PART I

THE EFFECT OF LIQUIDITY COSTS ON SECURITIES PRICES AND RETURNS

Introduction and Overview

We begin by considering the *direct* effects of trading costs on the values of financial assets. Investors require compensation for the trading costs they pay when they buy or sell securities. If two assets generate the same cash flows over time but one of them is less liquid (has higher trading costs), rational investors will pay less for the less liquid asset, which costs more to trade. Consequently, the less liquid asset will have a lower value and a higher required (expected) return. Overall, we should observe that the returns on financial assets are increasing in asset illiquidity or transaction costs. Just as risk-averse investors require a higher return to compensate for a higher asset risk, we propose that investors require a higher return to compensate for greater asset illiquidity or transaction costs. The following chapters study these relations.

First, we address the meaning of illiquidity or trading costs. Stated simply, trading costs are the direct and indirect costs associated with trading a security. The most easily measured component of trading costs is the direct costs: brokerage fees, transaction taxes, and other trade-processing fees. In addition, there are search and delay costs that arise because the buyers and sellers of a security are not continuously available to transact, so a seller needs to search for buyers, especially if he or she needs to liquidate a large-size position and, similarly, a buyer needs to find sellers at the time he or she wants to buy. Another component of trading costs is the bid–ask spread. In securities markets with quoted bid and offer (ask) prices, which are the buying and selling prices, the buy transaction is naturally executed at a higher price than the sell transaction, resulting in a bid–ask spread. However, the bid and ask prices apply only to limited trade quantities. Larger transactions have a greater impact on the transaction price: They raise the buying price and lower the selling price, resulting in a market impact cost. The market impact cost is greater for larger-size transactions, when there is greater information asymmetry between the two parties in the transaction, and when there is greater friction in accessing the market by traders willing to trade.

The nature of trading costs depends on the structure of the market where the security is traded. Most financial markets include market makers, intermediaries who buy or sell securities for their own accounts to close or narrow down the time gaps between purchases and sales. These market makers may be "official" market makers, such as the New York Stock Exchange (NYSE) specialists or Nasdaq market makers, or they may be traders who submit limit orders to the market and thus stand ready to buy and sell on demand. Market makers are willing to take a long or short position in the security so they will, for example, buy the security when a sell order arrives, hoping to sell it shortly thereafter at a profit.[1] As this happens, market makers accept inventory risk for which they expect to be compensated. As their inventory position (short or long) increases, their risk further increases, which means that the compensation they demand for each additional unit they trade increases as well. Therefore, market makers quote bid and ask prices at which they are willing to sell or buy limited quantities of the security, but beyond that limited quantity they further increase the price at which they are willing to sell or reduce the price at which they are willing to buy, depending on the quantities involved.

This is further exacerbated by adverse selection, which results from the fact that there is asymmetric information between traders about the value of the asset being traded. Adverse selection arises when a trader sells a security because he has private information that the security is overpriced. An uninformed market maker or trader on the other side of the transaction will try to protect himself by offering the seller a lower price. Similarly, a purchase may indicate positive information that buyers have about the security's price, and this will induce uninformed market makers and traders to ask for a higher selling price.[2] The greater the extent of asymmetric information, the lower will be the selling price and the higher will be the buying price. Because uninformed sellers and buyers are viewed as being possibly informed ones – other traders and market makers cannot

[1] See Amihud and Mendelson (1980) for a formal model of market making. In the NYSE, in addition to trading by market making, there was trading by auction. Amihud and Mendelson (1987) analyze the effects of the two trading methods – auction and market making – on stock return behavior. Mendelosn (1982) models securities price behavior in an auction market.

[2] See Glosten and Milgrom (1985) and Kyle (1985) for formal models.

distinguish between informed and uninformed sellers and buyers – they have to bear a price discount when selling and pay a premium when buying. These discounts and premia constitute illiquidity costs for these uninformed traders.

The result of these effects is that the trading cost of buying from or selling to a liquidity provider, such as market maker, has a fixed component and a variable component. The "round-trip" fixed component of the trading cost is given by the bid–ask spread (the difference between the ask and bid prices), which is the cost of buying and selling a small quantity of the security (for a single transaction, the fixed component is taken to be half the bid–ask spread). The variable component of trading costs is given by the market impact cost: the more the investor buys or sells, the greater the trading cost per unit he is trading.

The trading cost of traded stock often amounts to less than a percent of its value. Data provided by the Investment Technology Group (ITG) for the twenty-five quarters ending in the fourth quarter of 2011 show that the average trading cost of stocks in the United States was 0.52% of stock value (including commissions), ranging between a low of 0.33% in the third quarter of 2007 and a high of 0.94% in the last quarter of 2008. Average U.S. equity trading costs ranged from 0.43% for large-cap stocks to 1.01% for small-cap stocks. In fact, large stocks are more liquid. In the United Kingdom and Japan, ITG estimates 2011 trading costs to be slightly above 0.5%.[3]

Although trading costs of the order of 0.5% may seem small, their value effect is large because they are incurred *repeatedly* each time the security is traded. Therefore, we need to consider the cumulative effect of trading costs throughout the security's life. Consider, for example, an asset that pays out a riskless annual dividend of $4 in perpetuity and suppose the risk-free annual rate is 4%. Absent trading costs, the asset price is $100. However, if the asset incurs a trading cost of $0.50 (0.5% of its value) and is traded once a year, the cash flow stream associated with the trading costs has a net present value of $12.5 of the asset's value, meaning that the price of the asset drops to $100 – $12.5 = $87.5. Said differently, while a transaction cost of 0.5% is a small fraction of the asset's value, it should really be compared to the 4% dividend yield, because both dividends and transaction costs are "flows" that are incurred repeatedly. Since the transaction cost is one-eighth of the dividend yield, its present value is one-eighth of the present value of dividends ($12.5/$100 = 1/8). Furthermore, if the asset is traded every half-year, then after accounting for transaction costs, the asset's value will be about $75, a

3 Elkins McSherry LLC, which measures trading costs in multiple markets around the world, also reports that recent trading costs of stocks are about 0.5%.

Cambridge University Press
978-0-521-19176-0 - Market Liquidity: Asset Pricing, Risk, and Crises
Yakov Amihud, Haim Mendelson and Lasse Heje Pedersen
Excerpt
More information

discount of \$25. The value discount is translated to a return premium. The \$4 dividend constitutes a 5.3% return on the asset whose price is \$75, which means a return premium of 1.3% due to transaction cost of 0.5% compared to the return on the perfectly liquid asset with the same cash flow.

A similar analysis can be applied to any asset whose cash payments grow over time and whose trading costs are proportional to its price. Consider a stock whose next-period cash dividend, D, grows at a rate of g, and its required return is $r > g$. Absent transaction costs, its price (the present value of its cash payments), is given by what is known as Gordon's growth formula $P_0 = D/(r - g)$. If transaction costs in the stock are a fraction c of its price, they increase at the same rate as the stock price, namely g. Assume that transaction costs are incurred at the same frequency for which r and g apply. The present value of the transaction costs is thus $cP/(r - g)$, and because $g = r - D/P$, where D/P is the dividend yield, we obtain that the relative price discount is simply $c/(D/P)$. In this context, $1/(D/P)$ is the *transaction cost multiplier* which, when applied to the transaction cost c, gives the price discount. The dividend yield D/P on the S&P 500 stocks has recently been about 2%, corresponding to a transaction cost multiplier of $1/(D/P) = 50$. With this dividend yield, a transaction cost of $c = {}^1\!/_2\%$ translates into a price discount of 25% (if the stock trades once a year). Therefore, we observe that a small trading cost brings about a large (fifty-fold in this example) decline in the stock's price. Securities with higher trading costs will have lower values and they will have to generate higher returns to become attractive to investors. This implies a higher cost of capital for firms whose securities have higher trading costs. Trading costs thus significantly affect firms' ability to raise capital for investments, the capital allocation process, and the real economy.

Higher trading costs can be better borne by long-term investors who trade less frequently and, therefore, can depreciate them over a longer investment horizon. Frequently trading investors are willing to pay more for assets with low transaction costs. In equilibrium, there will be *liquidity clienteles:* Other things being equal, fewer liquid assets will be held by investors with a longer expected holding period. Therefore, while expected return is an increasing function of trading costs, it should be concave (increasing at a decreasing rate), reflecting the mitigating effect of long-term holding periods on the sensitivity of return to transaction costs.

The following chapters show, both theoretically and empirically, that differences in trading costs explain differences in securities values and returns across stocks and bonds. Furthermore, when stock liquidity improves, its value rises, as the theory predicts. The first article by Amihud and Mendelson

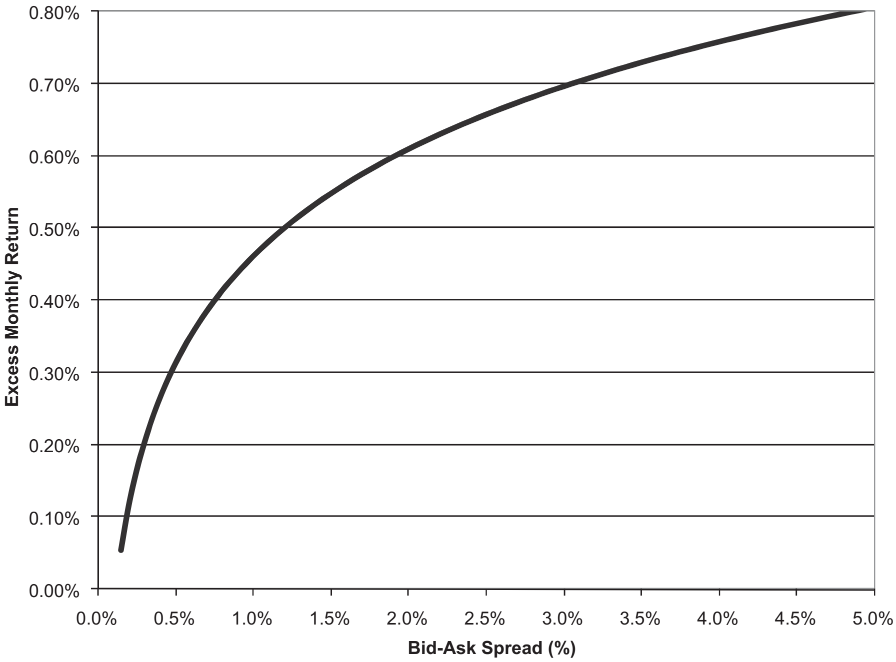

Figure PI.1. Relation between excess monthly return on NYSE stocks and their bid–ask spread, 1961–1980.

(1986) lays the theoretical foundations for the relation between asset returns and trading costs and shows empirically that this relation is economically and statistically significant across stocks that are traded on the NYSE and American Stock Exchange (AMEX). Figure PI.1 shows the estimated relation between excess monthly returns and stocks' bid–ask spreads. As suggested by the theory, the relation is increasing and concave.

Amihud and Mendelson (1986) measure stock illiquidity by the quoted bid–ask spread, a measure of trading costs that was available in the 1980s. A decade later, trade-by-trade data became available, which enabled Brennan and Subrahmanyam (1996) to estimate stock illiquidity using both the market impact cost and the bid–ask spread. The authors find that illiquidity increases the required return on stocks. Silber (1991) examines the effect of illiquidity on stock prices in the context of stocks whose trading is restricted. Consistent with the theory, the author finds that trading restrictions lower stock prices.

The finding that less liquid stocks generate higher (risk-adjusted) returns is supported in a recent paper by Amihud, Mendelson, and Goyenko (2010),

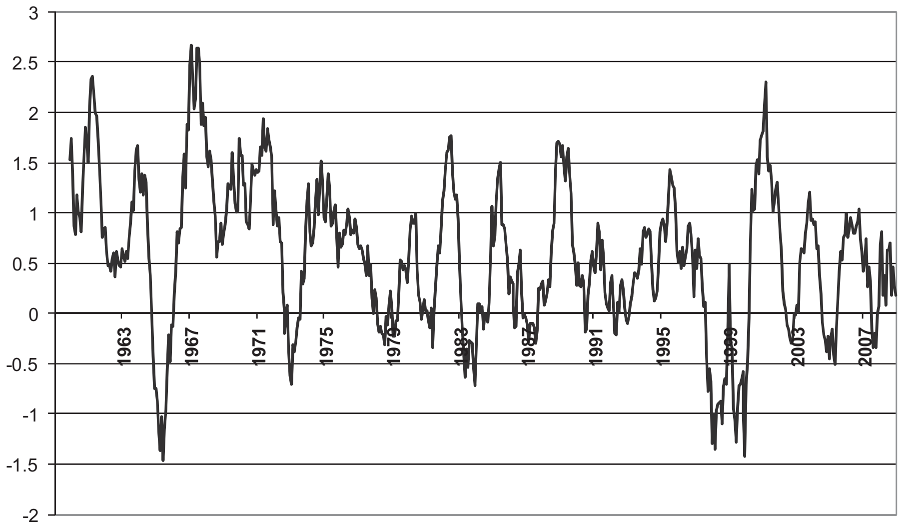

Figure PI.2. Twelve-month moving averages of monthly risk-adjusted liquidity premia (HMLI) for NYSE- and AMEX-traded stocks, 1960–2009 (in %). *Source:* Amihud, Mendelson and Goyenko (2010).

who study the return-liquidity relation over the past fifty years (1960–2009). Stocks on the NYSE and AMEX are sorted each month by their past illiquidity, using the illiquidity measure of Amihud (2002), conditional on their past volatility, and ranked into five equally sized portfolios. The monthly return on the high-minus-low illiquidity portfolios ([HMLI], top and bottom quintiles), measures the liquidity premium given by the excess return on illiquid stocks relative to liquid ones. The average HMLI liquidity premium over the fifty-year sample period is significantly positive. Adjusting for risk by a regression of HMLI on the four common risk factors of Fama and French (1993) and Carhart (1997) generates an alpha of 0.5% per month that is significantly positive. This means that in the past fifty years, the average risk-adjusted excess return of the HMLI portfolio is about 6% annually.

Figure PI.2 plots the twelve-month moving averages of the monthly risk-adjusted liquidity premia HMLI, given by the sum of the alpha coefficient and the residuals from the above regression of HMLI on the four risk factors. Figure PI.2 shows that while the risk-adjusted liquidity premium fluctuates over time, it is mostly positive. Although it becomes negative during the 1997–2000 period, it reverts to being positive during the last decade, including the period surrounding the recent financial crisis.

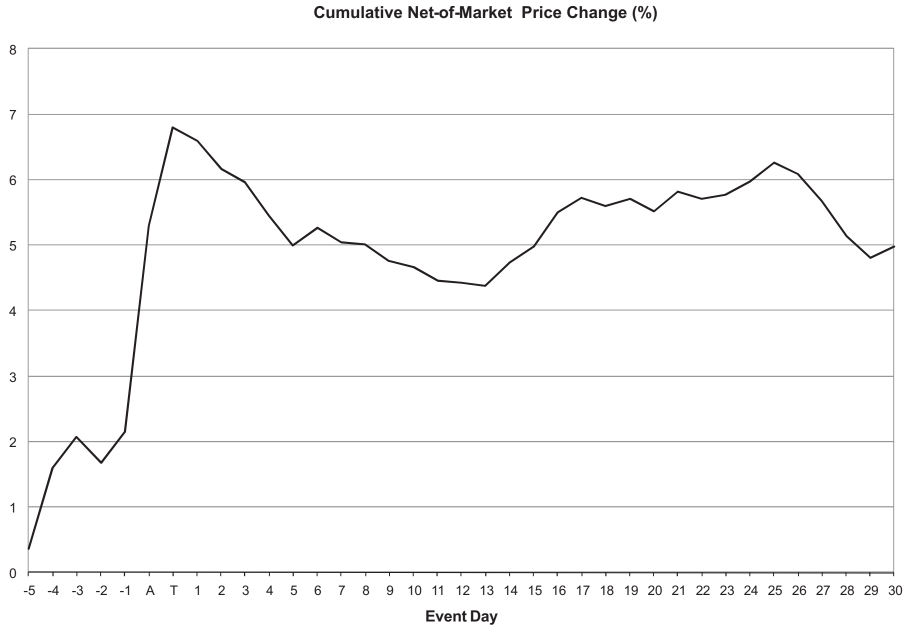

Figure PI.3. Cumulative price appreciation (net of market) for stocks transferred to a more liquid trading venue (in %). Day A is the day of the announcement on the stock being transferred to the new and more liquid trading venue; day T is the day the stock started trading in the more liquid trading venue.

The positive relation between asset expected returns and transaction costs also holds for fixed income securities. Amihud and Mendelson (1991; Chapter 2 in this book) study this relation for Treasury Bills and Notes. This work was extended by Chen, Lesmond, and Wei (2007) to the case of corporate bonds. For bonds, the tests examine the effects of liquidity on the yields to maturity. After controlling for risk and duration, both studies show that the yield to maturity is higher for less liquid bonds, as the theory predicts.

So far, this review has considered the cross-sectional relation between trading costs and securities returns (and values). A further step is taken by Amihud, Mendelson, and Lauterbach (1997; Chapter 3 in this book), who examine the effects of a *change* in stock liquidity of a given financial asset over time, showing that increased stock liquidity due to an improved trading system significantly raises stock prices. Whereas studies on the effect of liquidity on asset prices require controlling for other factors that affect asset returns, in this study the stocks and their underlying cash flows remained the same. Their liquidity, however, increased due to their transfer to a new venue and trading method. Importantly, the change in trading venue was exogenous and did not convey information about the stocks, because the decisions were made by the management of the stock exchange, without any company discretion. This comes as close as possible to a controlled experiment on the effects of changing liquidity on stock prices, where everything else remains unchanged. The stocks that were transferred to the more liquid trading venue enjoy a sharp and permanent price increase of nearly 6% on average. The evolution of stock prices around the time of the transfer to the new and more liquid trading venue is depicted in Figure PI.3.

CHAPTER 1

Asset Pricing and the Bid–Ask Spread

Summary and Implications

This article establishes the theory on the effect of liquidity on asset values and provides estimations of the relation between expected returns and liquidity across different stocks. The Amihud–Mendelson model gives rise to two major empirical predictions that are discussed in this chapter's introduction: expected asset returns increase in the assets' trading costs and the return–trading cost relation is concave. The first prediction results from the fact that investors demand a higher compensation for bearing higher trading costs. The second is due to the clientele effect: because less liquid assets are held in equilibrium by investors with longer holding periods, the additional compensation they require for an increase in trading costs is lower.

The Amihud–Mendelson model shows that, in equilibrium, the return on an asset whose trading is costly is equal to the return that would be earned on a similar-risk asset that is perfectly liquid (entailing zero trading costs) plus a return premium that compensates investors for the transaction costs they bear. That return premium is an increasing function of the expected trading cost per unit of time, which is the product of the asset's transaction cost by the frequency of asset sales. Consequently, higher trading costs lower asset prices because when discounting an asset's cash flow (dividend) at a higher rate of return due to higher trading costs, its value is lower. That is, higher trading costs produce an asset price discount.

Amihud and Mendelson show that there is a *clientele effect* whereby long-term investors tend to invest in assets that are less liquid (yielding higher returns) and short-term investors tend to invest in assets that are more liquid. Because of this specialization, the higher the trading costs, the smaller the effect of a marginal increase in these costs on the return required by investors. As a result, the required return on assets is not only an increasing function of transaction costs but also concave (increasing at a decreasing rate). Long-term investors can effectively depreciate their trading costs over a longer holding period, and thus require a smaller compensation in terms of per-period additional return than short-term investors.

Further, the Amihud–Mendelson model shows that the price discount due to trading costs consists of two components. The first component is the expected present value of all trading costs in the asset over its lifetime (for stocks, this is calculated by discounting the infinite stream of transaction cost cash flows). The second component reflects an additional discount in value that is needed to induce long-term investors to hold the less liquid assets. While all investors prefer assets with lower trading costs, long-term investors can outbid short-term investors on assets with any trading costs because long-term investors bear these costs less frequently. Long-term investors will not hold the less liquid assets unless offered more than a mere compensation for their higher expected transaction costs. To induce these investors to hold the less-liquid assets, their price net of expected trading costs must be lower than the net price of the more liquid assets. As a result, even after subtracting the present value of all trading costs, low-liquidity assets are still cheaper for their investors than liquid assets. Thus, the *net* return on assets, after subtracting the expected per-period cost of trading, is higher for assets with higher transaction costs.

Amihud and Mendelson test these predictions on the return–trading cost relation using data on stocks traded on the NYSE and AMEX over the period 1961–1980. Their measure of trading cost is the relative bid–ask spread, that is, the ratio of the dollar difference between the bid and ask prices to the stock price. The analysis groups stocks in each year into 49 (7×7) portfolios sorted on the previous year's relative spread and, within that, on past systematic risk (beta). The average bid–ask spreads on the lowest and highest spread portfolios are 0.49% and 3.2%, respectively, with the median-spread portfolio having an average bid–ask spread of 1.1%. The test procedure estimates a regression of the monthly return of each portfolio on the portfolio's bid–ask spread and beta. To test the clientele effect, the estimation allows the return–spread relation to be piecewise linear. Specifically, the estimation regresses the portfolio return in each year on dummy variables for each portfolio. The coefficients of these dummy variables provide the average return for each spread group and each beta risk group. The model also includes bid–ask spreads adjusted for the spread groups' mean, alowing a different coefficient for each of the seven spread