (1991). There are many reasons why having a lower bound on the value of an illiquid asset could be useful. The lower bound could serve as a reservation price in negotiations between sellers and prospective buyers and provide a benchmark in litigation and dispute resolution. The lower bound could also provide guidance to regulators and policymakers in making regulatory capital decisions or establishing limits on the collateral value of illiquid assets used to secure debt financing or in margin accounts.¹

The results provide a number of important insights into the potential effects of illiquidity on asset values. First, we show that the value of immediacy in financial markets is much higher than the value of future liquidity. For example, the discount for illiquidity for the first day of illiquidity is 2.4 times that for the second day, 4.2 times that for the fifth day, 6.2 times that for the tenth day, and 20.0 times that for the 100th day. These results suggest that immediacy is viewed as fundamentally different in its nature. This dramatic time asymmetry in the value of liquidity may also help explain the rapidly growing trend toward electronic execution and high-frequency trading in many financial markets.

Second, our results confirm that the values of illiquid assets can be heavily discounted in the market. We show that investors could discount the value of illiquid assets by as much as 10%, 20%, or 30% for illiquidity horizons of one, two, or five years, respectively. Although our results only provide lower bounds on the values of illiquid assets, the evidence in the empirical literature suggests that these bounds may be realistic approximations of the prices at which various types of thinly traded securities are sold in privately negotiated transactions. For example, Silber (1991) documents that restricted stocks—stocks that investors cannot trade for two years after they are acquired—are placed privately at an average discount of 34% relative to fully liquid shares, and many are placed at discounts in excess of 50%. Berkman and Eleswarapu (1998) find that when an exchange rule allowing forward trading was abolished, the decreased liquidity of the affected shares resulted in prices declining by 15%. Fleckenstein et al. (2014) show that portfolios of Treasury inflation protected securities have traded at discounts of more than 23% relative to portfolios of more-liquid Treasury bonds with identical cash flows. Aragon (2007) finds that hedge funds with lockups have annual returns that average 4%–7% more than hedge funds without lockups. Brenner et al. (2001) find

that thinly traded currency options with three- to six-month maturities are placed privately at roughly a 20% discount to fully liquid options. Chen and Xiong (2001) show that restricted institutional shares in China trade at average discounts of 78%–86% relative to otherwise identical common shares. There are many other similar examples in the empirical literature (for example, see the review article by Amihud et al. 2005).

Third, we find that the effects of illiquidity and volatility on asset prices are fundamentally entangled. Specifically, asset return variances and the degree of asset illiquidity are indistinguishable in their effects on discounts for illiquidity. This makes intuitive sense since investors are more likely to want to sell assets when prices have diverged significantly from their original purchase prices. This divergence, however, can arise both through the passage of time as well as through the volatility of asset prices. Because of this, assets with stable prices such as cash or short-term Treasury bills can be viewed as inherently more liquid than assets such as stocks even when all are readily tradable. This may also help explain why concerns about market liquidity become much more central during financial crises and periods of market stress.

Finally, the results indicate that the effect of illiquidity on asset prices is smaller for investments with higher dividends or cash payouts. An important implication of this is that investors in illiquid assets such as private equity, venture capital, leveraged buyouts, etc. have strong economic incentives to increase payouts. Thus, illiquidity may have the potential to be a fundamental driver of both dividend policy and capital structure decisions for privately held ventures or thinly traded firms.

2. Literature Review

The literature on the effects of illiquidity on asset valuation is far too extensive for us to be able to review in detail. Instead, we will simply summarize some of the key themes that have been discussed in this literature. For an in-depth survey of this literature, see the excellent review by Amihud et al. (2005) on liquidity and asset prices.

Many important papers in this literature focus on the role played by transaction costs and other financial frictions in determining security prices. Amihud and Mendelson (1986) present a model in which risk-neutral investors consider the effect of future transaction costs in determining current valuations for assets. Constantinides (1986) shows that while transaction costs can have a large effect on trading volume, investors optimally trade in a way that mitigates the effect of transaction costs on prices. Vayanos (1998) and Vayanos and Vila (1999) show that transaction costs can increase the value of liquid assets, but can have an ambiguous effect on the value of illiquid assets.

A number of recent papers recognize that liquidity is time varying and develop models in which liquidity risk is priced into asset valuations. Pastor and Stambaugh (2003) consider a model in which marketwide systemic liquidity risk is priced. Acharya and Pedersen (2005) show how time-varying liquidity risk affects current security prices and future expected returns. Gromb and Vayanos (2002) and Brunnermeier and Pedersen (2009) develop models in which changes in the abilities of dealers to fund their inventories translates into variation in the liquidity they can provide, which in turn results in a liquidity risk premium being embedded into asset values.

Another recent theme in the literature addresses the effects of search costs or the cost of being present in the market on liquidity and asset prices. Duffie et al. (2005, 2007), Vayanos and Wang (2007), Vayanos and Weill (2008), Duffie et al. (2009), and others consider models in which agents incur costs as they search for other investors willing to trade with them, and show how these costs affect security prices. Huang and Wang (2009, 2010) study asset pricing in a market where it is costly for dealers to be continuously present in the market and provide liquidity.

A number of papers in the literature view illiquidity from the perspective of a limitation on the ability of an agent to trade continuously. Lippman and McCall (1986) define liquidity in terms of the expected time to execute trading strategies. Longstaff (1995, 2001), Kahl et al. (2003), and Finnerty (2012) study the implications of trading restrictions on assets. Longstaff (2009) presents a general equilibrium asset pricing model in which agents must hold asset positions for a fixed horizon rather than being able to trade continuously.

Finally, several papers approach the valuation of liquidity from an option-theoretic perspective. In Copeland and Galai (1983), dealers take inventory risk by providing bid and ask quotes over some horizon in markets where investors may have private information. They show that the bid-ask spread compensating the dealers for bearing this risk can be modeled as an option straddle. Similarly, Chacko et al. (2008) value immediacy by modeling limit orders as American options. Ang and Bollen (2010) model the option to withdraw funds from a hedge fund as a real option. Ghaidarov (2014) models the option to sell equity securities as a forward-starting put option.

3. Modeling Illiquidity

In this section, we present a new approach to modeling thin trading or illiquidity in financial markets. This approach provides a simple framework that can be used to place lower bounds on the values of illiquid assets. Note that placing a lower bound on the value of the illiquid asset is equivalent to placing an upper bound on the size of the discount for illiquidity. For

clarity, we will generally couch the discussion in terms of the discount for illiquidity.

The concept of a stopping rule plays a central role in how we model illiquidity. Intuitively, a stopping rule can be viewed as a decision rule that determines the (potentially random) stopping time τ when the asset is to be sold, where τ depends only on information available in the market up to and including time τ . For example, a decision rule to sell the asset at a prespecified date T is a stopping rule. A decision rule to sell the asset via a limit order that is executed the first time the asset price reaches a value of, say, 50 is a stopping rule. In contrast, a decision rule to sell the asset when its price reaches its maximum value between time zero and time T is not a stopping rule since the time at which the maximum is attained is not known for certain prior to time T .²

The key insight underlying this modeling framework is that illiquidity can be viewed as a restriction on the set of stopping rules that an investor can follow in selling the asset. In particular, an investor that purchases a liquid asset can follow any stopping rule he chooses in selling the asset. In contrast, an investor that purchases an illiquid asset is restricted to a subset of stopping rules. If the investor's preferred stopping rule is not included in the subset, then the investor must choose a stopping rule that is suboptimal from his perspective. In this case, the investor suffers a welfare loss and may only be willing to purchase the illiquid asset at a discount relative to what he would be willing to pay for the fully liquid asset.

Specifically, let T denote the horizon over which an investor faces illiquidity constraints on his holdings of an asset.³ Let X_T denote the value of the investor's position at time T if the investor were able to follow his preferred stopping rule in selling the asset and then reinvesting the proceeds in the riskless asset. Similarly, let Y_T denote the value of the investor's position at time T by following the best stopping rule allowed him by the illiquidity of the asset and then reinvesting the proceeds in the riskless asset. Clearly, if an investor has preferences over stopping rules, then these two outcomes are not equivalent and the investor may be unwilling to pay as much for the illiquid asset.

Viewing illiquidity from this perspective suggests a very intuitive framework for placing bounds on the discount for illiquidity. Recall that ex ante, the investor would prefer to receive X_T at time T , but will only receive Y_T because of the illiquidity of the asset. However, if the investor were to be given an option that allowed him to exchange X_T for Y_T at time T (known as an exchange option), then the investor would be made completely whole on an ex post basis. In particular, an investor with a portfolio consisting of the illiquid asset and an exchange option with cash flow $\max(0, X_T - Y_T)$ at time T would end up with $Y_T + \max(0, X_T - Y_T) =$

$\max(X_T, Y_T)$. This cash flow, however, is greater than or equal to the cash flow X_T that the investor would have received had he purchased the liquid asset instead of the illiquid asset. A simple dominance argument implies that the investor would prefer the portfolio of the illiquid asset and the exchange option to owning the liquid asset. In turn, this implies that the sum of the values of the illiquid asset and the exchange option should be greater than or equal to that of the liquid asset, or alternatively, that the value of the exchange option represents an upper bound on the discount for illiquidity.

Recall that we designated Y_T as the value of the investor's position by following the best stopping rule from the subset A of stopping rules permitted by the illiquidity of the asset. It is important to observe, however, that the above analysis holds even when the stopping rule is chosen from a more restricted subset of A , which we designate as B . This follows since the investor still prefers $\max(X_T, Y_T)$ to X_T , where Y_T is now the value of the investor's portfolio following the best stopping rule in B . Thus, the value of the exchange option with payoff $\max(0, X_T - Y_T)$ also represents an upper bound on the discount for illiquidity. Furthermore, this is true for any subset $B \subset A$. We will use this result in the next section since by limiting B to the worst-case stopping rule (buy and hold), we can easily identify Y_T and solve for the value of the exchange option. Intuitively, this simply means that by compensating the investor for the worst-case illiquidity scenario, we are also fully compensating the investor for the actual illiquidity scenario he faces. Thus, the upper bound we derive is applicable even when the illiquidity of the asset does not limit the investor to following a buy-and-hold strategy.⁴

4. The Discount for Illiquidity

As discussed above, the task of finding the upper bound on the discount for illiquidity can be reduced to solving for the value of the exchange option. To do this, we first need to specify a valuation framework for the exchange option.

As the valuation framework for the exchange option, we adopt the familiar Black and Scholes (1973) option-pricing setting. Let S_t be the price per share of an asset, where the share is fully liquid and can be traded continuously in the financial markets without frictions. We assume that the dynamics of S_t are given by the following geometric Brownian motion process under the risk-neutral pricing measure,

$$dS = rS dt + \sigma S dZ, \quad (1)$$

where r is the constant riskless rate, σ is the volatility of continuously compounded returns, and dZ is the

increment of a standard Brownian motion. For simplicity, we assume for the present that the asset does not pay any dividends or cash flows before time T . This assumption, however, will be relaxed later.

As described above, we will solve for value of the exchange option under the worst-case illiquidity scenario. The resulting value of the exchange option will clearly provide an upper bound for the amount required to compensate investors who acquire assets with less-severe liquidity restrictions. Specifically, the worst-case scenario is that once the illiquid asset is purchased at time zero, it cannot be sold again until time T . Thus, the illiquid asset is completely nonmarketable from time zero to time T . In this worst-case scenario, an investor who buys the illiquid asset at time zero has only one stopping rule available—selling the asset at time T . As a result, the cash flow received at time T from following this stopping rule is simply $Y_T = S_T$.

Next, we need to identify X_T . To do so, we need several preliminary results. First, let τ , $0 \leq \tau \leq T$, denote the time at which the stopping rule chosen by the investor results in the liquid asset being sold. The cash flow received by the investor from selling the liquid asset at time τ and reinvesting the proceeds in the riskless asset is $X_T = S_\tau e^{r(T-\tau)}$.

Second, substituting these expressions for X_T and Y_T into the expression for the payoff from an exchange option implies that the cash flow at time T from the exchange option is given by

$$\max(0, S_\tau e^{r(T-\tau)} - S_T). \quad (2)$$

As shown, this cash flow depends on the asset price at both the stopping time τ and the final date T . An important implication of this is that once the stopping time τ is reached, the value of S_τ is known and is no longer stochastic. This means that as of time τ , the exchange option can be viewed as a simple put option on the asset value with a fixed strike price of $S_\tau e^{r(T-\tau)}$. Thus, as shown in the appendix (which provides the derivation for this and all other results in the paper), the value of the option at time τ is given by substituting in the current stock price S_τ and the strike price $S_\tau e^{r(T-\tau)}$ into the Black and Scholes (1973) formula for puts,

$$S_\tau [N(\sqrt{\sigma^2(T-\tau)}/2) - N(-\sqrt{\sigma^2(T-\tau)}/2)], \quad (3)$$

where $N(\cdot)$ is the standard cumulative normal distribution function. This expression for the value of the option as of the stopping time τ is true for any stopping rule.

Third, to solve for the initial or time-zero value of the exchange option, we take the present value of receiving a cash flow at time τ equal to the value of the put given in Equation (3),

$$E[e^{-r\tau} S_\tau [N(\sqrt{\sigma^2(T-\tau)}/2) - N(-\sqrt{\sigma^2(T-\tau)}/2)]], \quad (4)$$

where the expectation is taken with respect to the joint distribution of S_τ and the stopping time τ .

For any choice of τ , the value of the exchange option in Equation (4) provides an upper bound on the discount for illiquidity. Rather than attempting to solve for the value of the exchange option for a specific stopping rule, our approach will simply be to solve for the maximum value of the exchange option over all possible stopping rules. Clearly, this approach will result in a value for the exchange option that dominates the value of the exchange option resulting from any other stopping rule that might be chosen by the investor. Intuitively, the maximized value of the exchange option can be viewed as an “upper bound on an upper bound.”

Given this structure, we can now determine the maximized value of the exchange option. This can be obtained by solving for the stopping rule that maximizes the expression given above,

$$\max_{\tau} E[e^{-r\tau} S_{\tau} [N(\sqrt{\sigma^2(T-\tau)}/2) - N(-\sqrt{\sigma^2(T-\tau)}/2)]] \quad (5)$$

As shown in the appendix, the stopping rule that maximizes the value of the exchange option has a surprisingly simple form—the maximizing stopping rule is simply to sell the liquid asset immediately at time zero. Thus, $\tau = 0$. The intuition for this result is easily understood. By compensating an investor for illiquidity in a way that allows them to attain the maximum of X_T and Y_T , the investor has a strong incentive to ensure that X_T and Y_T are as different as possible. By stopping at time zero, the exchange option allows the investor to choose between payoffs linked to the most temporally divergent values of the asset price possible: S_0 and S_T .

Finally, to obtain the maximized value of the exchange option, we substitute the maximizing stopping rule $\tau = 0$ into Equation (4). It is easily shown that the resulting value for the exchange option and upper bound on the discount for illiquidity is given by

$$S_0 [N(\sqrt{\sigma^2 T}/2) - N(-\sqrt{\sigma^2 T}/2)]. \quad (6)$$

This closed-form solution for the value of the exchange option has a very simple structure. In particular, the value of the exchange option is an explicit function of both the length of the illiquidity horizon T and the volatility of the liquid asset as measured by σ .

The discount for illiquidity is easily shown to be an increasing function of both the illiquidity horizon T and the volatility parameter σ . These comparative statistics results are intuitive since an increase in T restricts the stopping rules available to the investor further, while an increase in σ increases the opportunity cost of not being able to trade. The dependence of the discount on volatility and the horizon T parallels the results in Copeland and Galai (1983) and others who use option-theoretic frameworks to model illiquidity.

Given the closed-form solution for the exchange option, the lower bound on the value of the illiquid asset is given by subtracting the value of the exchange option from the value of an equivalent liquid asset.

5. Discussion

These results for the lower bound on the value of illiquid or thinly traded securities have many interesting implications. To illustrate, Table 1 reports the lower bounds for illiquidity horizons ranging from one day to 30 years, and for volatilities ranging from 10% to 50%.

Table 1 shows that illiquidity can have a dramatic effect on asset values. In particular, the price of an investment could be discounted by as much as 20%–40% for illiquidity horizons ranging from two to five years. Furthermore, asset prices could be discounted by more than 50% for illiquidity horizons of 10 years or longer. Although our results provide only lower bounds on the values of illiquid or thinly traded assets, these lower bounds are actually consistent with empirical evidence about discounts for illiquidity.

The effects of illiquidity can also be substantial even for relatively short horizons. Table 1 shows that a one-day illiquidity horizon implies a lower bound on the value of an illiquid asset ranging from 99.75% to 98.74% of the value of the liquid asset. Similarly, for a one-week horizon, the lower bound ranges from 99.45% to 97.23% of the value of the liquid asset.

These results imply that the discounts for illiquidity horizons measured in days are surprisingly large. For example, an investor who bought an illiquid asset at a discount of 1%, but was then able to sell a day later at the fully liquid price would realize a huge annualized rate of return on the transaction. This suggests that the value of immediacy (the ability to sell immediately) could represent one of the largest types of risk premia in financial markets.

To explore this, we compute the annualized discounts for illiquidity by dividing the discounts implied

Table 1. Percentage Lower Bounds for Illiquid Asset Values

Illiquidity horizon	Volatility				
	10%	20%	30%	40%	50%
1 Day	99.748	99.495	99.243	98.991	98.739
1 Week	99.447	98.894	98.340	97.787	97.234
1 Month	98.848	97.697	96.546	95.396	94.247
1 Year	96.012	92.034	88.076	84.148	80.259
2 Years	94.363	88.754	83.200	77.730	72.367
5 Years	91.098	82.306	73.732	65.472	57.615
10 Years	87.437	75.183	63.526	52.709	42.920
20 Years	82.306	65.472	50.233	37.109	26.355
30 Years	78.419	58.388	41.131	27.332	17.090

Notes. This table reports the lower bound on the value of an illiquid asset expressed as a percentage of the price of an equivalent liquid asset. Volatility denotes the volatility of returns for an equivalent liquid asset.

Table 2. Annualized Percentage Discounts for Illiquidity

Illiquidity horizon	Volatility				
	10%	20%	30%	40%	50%
1 Day	63.075	126.150	189.225	252.300	315.375
1 Week	28.766	57.533	86.299	115.060	143.811
1 Month	13.819	27.636	41.447	55.248	69.038
1 Year	3.988	7.966	11.924	15.852	19.741
2 Years	2.819	5.623	8.399	11.135	13.816
5 Years	1.780	3.539	5.254	6.906	8.477
10 Years	1.256	2.482	3.647	4.729	5.708
20 Years	0.885	1.726	2.488	3.145	3.682
30 Years	0.719	1.387	1.962	2.422	2.764

Notes. This table reports the annualized percentage discount for illiquidity where this value is computed as the ratio of the percentage discount for illiquidity divided by the length of the illiquidity horizon measured in years. Discounts for illiquidity are expressed as a fraction of the value of an equivalent liquid asset. Volatility denotes the annualized volatility of returns for an equivalent liquid asset.

by the results in Table 1 by the illiquidity horizons. These annualized discounts are reported in Table 2. As shown, the annualized discounts for short horizons such as a day or a week can be orders of magnitude greater than those for longer horizons. For example, the annualized discount for illiquidity for a one-day horizon is approximately 16 times as large as that for a one-year horizon.

As an alternative way of viewing these results, Table 3 reports the marginal discount of illiquidity as the illiquidity horizon ranges from one to 20 days. Specifically, we report the discount for a one-day horizon, the marginal or incremental increase in the discount as the horizon is increased to two days from one day, the marginal or incremental increase in the discount as the horizon is increased to three days from two days, and so forth.

As illustrated in Table 3, the discount for illiquidity for the first day is much larger than for the second, third, etc. days. In particular, the discount for the first day of illiquidity is 2.41 times that for the second day, 3.15 times that for the third day, 8.83 times that for the 20th day, and 32.33 times the discount for the day one year later. Clearly, liquidity today is worth much more than liquidity tomorrow.

These results are consistent with the literature on the value of immediacy. For example, Demsetz (1968) defines immediacy as the price concession that would be needed to transact immediately. Our results indicate that this price concession could be relatively large, particularly for assets with higher return volatilities such as stocks. Other authors who focus on the valuation of immediacy include Stoll (2000), who develops a regression-based model of the price of immediacy, and Grossman and Miller (1988), who model market liquidity as being determined by the supply and demand for immediacy.

Table 3. Marginal Percentage Discounts for Illiquidity

Day	Volatility				
	10%	20%	30%	40%	50%
1	0.252	0.505	0.757	1.009	1.262
2	0.105	0.209	0.314	0.418	0.522
3	0.080	0.160	0.241	0.321	0.401
4	0.068	0.135	0.203	0.270	0.338
5	0.060	0.119	0.179	0.238	0.298
6	0.054	0.108	0.162	0.215	0.269
7	0.050	0.099	0.149	0.198	0.247
8	0.046	0.092	0.138	0.184	0.230
9	0.043	0.087	0.130	0.173	0.216
10	0.041	0.082	0.123	0.164	0.204
11	0.039	0.078	0.117	0.156	0.194
12	0.037	0.074	0.112	0.149	0.186
13	0.036	0.071	0.107	0.143	0.178
14	0.034	0.069	0.103	0.137	0.171
15	0.033	0.066	0.099	0.132	0.165
16	0.032	0.064	0.096	0.128	0.160
17	0.031	0.062	0.093	0.124	0.155
18	0.030	0.060	0.090	0.120	0.150
19	0.029	0.059	0.088	0.117	0.146
20	0.029	0.057	0.086	0.114	0.143

Notes. This table reports the marginal or incremental change in the discount for illiquidity for horizons ranging from 1 to 20 days. Volatility denotes the annualized volatility of returns for an equivalent liquid asset.

The lower bounds in Table 1 also illustrate an interesting symmetry between the length of the illiquidity horizon and the volatility of the asset. In particular, doubling the volatility has essentially the same effect on the lower bound as quadrupling the length of the illiquidity horizon. The reason for this symmetry is easily seen from the expression for the lower bound in Equation (6). As shown, the lower bound depends on volatility and length of the illiquidity horizon only through the product $\sigma^2 T$.

From an intuitive perspective, this means that neither volatility nor the length of the illiquidity horizon are fundamental in determining the lower bound. Rather, it is the total realized variance $\sigma^2 T$ before the asset can be traded again that matters. This implies that volatility and the timing of illiquidity are fundamentally entangled in the sense that their effects are indistinguishable from each other.

This notion of entanglement may also help explain why some assets such as Treasury bills are viewed as inherently more liquid than stocks even when orders for either can be executed within seconds. Since stocks have higher volatility, their lower bounds will always be smaller than is the case for Treasury bills even when both are tradable at the same frequency. Liquidity is not simply a function of market microstructure. Rather, it depends also on the inherent riskiness of the underlying asset. These considerations may help explain why concerns about liquidity become

particularly acute during volatile high-stress periods in the financial markets.

6. Extension to Dividends

In this section, we extend the analysis to the situation in which the asset pays dividends, coupons, or other cash payouts over time. This situation differs from the earlier case in that when an investor sells the liquid asset, the investor no longer receives the stream of dividends. In contrast, the holder of an illiquid asset continues to receive dividends until the illiquidity horizon is reached. For symmetry, we will assume that dividends are reinvested in the riskless asset as they are received.

Given this structure, the value X_T of the investor's position at time T from following the optimal stopping rule is $S_\tau e^{r(T-\tau)} + \int_0^\tau \rho_s e^{r(T-s)} ds$ where ρ_s is the dividend (assumed continuous). The value Y_T of the investor's position at time T from following the restricted stopping rule is $S_T + \int_0^T \rho_s e^{r(T-s)} ds$. As before, the upper bound on the discount for illiquidity is given by the value of the exchange option with cash flow at time T of $\max(0, X_T - Y_T)$.

Despite the introduction of dividends into the framework, the appendix shows that the stopping strategy that maximizes the value of this exchange option is identical. Specifically, the maximizing stopping rule is to sell the liquid asset immediately at time zero, $\tau = 0$. The intuition for this result is the same as before; the value of the exchange option is maximized when the value of X_T and Y_T are as temporally divergent as possible.

The specific functional form of the exchange option will clearly depend on the nature of the dividend stream ρ_t . To provide some examples of the effect of dividends on the discount for illiquidity, we will make the standard assumption that the underlying asset has a constant dividend yield. Specifically, we assume that the dividend is ρS_t , where ρ is a constant. Given this assumption, the asset price dynamics in Equation (1) imply that dividends are random and conditionally log normally distributed. Sums of log normals, however, are not log normal, which implies that the exchange option does not have a simple closed-form solution. Accordingly, we will solve for the value of the exchange option via straightforward simulation. Table 4 presents lower bounds for the value of the illiquid asset for dividend yields ranging from 0 to 8%, and where volatility is held fixed at 30%.

As shown, dividends can have a major effect on the discount for illiquidity, particularly for longer horizons. Furthermore, the discount for illiquidity decreases as the dividend yield increases. This result is intuitive since by receiving dividends, an investor in an illiquid asset is able to convert some of his position

Table 4. Percentage Lower Bounds for Illiquid Asset Values When the Asset Pays a Continuous Proportional Dividend

Illiquidity horizon	Dividend yield				
	0%	2%	4%	6%	8%
1 Day	99.243	99.243	99.243	99.243	99.243
1 Week	98.340	98.340	98.340	98.341	98.341
1 Month	96.546	96.549	96.552	96.555	96.558
1 Year	88.076	88.195	88.311	88.426	88.538
2 Years	83.200	83.527	83.844	84.151	84.446
5 Years	73.732	74.976	76.119	77.170	78.139
10 Years	63.526	66.875	69.696	72.080	74.108
20 Years	50.233	58.523	64.351	68.584	71.764
30 Years	41.131	54.567	62.659	67.927	71.604
100 Years	13.361	49.384	61.161	67.393	71.380
1,000 Years	0.001	49.222	61.075	67.375	71.331

Notes. This table reports the percentage lower bounds on the value of an illiquid asset where the asset pays a continuous dividend at the indicated dividend yield. The lower bound is expressed as a percentage of the price of an equivalent liquid asset. Asset return volatility is fixed at 30%.

into cash sooner than if the asset did not pay dividends. In essence, by paying dividends or other cash flows, an illiquid asset partially liquidates itself. Thus, the illiquidity constraint is relaxed to some extent by the payment of dividends.

These results have many important implications for illiquid investments such as partnerships, private equity, venture capital, closely held firms, etc. Specifically, these results suggest that investors in these types of assets have strong incentives to accelerate the payment of distributions, dividends, and other cash flows to reduce the impact of illiquidity on their holdings.

7. Robustness

For simplicity, we have focused on the case where an investor faces illiquidity constraints until a fixed time T . For many situations, this characterization of illiquidity is very realistic. For example, most of the examples given in the introduction such as retirement and pension assets, life insurance and annuities, private equity and venture capital, restricted options and shares, fixed income investments, etc. clearly have finite horizons. On the other hand, however, there are clearly other forms of illiquidity that may not have a fixed horizon. Thus, it is important to consider how robust the results are to this simplifying assumption.

As one way of addressing this issue, we examine how the results change when the length of the illiquidity horizon increases unboundedly. This can be seen in Table 4, which reports the lower bounds for illiquid asset values for values of T as large as 1,000 years. As shown, with the exception of the zero-dividend-yield case, the lower bounds for 100- and 1,000-year horizons are very similar to those for much shorter horizons. For example, when the dividend yield is 4%, the

lower bound is 62.659% for $T = 30$, 61.161% for $T = 100$, and 61.075% for $T = 1,000$. These results suggest that our modeling approach can still provide useful bounds even in cases where the assets may be illiquid for arbitrarily long horizons—that the results are robust to large changes in the value of T .

As an alternative way of modeling illiquidity, we extend the basic framework to the situation where the asset is illiquid until time T , but then can only be sold at a discount (reflecting ongoing illiquidity beyond time T). Let γS_T denote the price at which the asset can be sold at time T , where $0 < \gamma \leq 1$. Following the same approach as earlier, it is easily shown that the value of the exchange option compensating the investor for the illiquidity of the asset is now given by

$$S_0[N(-a_2) - \gamma N(-a_1)], \quad (7)$$

where

$$a_1 = \ln \gamma / \sqrt{\sigma^2 T} + \sqrt{\sigma^2 T} / 2, \quad (8)$$

$$a_2 = a_1 - \sqrt{\sigma^2 T}. \quad (9)$$

These results raise an intriguing question: Is there a recursive or “fixed point” value of γ such that if the asset could only be sold at a discounted price of γS_T at time T , then the lower bound at time zero would have a similarly discounted value of γS_0 ? Strictly speaking, the answer is no since it is easily shown that γS_0 represents an upper bound. On the other hand, it turns out that γ can be chosen so that the lower bound at time zero is arbitrarily close to γS_0 . This is illustrated in Table 5, which shows the values of γ such that the discount at time zero equals the discount at time T , within an accuracy of 0.10%. Thus, for all practical purposes, a recursive solution of this nature is feasible.⁵ These results illustrate that the basic modeling framework can easily be extended to allow for less than perfect liquidity at time T .

Table 5. Recursive Lower Bounds for Illiquid Asset Values

Illiquidity horizon	Volatility				
	10%	20%	30%	40%	50%
1 Day	0.9959	0.9871	0.9770	0.9662	0.9549
1 Week	0.9852	0.9619	0.9367	0.9106	0.8844
1 Month	0.9599	0.9064	0.8518	0.7984	0.7470
1 Year	0.8268	0.6550	0.5146	0.4033	0.3164
2 Years	0.7521	0.5365	0.3802	0.2700	0.1930
5 Years	0.6190	0.3596	0.2102	0.1252	0.0765
10 Years	0.4946	0.2302	0.1108	0.0561	0.0301
20 Years	0.3596	0.1252	0.0480	0.0206	0.0100
30 Years	0.2820	0.0803	0.0268	0.0107	0.0051

Notes. This table reports the value of γ such that if the asset can be sold at γ times the fully liquid value at the end of the illiquidity horizon, then the lower bound at time zero is γ times the current fully liquid value (within an accuracy of 0.10%). Volatility denotes the volatility of returns for an equivalent liquid asset.

Finally, in the next section, we show that the model can also be generalized to the case where the horizon T may be stochastic. This again argues that the basic nature of the lower bound results is robust to changes in the modeling assumptions.

8. Random Illiquidity Horizons

To allow for a random illiquidity horizon, let us assume that the asset remains illiquid until the realization of a random event, which we model using a Poisson process with intensity λ . Thus, the illiquidity horizon T is exponentially distributed with density $\lambda \exp(-\lambda T)$. In turn, this implies that the mean time that the asset is illiquid is $1/\lambda$.

Since the exchange option value in Equation (6) provides the upper bound given a specific horizon T , we can solve for the upper bound in this random horizon case by taking the expectation of the exchange option value with respect to the exponential density of T . As shown in the appendix, this leads to the following remarkably simple expression for the expected value of the exchange option,

$$\frac{S_0}{\sqrt{1 + 8\lambda/\sigma^2}}. \quad (10)$$

Table 6 illustrates the lower bounds on thinly traded assets implied by this formula for various combinations of the mean illiquidity horizon (measured by $1/\lambda$) and asset volatility. As shown, the lower bound when the mean illiquidity horizon is one day is higher than the lower bound for a fixed illiquidity horizon of one day—as shown in Table 1—and similarly for all of the other horizons. At first glance, this result may seem

Table 6. Percentage Lower Bounds for Illiquid Asset Values When the Illiquidity Horizon is Stochastic

Mean illiquidity horizon	Volatility				
	10%	20%	30%	40%	50%
1 Day	99.776	99.553	99.329	99.106	98.882
1 Week	99.510	99.020	98.529	98.039	97.549
1 Month	98.979	97.959	96.940	95.921	94.904
1 Year	96.467	92.947	89.453	85.997	82.592
2 Years	95.006	90.050	85.166	80.388	75.746
5 Years	92.119	84.383	76.923	69.849	63.239
10 Years	88.889	78.178	68.200	59.175	51.205
20 Years	84.383	69.849	57.143	46.548	37.983
30 Years	80.988	63.884	49.767	38.763	30.439

Notes. This table reports the lower bound on the value of an illiquid asset expressed as a percentage of the price of an equivalent liquid asset. In this case, the asset remains illiquid until the realization of a Poisson process with intensity λ . The mean illiquidity horizon represents the average time that the asset is illiquid (determined as $1/\lambda$). Volatility denotes the volatility of returns for an equivalent liquid asset.

counterintuitive since it implies that the discount for illiquidity is less when the illiquidity horizon is random than when it is fixed—almost as if investors preferred uncertainty about the illiquidity horizon. The rationale for this result, however, is a simple artifact of Jensen’s inequality. Specifically, the exchange option is a convex function of the length of the illiquidity horizon. Thus, uncertainty over the illiquidity horizon results in the expected value being higher than the value of the exchange option evaluated at the expected length of the illiquidity horizon.

9. Conclusion

We model illiquidity as a restriction on the stopping rules that an investor can follow in selling asset holdings. We use this framework to derive realistic lower bounds on the value of illiquid and thinly traded investments.

A number of economic insights emerge from this analysis. For example, we show that immediacy plays a unique role and is much more highly valued than ongoing liquidity. In addition, we show that illiquidity can reduce the value of an asset substantially. For illiquidity horizons on the order of those common in private equity, the discount for illiquidity can be as much as 30%–50%. Although large in magnitude, these discounts are consistent with the empirical evidence on the valuation of thinly traded assets. Thus, these lower bounds could be useful in determining reservation prices and providing conservative valuations in situations where other methods of valuation are not available.

Finally, we find that the discount for illiquidity decreases as the cash flow generated by the underlying asset increases. Thus, investors in private ventures may have strong incentives to increase dividends and other cash flows to reduce the impact of illiquidity on their holdings. This implies that the illiquid nature of investments in partnerships, private equity, venture capital, LBOs, etc. has the potential to introduce agency conflicts as cash flow policy is impacted.

Acknowledgments

The author is grateful for helpful discussions with Robert Brooks, Maureen Chakraborty, Stillian Ghaidarov, and Stephen Schurman, and for the comments of participants at the 2015 Western Finance Association meetings. The author is particularly grateful for the comments and suggestions of department editor Gustavo Manso and an anonymous referee. All errors are the author’s responsibility.

Appendix

The value of the investor’s portfolio at time T if he is allowed to follow the optimal stopping rule is $X_T = S_\tau e^{r(T-\tau)}$. The value of the investor’s portfolio at time T if he is not allowed to sell until time T is $Y_T = S_T$. Substituting these expressions

into the payoff function $\max(0, X_T - Y_T)$ for the exchange option gives Equation (2).

The Black and Scholes (1973) formula for the time- t value of a European put with strike price K and time until expiration of $T - t$ is given by

$$Ke^{-r(T-t)}N(-d_2) - S_tN(-d_1), \quad (\text{A.1})$$

where

$$d_1 = \frac{\ln(S_t/K) + (r + \sigma^2/2)(T - t)}{\sqrt{\sigma^2(T - t)}}, \quad (\text{A.2})$$

$$d_2 = d_1 - \sqrt{\sigma^2(T - t)}. \quad (\text{A.3})$$

At the stopping time τ , the value of S_τ is known and is no longer stochastic. Thus, the value of the exchange option as of time τ is simply the present value of a put option on with strike price $S_\tau e^{r(T-\tau)}$ and time until expiration of $T - \tau$. Substituting these values into the Black–Scholes formula above gives the value of the exchange option at time τ ,

$$S_\tau [N(\sqrt{\sigma^2(T - \tau)}/2) - N(-\sqrt{\sigma^2(T - \tau)}/2)], \quad (\text{A.4})$$

which is Equation (3).

Standard results now imply that the value at time zero of the exchange option can be obtained by discounting the value in Equation (A.4),

$$E[e^{-r\tau} S_\tau [N(\sqrt{\sigma^2(T - \tau)}/2) - N(-\sqrt{\sigma^2(T - \tau)}/2)]], \quad (\text{A.5})$$

where the expectation is taken with respect to the joint distribution of S_τ and the stopping time τ under the risk-neutral measure.

To find the stopping rule that maximizes the value of the exchange option at time zero, we rewrite Equation (A.5) as

$$\max_\tau E[e^{-r\tau} S_\tau [N(\sqrt{\sigma^2(T - \tau)}/2) - N(-\sqrt{\sigma^2(T - \tau)}/2)]], \quad (\text{A.6})$$

where the inner expectation is taken with respect to the distribution of S_τ conditional on τ . From the dynamics of S given in Equation (1), it is readily shown that

$$e^{-r\tau} S_\tau = e^{-r\tau} S_0 \exp((r - \sigma^2/2)\tau + \sigma Z_\tau), \quad (\text{A.7})$$

$$= S_0 \exp(-\sigma^2\tau/2 + \sigma Z_\tau). \quad (\text{A.8})$$

The expression in Equation (A.8), however, is an exponential martingale. Thus, $E[e^{-r\tau} S_\tau] = S_0$ for all τ because of the strong Markov property of S_t .

Substituting this last result into Equation (A.6) gives

$$S_0 \max_\tau E[[N(\sqrt{\sigma^2(T - \tau)}/2) - N(-\sqrt{\sigma^2(T - \tau)}/2)]]. \quad (\text{A.9})$$

From the properties of the standard normal distribution function, however, it is easily shown that $N(x) - N(-x)$, where $x > 0$, is an increasing function of x . Thus, the expression in Equation (A.9) is maximized when τ takes the lowest value possible. In turn, this implies that the the stopping rule that maximizes the value of the exchange option in Equation (A.9) is to stop immediately, $\tau = 0$. Substituting this result into Equation (5) leads to the maximized value of the exchange option given in Equation (6).

As an alternative derivation of this last result, we could proceed recursively to show that at time $T - \epsilon$, the value of

the exchange option is maximized by stopping rather than waiting and stopping at time T . Similarly, the value of the exchange option is maximized by stopping at time $T - 2\epsilon$ rather than at time $T - \epsilon$, and so forth. This recursive argument again shows that the maximizing stopping rule is $\tau = 0$.

Differentiating the exchange option value in Equation (6) with respect to σ gives

$$S_0 \exp(\sigma^2 T/8) \sqrt{T}, \quad (\text{A.10})$$

which is positive. Similarly, differentiating the exchange option value with respect to T gives

$$S_0 \exp(\sigma^2 T/8) \frac{\sigma}{2\sqrt{T}}, \quad (\text{A.11})$$

which is positive.

Turning to the case with dividends, the asset price dynamics are given by

$$dS = (r - \rho)S dt + \sigma S dZ. \quad (\text{A.12})$$

The exchange option payoff function at time T is given by

$$\max(0, X_T - Y_T) \quad (\text{A.13})$$

$$= \max\left(0, S_T e^{r(T-\tau)} + \int_0^\tau \rho S_t e^{r(T-t)} dt - S_T - \int_0^T \rho S_t e^{r(T-t)} dt\right), \quad (\text{A.14})$$

$$= \max\left(0, S_\tau e^{r(T-\tau)} - S_T - \int_\tau^T \rho S_t e^{r(T-t)} dt\right), \quad (\text{A.15})$$

$$= S_\tau \max\left(0, e^{r(T-\tau)} - \frac{S_T}{S_\tau} - \int_\tau^T \rho \frac{S_t}{S_\tau} e^{r(T-t)} dt\right). \quad (\text{A.16})$$

The present value of this payoff function as of time τ is

$$S_\tau E_\tau \left[e^{-r(T-\tau)} \max\left(0, e^{r(T-\tau)} - \frac{S_T}{S_\tau} - \int_\tau^T \rho \frac{S_t}{S_\tau} e^{r(T-t)} dt\right) \right], \quad (\text{A.17})$$

which becomes

$$S_\tau E_\tau \left[e^{-r(T-\tau)} \max(0, e^{r(T-\tau)} - \exp((r - \rho - \sigma^2/2)(T - \tau)) + \sigma(Z_T - Z_\tau) - \rho \int_\tau^T \exp((r - \rho - \sigma^2/2)(t - \tau)) + \sigma(Z_t - Z_\tau) e^{r(T-t)} dt) \right], \quad (\text{A.18})$$

after substituting in the solution for the asset prices. In turn, this reduces to

$$S_\tau E_\tau \left[\max(0, 1 - \exp((-\rho - \sigma^2/2)(T - \tau) + \sigma(Z_T - Z_\tau)) - \rho \int_\tau^T \exp((-\rho - \sigma^2/2)(t - \tau) + \sigma(Z_t - Z_\tau)) dt) \right]. \quad (\text{A.19})$$

This last equation can also be expressed as

$$S_\tau E_\tau [\max(0, 1 - W)], \quad (\text{A.20})$$

where W is a martingale and is independent of S_τ . It is readily seen that the variance of W is a decreasing function of τ because of the independence of Brownian increments. From

Theorem 8 of Merton (1973), this implies that the value of the exchange option at time τ is a decreasing function of τ . Following a similar line of reasoning as above, this implies that the time-zero value of the exchange option is maximized by setting $\tau = 0$.

Finally, to solve for the expected value of the exchange option when the illiquidity horizon T is exponentially distributed, we note that Equation (6) can be expressed as

$$S_0 \operatorname{erf}\left(\sqrt{\frac{\sigma^2 T}{8}}\right), \quad (\text{A.21})$$

where $\operatorname{erf}(\cdot)$ is the error function described in Abramowitz and Stegun (1964, Chap. 7). Using Gradshteyn and Ryzhik (1980, Equation 6.283.2), the expectation

$$S_0 \lambda \int_0^\infty \operatorname{erf}\left(\sqrt{\frac{\sigma^2 t}{8}}\right) e^{-\lambda t} dt, \quad (\text{A.22})$$

reduces to Equation (10).

Endnotes

¹For example, Statement of Financial Accounting Standards (SFAS) 157 allows for the use of unverifiable inputs in the valuation of a broad category of illiquid assets that are designated as level 3 investments. The lower bound presented in this paper provides a conservative but much more objective standard for valuing these types of illiquid assets.

²More formally, let $I = [0, T]$, $T < \infty$, define the set of times where stopping is possible in the continuous time framework of this paper. Let $(\Omega, \mathcal{F}, \mathcal{F}_{t \in I}, P)$ be a filtered probability space. A random variable $\tau: \Omega \rightarrow I$ defined by a stopping rule is a stopping time if the event $\{\tau \leq t\}$ belongs to the σ -field \mathcal{F}_t for all t in I . For a discussion of stopping times, see Karatzas and Shreve (2000).

³By focusing on investors who are willing to consider holding thinly traded assets, we implicitly make the standard assumption that preferences are defined in terms of the portfolio value at time T . Thus, we rule out pathologies such as where the inability to sell at some time prior to T would result in an unboundedly negative utility, since, in that case, the investor would clearly never consider holding the illiquid asset in the first place.

⁴I am grateful to the referee for raising this issue.

⁵I am grateful to the referee for suggesting this alternative specification and raising the issue of whether a recursive solution can be found.

References

- Abramowitz M, Stegun IA (1964) *Handbook of Mathematical Functions* (Dover Publications, New York).
- Acharya VV, Pedersen LH (2005) Asset pricing with liquidity risk. *J. Financial Econom.* 77(2):375–410.
- Amihud Y, Mendelson H (1986) Asset pricing and the bid-ask spread. *J. Financial Econom.* 17(2):223–249.
- Amihud Y, Mendelson H, Pedersen LH (2005) Liquidity and asset prices. *Foundations Trends in Finance* 1:269–364.
- Ang A, Bollen N (2010) Locked up by a lockup: Valuing liquidity at a real option. *Financial Management* 39(3):1069–1095.
- Aragon G (2007) Share restrictions and asset pricing: Evidence from the hedge fund industry. *J. Financial Econom.* 83(1):33–58.
- Berkman H, Eleswarapu VR (1998) Short-term traders and liquidity: A test using Bombay stock exchange data. *J. Financial Econom.* 47(3):339–355.

- Black F, Scholes M (1973) The pricing of options and corporate liabilities. *J. Political Econom.* 81(3):637–654.
- Brenner M, Eldor R, Hauser S (2001) The price of options illiquidity. *J. Finance* 56(2):789–805.
- Brunnermeier M, Pedersen LH (2009) Market liquidity and funding liquidity. *Rev. Financial Stud.* 22(6):2201–2238.
- Chacko GC, Jurek JW, Stafford E (2008) The price of immediacy. *J. Finance* 63(3):1253–1290.
- Chen Z, Xiong P (2001) Discounts on illiquid stock: Evidence from china, Working paper, Yale University, New Haven, CT.
- Constantinides GM (1986) Capital market equilibrium with transaction costs. *J. Political Econom.* 94(3):842–862.
- Copeland TE, Galai D (1983) Information effects on the bid-ask spread. *J. Finance* 38(5):1457–1469.
- Demsetz H (1968) The cost of transacting. *Quart. J. Econom.* 82(1):33–53.
- Duffie D, Gârleanu N, Pedersen LH (2005) Over the counter markets. *Econometrica* 73(6):1815–1847.
- Duffie D, Gârleanu N, Pedersen LH (2007) Valuation in over-the-counter markets. *Rev. Financial Stud.* 20(6):1865–1900.
- Duffie D, Malamud S, Manso G (2009) Information percolation with equilibrium search dynamics. *Econometrica* 77(5):1513–1574.
- Finnerty JD (2012) An average strike put option model of the marketability discount. *J. Derivatives* 19(Summer, 4):53–69.
- Fleckenstein M, Longstaff FA, Lustig H (2014) The TIPS–treasury puzzle. *J. Finance* 69(5):2151–2197.
- Ghaidarov S (2014) Analytical bound on the cost of illiquidity for equity securities subject to sale restrictions. *J. Derivatives* 21(Summer, 4):31–48.
- Gradshteyn IS, Ryzhik I (1980) *Table of Integrals, Series, and Products*, 4th ed. (Academic Press, San Diego).
- Gromb D, Vayanos D (2002) Equilibrium and welfare in markets with financial constrained arbitrageurs. *J. Financial Econom.* 66(2–3):361–407.
- Grossman SJ, Miller MH (1988) Liquidity and market structure. *J. Finance* 43(3):617–633.
- Hansen LP, Jagannathan R (1991) Implications of security market data for models of dynamic economies. *J. Political Econom.* 99(2):225–262.
- Huang J, Wang J (2009) Liquidity and market crashes. *Rev. Financial Stud.* 22(7):2607–2643.
- Huang J, Wang J (2010) Market liquidity, asset prices, and welfare. *J. Financial Econom.* 95(1):107–127.
- Kahl M, Liu J, Longstaff FA (2003) Paper millionaires: How valuable is stock to a stockholder who is restricted from selling it? *J. Financial Econom.* 67(3):385–410.
- Karatzas I, Shreve SE (2000) *Brownian Motion and Stochastic Calculus*, 2nd ed. (Springer, New York).
- Lippman SA, McCall JJ (1986) An operational measure of liquidity. *Amer. Econom. Rev.* 76(1):43–55.
- Longstaff FA (1995) How much can marketability affect security values? *J. Finance* 50(5):1767–1774.
- Longstaff FA (2001) Optimal portfolio choice and the valuation of illiquid securities. *Rev. Financial Stud.* 14(2):407–431.
- Longstaff FA (2009) Portfolio claustrophobia: Asset pricing in markets with illiquid assets. *Amer. Econom. Rev.* 99(4):1119–1144.
- Merton RC (1973) Theory of rational option pricing. *Bell J. Econom. Management Sci.* 4(4):141–183.
- Pastor L, Stambaugh R (2003) Liquidity risk and expected stock returns. *J. Political Econom.* 111(3):642–685.
- Silber WL (1991) Discounts on restricted stock: The impact of illiquidity on stock prices. *Financial Analysts J.* 47(4):60–64.
- Stoll HR (2000) Friction. *J. Finance* 55(4):1479–1514.
- Vayanos D (1998) Transaction costs and asset prices: A dynamic equilibrium model. *Rev. Financial Stud.* 11(1):1–58.
- Vayanos D, Vila J-L (1999) Equilibrium interest rate and liquidity premium with transaction costs. *Econom. Theory* 13(3):509–539.
- Vayanos D, Wang T (2007) Search and endogenous concentration of liquidity in asset markets. *J. Econom. Theory* 136(1):66–104.
- Vayanos D, Weill P-O (2008) A search-based theory of the on-the-run phenomenon. *J. Finance* 63(3):1361–1398.