Netspar
Network for Studies on Pensions, Aging and Retirement

Netspar DISCUSSION PAPERS

Andrew Ang, Dimitris Papanikolaou and Mark Westerfield

Portfolio Choice with Illiquid Assets

DP 10/2010-092 (revised version August 2013)

Portfolio Choice with Illiquid Assets*

Andrew Ang
Columbia University and NBER

Dimitris Papanikolaou
Northwestern University and NBER

Mark M. Westerfield
University of Washington

August 2013

Abstract

We present a model of optimal allocation to liquid and illiquid assets, where illiquidity risk results from the restriction that an asset cannot be traded for intervals of uncertain duration. Illiquidity risk leads to increased and state-dependent risk aversion, and reduces the allocation to both liquid and illiquid risky assets. Uncertainty about the length of the illiquidity interval, as opposed to a deterministic non-trading interval, is a primary determinant of the cost of illiquidity. We allow market liquidity to vary from 'normal' periods, when all assets are fully liquid, to 'illiquidity crises', when some assets can only be traded infrequently. The possibility of a liquidity crisis leads to limited arbitrage in normal times. Investors are willing to forego 2% of their wealth to hedge against illiquidity crises occurring once every ten years.

JEL Classification: G11, G12

Keywords: Asset Allocation, Liquidity, Alternative Assets, Liquidity Crises

*We thank Andrea Eisfeldt, Will Goetzmann, Katya Kartashova, Leonid Kogan, Francis Longstaff, Jun Liu, Chris Mayer, Lubos Pastor, Liang Peng, Eduardo Schwartz, Dimitri Vayanos, Pietro Veronesi, and seminar participants at the Bank of Canada, the Brazilian Finance Society, the Financial Risks International Forum, Oxford, LBS, the Pacific Northwest Finance Conference, Texas A&M, UC Irvine, University of Florida, UNC, USC, the USC-UCLA-UCI Finance Day, the Q-group, and the Western Finance Association, for comments and helpful discussions. We thank Sarah Clark for providing data on illiquid assets for calibration.

1 Introduction

Many financial assets are illiquid, and this lack of liquidity can often be traced to difficulty in finding a counterparty with whom to trade. In several markets, appropriate counterparties need to have specialized abilities and capital which are in limited supply. Importantly, the waiting time until the next opportunity to trade is uncertain.[1] Further, systematic variation in the level of market liquidity is possible as financial intermediaries receive negative shocks and withdraw from market making. In this paper, we investigate the effects of this illiquidity and illiquidity risk on asset allocation.

We develop a tractable model of illiquidity. An illiquid asset can only be traded contingent on the arrival of a randomly occurring trading opportunity – a liquidity event – modeled as an i.i.d. Poisson process. We interpret these random trading times as the reduced-form outcome of a search process to find an appropriate counterparty (e.g. Diamond, 1982). Since the illiquid asset cannot be traded for an *uncertain* period of time, the investor is exposed to risk that cannot be hedged.

Illiquidity risk affects the portfolio choice problem in two ways. First, liquid and illiquid wealth are imperfect substitutes. The investor's immediate obligations – consumption or payout – can only be financed through liquid wealth. If the investor's liquid wealth drops to zero, these obligations cannot be met until after the next liquidity event. Thus, the investor reduces her allocation to *both* the liquid and illiquid risky assets in order to reduce the probability that a state with zero liquid wealth – as opposed to only zero total wealth – is reached. This increase in effective risk aversion corresponds to real-world situations where investors or investment funds are insolvent, not because their assets under management have hit zero, but because they cannot fund their immediate obligations. The resulting underinvestment in illiquid assets relative to the Merton benchmark is substantial.[2] Second, since the investor's ability to fund intermediate consumption depends on her liquid wealth, fluctuations in the share of illiquid assets in the portfolio induce endogenous *time-varying* risk aversion.

We extend our baseline model of illiquidity to include time-varying arrival rates of liquidity events. Specifically, we allow for infrequently occurring illiquidity crises. There are two distinct regimes. The first regime represents 'normal' times, where all assets are fully liquid, as in the Merton benchmark.

[1] Examples include structured credit products, private equity and venture capital; small equity and bond issues, or large real estate and infrastructure projects. In some of these markets, the waiting time until the next opportunity to trade is uncertain because the number of participants is small. In other cases, for instance in private equity and venture capital limited partner investments, the time of exit and re-investment is stochastic and depends on the IPO or M&A markets.

[2] A standard calibration indicates that if the expected time between liquidity events is once a year, the investor should cut her investment in the illiquid asset by 33% relative to an otherwise identical but fully liquid asset. Further, the investor should be prepared for large, skewed changes in the relative value of illiquid to liquid holdings in her portfolio.

The second regime represents a 'drying up' of liquidity – a liquidity crisis – where now one of the risky assets becomes illiquid and can only be traded at infrequent intervals. This model is thus applicable to a wide range of assets that are normally liquid, but are subject to occasional market freezes.[3]

The possibility of a liquidity crisis leads to limited arbitrage in normal times. We consider the case in which there are two perfectly correlated securities with different Sharpe ratios. In a portfolio choice model without illiquidity risk, this case presents an arbitrage opportunity; the investor takes positions of plus or minus infinity in the two different assets. When one asset is illiquid, the investor allocates a finite amount – and will not use leverage – in the 'arbitrage' trades, even in normal times. While both securities are fully liquid during normal times, 'arbitrage' trades entail a hidden cost because in the event of a crisis they become imperfect substitutes. The inability to immediately de-lever at the onset of a crisis implies that potential arbitrageurs avoid leverage in normal times. Hence, investment in apparent arbitrage opportunities is limited by the wealth of arbitrageurs, even when realizing the arbitrage involves no short positions.

We derive the risk premium associated with a systematic liquidity crisis.[4] Following an approach similar to the ICAPM Merton (1973), we examine the investor's marginal value of wealth, taking the prices of other risky assets as given. A transition from the liquid into the illiquid state raises the marginal value of liquid wealth, implying a negative risk premium. Hence, assets that pay off in the onset of a liquidity crisis earn lower risk premia. For typical parameter values, the investor is willing to pay 0.5% to 2% per annum over the actuarial probability of a crisis to receive liquid funds at the onset a deterioration of market liquidity. Hence, our model provides a theoretical justification of empirical specifications of the stochastic discount factor that load on measures of illiquidity.

Finally, we explore the determinants of the cost of illiquidity by varying the baseline model. First, we compute portfolio policies and the utility cost of illiquidity in a version of our model without intermediate consumption. We confirm that the sub-optimal diversification that results from infrequent trade is not sufficient by itself to generate large utility costs of illiquidity. Absent the motive to smooth intermediate consumption, the effect of illiquidity on portfolio policies and welfare is minimal. Second, we compare welfare and portfolio policies between our setting and a

[3]For example, Brunnermeier (2009), Gorton (2010), Tirole (2011), and others have highlighted market freezes as a stylized fact of the 2008-2009 financial crisis. This is not simply a question of a seller reducing prices to a level where a buyer is willing to step in. As Tirole (2011) and Krishnamurthy, Nagel, and Orlov (2012) comment, there were no bids, at any price, representing "buyers' strikes" in certain markets where whole classes of investors simply exited previously liquid markets.

[4]This 'illiquidity risk premium' refers to the risk premium of an Arrow-Debreu security that pays off at the onset of a liquidity crisis. This risk premium is distinct from the 'illiquidity premium', defined as the price discount of an illiquid security.

version of our model with a deterministic rebalancing interval. The effect of illiquidity on portfolio choice is dramatically larger when the length of the illiquidity period is uncertain. Third, we verify that our quantitative results are not driven by the possibility that the investor reaches states with zero liquid wealth – and thus infinite marginal utility – by allowing the investor to pay a fixed cost to transact immediately. Last, we disentangle the effect of preference parameters; we find that the utility cost of illiquidity is highest for agents that unwilling to substitute across time (low elasticity of intertemporal substitution) but are willing to substitute across states (low risk aversion).

2 Illiquidity in Asset Markets

We motivate and quantify our notion of illiquidity based on a number of stylized facts, and we contrast our model with previous approaches in the literature.

Stylized Fact 1 *Most asset classes are illiquid, in the sense that trading is infrequent.*

Table 1 shows that most assets markets are characterized by long times between trades, low turnover, and trade in over-the-counter markets in which it is difficult to find counterparties. Except for 'plain vanilla' fixed income securities and public equities, investors need to wait for indeterminate periods before they can re-balance illiquid assets, and sometimes the time between liquidity events can extend to decades. Even within the fixed income and public equity markets, there are sub-classes that are illiquid. For instance, while the public equity market has a turnover well over 100%, corporate bonds have a turnover around 25-35%. The average municipal bond trades only twice per year; the entire market has a turnover of less than 10% per year. Further, transactions times for many over-the-counter equities, such as those traded on the pinksheet or NASDAQ OTC-BB markets, are often longer than a week with a turnover of approximately 35%.

In real estate markets, Levitt and Syverson (2008) report, for example, a typical time to sale between 110-135 days after initial listing of a house. The standard deviation of the time to sale is even larger than the mean and Levitt and Syverson note that some houses never sell. Lastly, typical holding periods for venture capital and private equity portfolios are 3 to 10 years. Even though the investment horizon is nominally fixed, partnerships often return investor's money prior to the partnership's formal 10-year end. Further, these times are stochastic. For example, in private equity, the median investment duration is four years with 16% returned before two years and 26% returned after six years (see e.g. Lopez-de Silanes, Phalippou, and Gottschalg, 2010).[5]

[5]The turnover from trade of private equity investments on the secondary market is much lower. While data on private equity portfolio turnover is not typically reported. Kensington, a Canadian private equity fund, reports a 2% turnover in 2008. Alpinvest, a large private equity fund-of-funds reports flows that imply a turnover of approximately 6%. This compares with typical turnover of over 70% for mutual funds (Wermers, 2000).

In most of these markets, illiquidity is characterized by the need to find counterparties to trade. Our notion of illiquidity in which an illiquid asset can only trade when there is a liquidity event – when an appropriate counterparty is found – puts us squarely in the tradition of Diamond (1982). A number of authors have used models with search frictions to consider the impact of illiquidity risk. Duffie, Gârleanu, and Pedersen (2005, 2007) consider risk-neutral and CARA utility cases and restrict asset holdings at two levels: zero and one. In Vayanos and Weill (2008), agents can only go long or short one unit of the risky asset. Gârleanu (2009) and Lagos and Rocheteau (2009) allow for unrestricted portfolio choice, but Gârleanu considers only CARA utility and Lagos and Rocheteau focus on proving the existence of equilibrium with search frictions. In contrast to these models, we focus on the investors' portfolio choice problem with CRRA and Epstein-Zin utility, which allows us to consider more general preference specifications that allow for wealth effects.[6] An alternative, and equally plausible, micro-foundation for our notion of liquidity can be found in models with adverse selection. For instance, Daley and Green (2012, 2013) show that adverse selection can generate significant delays in trading times and time-varying costs of liquidity.

Our notion of liquidity is conceptually distinct from two other common conceptualizations previously studied in the portfolio choice literature. The first is the idea that many assets are expensive to trade: securities are partially marketable and can be traded at posted prices, but with transactions costs (e.g. Constantinides (1986); Grossman and Laroque (1990); Vayanos (1998); Lo, Mamaysky, and Wang (2004)). In these models, liquidity can *always* be generated by paying a cost. Our results imply that this uncertain waiting time plays an important role in portfolio decisions.

The second type of liquidity studied in portfolio choice allows assets to be freely traded at posted prices, but only in limited quantities (e.g. Longstaff (2001)) or at deterministic times (e.g. Kahl, Liu, and Longstaff (2003); Koren and Szeidl (2003); Schwartz and Tebaldi (2006); Longstaff (2009); De Roon, Guo, and Ter Horst (2009); Dai, Li, Liu, and Wang (2010)). In these models, trade can always be generated at a known rate simply by waiting.[7] Similarly, the margin-based asset pricing literature (e.g. Gârleanu and Pedersen (2011)) develops models with differential borrowing costs and portfolio constraints but then allows for continuous trade within those constraints. In contrast, we

[6]Our work also relates to models with unhedgeable human capital risk (e.g. Heaton and Lucas, 2000; Santos and Veronesi, 2006). An important distinction is that our illiquid asset is infrequently traded, unlike human capital which is never traded.

[7]The only other model that features random opportunities to trade is Rogers and Zane (2002), who solve a model with random trading opportunities and no liquid risky asset using asymptotic expansions near the Merton benchmark $1/\lambda \to 0$. However, Rogers and Zane do not prove that these expansions are valid. In contrast to Rogers and Zane, we solve the ODEs characterizing the investors' problem numerically and intentionally consider realistic cases where $1/\lambda$ is large, as is the case for many illiquid asset markets (see Table 1). The behavior of the model as $1/\lambda \to 0$ can be very different from the Merton benchmark. In particular, even as $1/\lambda \to 0$, the investor is still trading on a set of measure zero, hence would never take a short position in either liquid or illiquid wealth.

show that illiquidity risk leads to investors behaving as if they were subject to portfolio constraints – e.g. not taking short positions in illiquid or potentially illiquid assets – as a response to the illiquidity friction.

Stylized Fact 2 *Illiquid asset classes are large.*

The illiquid asset markets listed in Table 1 are large and rival the size of the public equity market. For instance, the market capitalization of NYSE and NASDAQ is approximately $17 trillion. The estimated size of the residential real estate market is $16 trillion and the estimated size of the (direct) institutional real estate market is $9 trillion.[8] Further, the share of illiquid assets in many investors' portfolios is very large. Kaplan and Violante (2010) show that individuals hold the majority of their net wealth in illiquid assets, with 91% and 81% of households' net portfolios tied up in illiquid positions, mostly housing, taking median and mean values, respectively. High net worth individuals in the U.S. allocate 10% of their portfolios to "treasure" assets like fine art, jewelry, and the share of treasure assets rises to almost 20% in other countries.[9] The share of illiquid assets in institutions' portfolios has also dramatically increased over the last 20 years. Pension funds increased their holdings in illiquid ("other") asset classes from 5% in 1995 to close to 20% in 2010, as reported in the "Global Pension Asset Study 2011" by Towers Watson. Data from the National Association of College and University Business Officers (NACUBO) show that the (dollar-weighted) average share of illiquid "alternatives" in university endowment portfolios rose from 25% in 2002 to 52% in 2010.

Stylized Fact 3 *Normally liquid asset classes sometimes become illiquid.*

An important feature of feature of financial markets is that sometimes liquidity dries up in markets that are normally liquid. For instance, Krishnamurthy et al. (2012) document that in the market for money market funds, a usually liquid market, there were instances of "buyers' strikes" during the recent financial crisis, where investors were unwilling to trade at any price. Anderson and Gascon (2009) note that the commercial paper market froze not only in the 2008-2009 financial crisis, but it also froze in 1970 when the Penn Central railroad collapsed. In both cases, the Federal Reserve stepped in to help restore liquidity.

[8]NYSE and NASDAQ market capitalizations are approximately $12 trillion and $5 trillion as of July 2012 from nyxdata.com and nasdaqtrader.com. The estimated size of the U.S. residential real estate market is at December 2011 and is estimated by Keely, van Ark, Levanon, and Burbank (2012), down from a peak of $23 trillion in 2006. The estimate of the U.S. institutional real estate market is by Florance, Miller, Spivey, and Peng (2010), with the institutional real estate market losing $4 trillion from 2006 to 2010. The direct real estate market dwarfs the traded REIT market, with the FTSE NAREIT All Equity REITs Index having a total market capitalization of approximately $500 billion at the end of June 2012.

[9]Reported in "Profit or Pleasure? Exploring the Motivations Behind Treasure Trends," Wealth Insights, Barclays Wealth and Investment Management, 2012.

These illiquidity crises occur regularly in many asset markets.[10] Other examples include the repo market (Gorton and Metrick, 2012); residential and commercial mortgage-backed securities (Gorton, 2009; Acharya and Schnabl, 2010; Dwyer and Tkac, 2009); structured credit (Brunnermeier, 2009); and the auction rate security market (McConell and Saretto, 2010), which became illiquid in 2008 and is still frozen in 2013.[11] Borio (2004) notes that liquidity also dries up during periods of severe market distress: the Latin American debt crisis in the 1980s, the Asian emerging market crisis in the 1990s, Russian default crisis in 1998 (see also Elul, 2008), and of course the financial crisis over 2008-2009. Major liquidity crises have occurred at least once every ten years, many occurring in tandem with large downturns in asset markets.

We extend our baseline model to allow for infrequent illiquidity crises which has one normal regime, where all assets are liquid, and a regime representing a liquidity crisis, where the illiquid asset can only be traded at infrequently occurring liquidity events. In our model, investors underinvest in arbitrage opportunities during normal times because they carry illiquidity risk if a liquidity crisis arrives. Other models also generate limited arbitrage, like (e.g. Shleifer and Vishny, 1997; Gromb and Vayanos, 2002), but in these models investors underinvest in arbitrage opportunities because demand shocks to other investors can push prices further away from fundamentals. Limited arbitrage occurs in our setting, in contrast, due to the risk of a market freeze. At the onset of a liquidity crisis, arbitrageurs cannot reduce leveraged positions to prevent states with zero consumption. Hence they are unwilling to employ leverage in normal times, and arbitrage activity is limited by their wealth.

Last, our model provides a theoretical framework to study the pricing of liquidity risk. Recent empirical work has documented that several measures of liquidity are priced in the cross-section of asset returns (e.g. Pastor and Stambaugh, 2003; Acharya and Pedersen, 2005; Sadka, 2006; Korajczyk and Sadka, 2008). Specifically, stocks with heterogenous exposure to various measures of market liquidity earn different average stock returns, often controlling for their own level of liquidity. To the extent that these empirical measures of liquidity are correlated with the difficulty of finding a counterparty to trade, our model provides a framework to quantify the magnitude of this liquidity. In the spirit of Merton's ICAPM, we derive the magnitude of this risk premium directly from the first order conditions of an optimizing investor faced with the possibility of a liquidity crisis. Relative

[10] Practitioners refer to these events as "liquidity black holes." See (Taleb, 1997; Persaud, 2001; Morris and Shin, 2010).

[11] As noted by the SEC, "Report on the Municipal Securities Market," July 31, 2012: in 2008, the auction rate securities (ARS) market totaled approximately $200 billion; in February 2008, the market froze because there were no bidders in the primary auctions, where floating interest rates are set. As there was no secondary market, thousands of customers were unable to sell their ARS holdings. In 2011, there were no new issues of ARS. Other intermediated fund vehicles also became more illiquid during this time: hedge funds, for example, imposed 'gates' provisions that did not allow for investors to withdraw capital (see Ang and Bollen, 2010).

to Acharya and Pedersen (2005), who also propose a theoretical model with liquidity risk, the utility cost of illiquidity is endogenous in our model.

3 The Baseline Model

3.1 Information

The information structure obeys standard technical assumptions. There exists a complete probability space $(\Omega, \mathcal{F}, \mathcal{P})$ supporting the vector of two independent Brownian motions $Z_t = (Z_t^1, Z_t^2)$ and an independent Poisson process (N_t). $\mathcal{P}$ is the corresponding measure and $\mathcal{F}$ is a right-continuous increasing filtration generated by $Z \times N$.

3.2 Assets

There are three assets in the economy: a risk-free bond B, a liquid risky asset S, and an illiquid risky asset P. The riskless bond B appreciates at a constant rate r:

$$dB_t = r\, B_t\, dt \tag{1}$$

The second asset S is a risky asset whose price follows a geometric Brownian motion with drift μ and volatility σ:

$$\frac{dS_t}{S_t} = \mu\, dt + \sigma\, dZ_t^1. \tag{2}$$

The first two assets B and S are liquid and holdings can be re-balanced continuously.

The third asset P is an illiquid risky asset; its fundamental value evolves according to a geometric Brownian motion with drift ν and volatility ψ:

$$\frac{dP_t}{P_t} = \nu\, dt + \psi\rho\, dZ_t^1 + \psi\sqrt{1-\rho^2}\, dZ_t^2, \tag{3}$$

where ρ captures the correlation between the returns on the two risky assets. The illiquid asset P differs from the first two assets B and S because it can only be traded at stochastic times τ, which coincide with the arrival of a Poisson process with intensity λ. The parameter λ captures the severity of the illiquidity friction; the expected time between liquidity events is $1/\lambda$. When a trading opportunity arrives, the investor can trade at the price P_t without any other frictions.

In addition, the illiquid asset P cannot be pledged as collateral. Investors can issue non-state contingent debt by taking a short position in the riskless bond B; however, they cannot issue risky debt using the illiquid asset as collateral. If investors were allowed to do so, they could convert the illiquid asset into liquid wealth and thus implicitly circumvent the illiquidity friction.

3.3 The Investor

The investor has CRRA utility over sequences of consumption, C_t, given by:

$$\mathrm{E}\left[\int_0^\infty e^{-\beta t}\frac{C^{1-\gamma}}{1-\gamma}dt\right], \tag{4}$$

where β is the subjective discount factor and $\gamma > 1$.[12]

The agent's wealth has two components, liquid and illiquid wealth. Liquid wealth includes the amount invested in the liquid risky asset and the risk-free asset. Illiquid wealth, which equals the amount invested in the illiquid asset, cannot be immediately consumed nor converted into liquid wealth. The joint evolution of the investor's liquid, W_t, and illiquid wealth, X_t, is given by:

$$\frac{dW_t}{W_t} = \left(r + (\mu - r)\,\theta_t - c_t\right)dt + \theta_t\sigma dZ_t^1 - \frac{dI_t}{W_t} \tag{5}$$

$$\frac{dX_t}{X_t} = \nu dt + \psi\rho dZ_t^1 + \psi\sqrt{1-\rho^2}dZ_t^2 + \frac{dI_t}{X_t}. \tag{6}$$

The agent invests a fraction θ of her liquid wealth into the liquid risky asset, while the remainder $(1-\theta)$ is invested in the bond. Following Dybvig and Huang (1988) and Cox and Huang (1989), we restrict the set of admissible trading strategies, θ, to those that satisfy the standard integrability conditions. All policies are appropriately adapted to $\mathcal{F}_t$. The agent consumes out of liquid wealth, so liquid wealth decreases at rate $c_t = C_t/W_t$. When a trading opportunity arrives, the agent can transfer an amount dI_τ from her liquid wealth to the illiquid asset.

Finally, we assume the standard discount rate restriction, as in the Merton two-risky-asset model,

$$\beta > (1-\gamma)r + \frac{1-\gamma}{2\gamma(1-\rho^2)}\left(\left(\frac{\mu-r}{\sigma}\right)^2 - 2\rho\left(\frac{\mu-r}{\sigma}\right)\left(\frac{\nu-r}{\psi}\right) + \left(\frac{\nu-r}{\psi}\right)^2\right), \tag{7}$$

and that the illiquid asset has at least as high a Sharpe ratio as the liquid asset,

$$\frac{\nu-r}{\psi} \geq \frac{\mu-r}{\sigma}. \tag{8}$$

3.4 Discussion of Assumptions

Our assumption that the illiquid asset P cannot be collaterilized is motivated by the difficulty of finding a counterparty who is willing to lend cash using illiquid assets as collateral.[13] Alternatively, we

[12] The $\gamma = 1$ (log) case is qualitatively similar to the $\gamma > 1$ case. Quantitatively, the results for the log case can actually be more extreme because the low level of risk aversion results in a very high percentage of wealth invested in risky assets in the benchmark Merton economy. We discuss variation in risk aversion in Section 5.4, and a characterization of the Hamilton-Jacobi-Bellman equation in the log case is available from the authors on request.

[13] For instance, Krishnamurthy et al. (2012) find evidence suggesting that money market mutual funds, which are the main providers of repo financing, were unwilling to accept private asset-backed securities as collateral between the third quarter of 2008 and the third quarter of 2009. Even when illiquid assets like real estate, private equity, and

could re-interpret P as the fraction of illiquid wealth that cannot be collateralized. This interpretation is similar to the portfolio constraints analyzed in Gârleanu and Pedersen (2011).[14] The key difference is that we introduce an *additional* stochastic trading friction on the non-collateralizable portion of wealth, namely that it can be traded only infrequently.

Our assumption of an infinite horizon for the investor is conservative; any effects of illiquidity are magnified with finite horizons. For instance, if opportunities to trade arise every 10 years, on average, then an investor with a one-year horizon views the illiquid asset as a very unattractive asset. Thus, the portfolio weights, effects on consumption policies, and certainty equivalent compensations for bearing illiquidity risk should all be viewed as conservative bounds for finite-horizon investors.

4 The Solution to the Baseline Model

Markets are not dynamically complete, hence we use dynamic programming techniques to solve the investor's problem. First, we establish some basic properties of the solution. Then, we compute the investor's value function and optimal portfolio and consumption policies.

4.1 The Value Function

The agent's value function is equal to the discounted present value of her utility flow,

$$F(W_t, X_t) = \max_{\{\theta, I, c\}} \mathrm{E}_t \left[\int_t^\infty e^{-\beta(s-t)} U(C_s) ds \right]. \tag{9}$$

Problem 1 (Baseline) *The investor performs the maximization in (9), subject to the two intertemporal budget constraints (5) and (6), with re-balancing ($dI_t \neq 0$) only when the Poisson process N_t^λ jumps.*

Our first step is to establish that the investor does not use leverage – that liquidity risk eliminates any willingness by the investor either to short the illiquid asset or to fund long purchases of the illiquid asset using a net short position in liquid wealth:

Proposition 1 *Any optimal policies in Problem 1 will have $W > 0$ and $X \geq 0$ a.s.*

even art can be used as collateral, an investor cannot borrow an amount equal to the whole value of the illiquid asset position and thus our decomposition into liquid and illiquid assets is still valid.

[14] In the case of real estate, we could interpret the illiquid asset P as the fraction of the value of the property that cannot be used as collateral against a mortgage or a home equity line. Our interpretation assumes that the amount that the asset can be collateralized does not vary over time and that the constraint is always binding. We could extend the model to allow the investor to endogenously choose the amount of collateralized borrowing every period, up to a limit. This model is equivalent to a hybrid model of infrequent trading and transaction costs, with similar qualitative effects as the extension we introduce in Section 5.3.

Thus, without loss of generality, we restrict our attention to solutions with $W_t > 0$ and $X_t \geq 0$.

Second, the value function is bounded below by the problem in which the illiquid asset does not exist, and it is bounded above by the problem in which the entire portfolio can be continuously re-balanced: the Merton one- and two-stock problems.

Third, the utility function is homothetic and the return processes have constant moments, and so it must be the case that F is homogeneous of degree $1-\gamma$:

$$F(W,X) = (W+X)^{1-\gamma} H(\xi), \quad \text{where} \quad \xi \equiv \frac{X}{X+W}. \tag{10}$$

Thus, the investor's value function can be represented as a power function of total wealth times a function $H(\xi)$ of the fraction of her portfolio held in illiquid assets, ξ.

Our fourth step is to characterize the value function at the instant when the agent can re-balance between her liquid and illiquid wealth. When the Poisson process hits and the agent re-balances her portfolio, the value function will jump discretely. Denote the new, higher, value function as $F^*(W_t, X_t)$, so that the total amount of the jump is $F^* - F$. At the Poisson arrival, the agent is free to make changes to her entire portfolio, and thus

$$F^*(W_t, X_t) = \max_{I \in [-X_t, W_t)} F(W_t - I, X_t + I). \tag{11}$$

Since F^* must also be homogeneous of degree $1-\gamma$, there exists a function H^* such that $F^* = (W+X)^{1-\gamma} H^*(\xi)$. Since trading the illiquid asset is costless when a liquidity event arrives, the investor re-balances her portfolio so that the ratio of illiquid to total wealth equals $\xi^* = \arg\max H(\xi)$; hence, H^* is a constant function and equal to $H(\xi^*)$.

We can now solve the Baseline investor's problem:

Proposition 2 (Baseline) *In Problem 1, the investor's value function can be written as in (10), where $H(\xi)$ exists and is finite, continuous, and concave for $\xi \in [0,1)$. $H(\xi)$ obtains its maximum for some $\xi \in [0,1)$. Define $H^* = \max_\xi H(\xi)$ and $\xi^* = \arg\max_\xi H(\xi)$. When a trading opportunity occurs at time τ, the trader selects I_τ so that $\frac{X_\tau}{X_\tau + W_\tau} = \xi^*$. In addition, $H(\xi)$ is characterized between liquidity events by*

$$\begin{aligned} 0 = \max_{c,\theta} \Bigg[& \frac{1}{1-\gamma} c^{1-\gamma} (1-\xi)^{1-\gamma} - \beta H(\xi) + \lambda (H^* - H(\xi)) + H(\xi) A(\xi, c, \theta) + \frac{\partial H(\xi)}{\partial \xi} B(\xi, c, \theta) \\ & + \frac{1}{2} \frac{\partial^2 H(\xi)}{\partial \xi^2} C(\xi, c, \theta) \Bigg]. \end{aligned} \tag{12}$$

where the functions A, B, and C are defined as

$$
\begin{aligned}
A(\xi,c,\theta) \equiv & (1-\gamma)\Big(r+(1-\xi)((\mu-r)\theta-c)+\xi(\nu-r)-\frac{1}{2}\gamma\left(\xi^2\psi^2+(1-\xi)^2\sigma^2\theta^2+2\xi(1-\xi)\psi\rho\sigma\theta\right)\Big) \\
B(\xi,c,\theta) \equiv & \,\xi(1-\xi)\left(\nu-(r+(\mu-r)\theta-c)+\gamma\psi\theta\rho\sigma(2\xi-1)+\gamma\theta^2\sigma^2(1-\xi)+\gamma\psi^2\xi\right) \\
C(\xi,c,\theta) \equiv & \,\xi^2(1-\xi)^2\left(\theta^2\sigma^2+\psi^2-2\psi\theta\rho\sigma\right)
\end{aligned} \tag{13}
$$

The investor's value function has two parts. The first part $(W+X)^{1-\gamma}$ captures the effect of total wealth on the continuation utility. The second component $H(\xi)$ captures the effect of wealth composition between liquid and illiquid wealth. We hold total wealth constant and plot the function $H(\xi)$ in Panel A of Figure 1. We can interpret $H(\xi)$ as a composition penalty function: the investor has an optimal portfolio composition ξ^* to which she returns whenever she is able to re-balance. Between liquidity events, she experiences a welfare loss for two reasons as her portfolio composition deviates from the optimum. First, there is the standard effect from lack of optimal diversification. Second, there is an asymmetric effect arising from the fact that consumption is funded by liquid wealth only. Examining the slope of $H(\xi)$ as $\xi \to 1$, we see that this second effect is the main mechanism in our model.

4.2 Imperfect Substitutability of Liquid and Illiquid Wealth

In our model, liquid and illiquid wealth are imperfect substitutes. Illiquid wealth can be used to fund consumption only after the next trading time $t = \tau$. In contrast, liquid wealth can fund consumption both before and after τ. To quantify this non-substitutability, consider a fictitious market that lets the investor exchange one unit of illiquid wealth for q units of liquid wealth. Between liquidity events, the investor would be indifferent in participating in this fictitious market as long as

$$
q = \frac{F_X}{F_W}. \tag{14}
$$

When the investor has the opportunity to re-balance, $q = 1$. Between liquidity events, the relative price q differs from one, depending on whether the investor has too much, or too little illiquid wealth X relative to her desired allocation. In Panel B of Figure 1 we see that the relative price of illiquid wealth rapidly declines as the investor's allocation to illiquid assets ξ increases. When the illiquid endowment is large $X \gg W$, liquid wealth W is only used to fund immediate consumption, while illiquid wealth is used to fund future consumption. In this case, variation in liquid wealth becomes unimportant for long-run consumption and the value function becomes separable in X and W. Hence, even though liquid and illiquid wealth may be correlated, that correlation becomes a secondary concern for portfolio allocation.

4.3 Parameter Values

In our numerical solutions, we select our parameter values so that the liquid asset can be interpreted as an investment in the aggregate stock market. We set the parameters of the liquid asset to be $\mu = 0.12$, $\sigma = 0.15$, and set the risk-free rate to be $r = 0.04$. Table 2 shows that this set of parameters closely matches the performance of the S&P500 before the financial crisis.[15] We work mostly with the risk aversion case $\gamma = 6$, which for an investor allocating money between only the S&P500 and the risk-free asset produces an equity holding close to a classic 60% equity, 40% risk-free bond portfolio used by many institutional investors.

For most of our analysis, we take a conservative approach and set the parameters of the illiquid asset, $\nu = 0.12$ and $\psi = 0.15$, to be the same as those of the liquid asset. This has the advantage of isolating the effects of illiquidity rather than obtaining results due to the higher Sharpe ratio of the illiquid asset. Further, even for individual funds this assumption is not unrealistic, at least for some illiquid asset classes.[16] These parameters imply that our illiquid asset can be interpreted as any composite investment with the same Sharpe ratio as public equities, for example a composite illiquid risky bond investment. Further, to isolate the effect of illiquidity, in the baseline case we assume that the two risky assets are uncorrelated, $\rho = 0$; we explore the effect of correlation by subsequently varying ρ between 0 and 1.

Regarding the severity of the illiquidity friction, we take a baseline case of $\lambda = 1$, implying on average one year between transactions. For comparison, individual private equity, buyout, and venture capital funds can have average investment durations of approximately four years, which corresponds to $\lambda = 1/4$; an appropriate horizon for a single large real estate investment by institutions is 10 years ($\lambda = 1/10$) (see e.g. Table 1). Since λ is an important parameter, we take special care to show the portfolio and consumption implications for a broad range of λ. The economics behind the solution are immune to the particular parameter values chosen.

4.4 Optimal Portfolio Policies

In this section we characterize the investor's optimal asset allocation and consumption policies. Even though the investment opportunity set is constant, the optimal policies vary over time as a function

[15] The mean of the S&P500 including 2008-2010 falls to 0.10 and slightly more volatile, at 0.18, but our calibrated values are still close to these estimated values.

[16] Kaplan and Schoar (2005), Driessen, Lin, and Phalippou (2008) and Phalippou and Gottschalg (2009), for example, estimate private equity fund alphas, with respect to equity market indexes, close to zero. Table 2 shows that the reported returns on a composite illiquid investment in private equity, buyout, and venture capital has similar characteristics to equity. For example, over the full sample (1981-2010), the mean log return on the illiquid investment is 0.11 with a volatility of 0.17. This is close to the S&P500 mean and volatility of 0.10 and 0.18, respectively, over that period. Table 2 shows that the returns on liquid and illiquid investments are even closer in terms of means and volatilities before the financial crisis.

of the amount of illiquid assets held in the investor's portfolio.

Participation

Before characterizing the optimal allocation, we first show the sufficient conditions for the investor to have a non-zero holding of the illiquid asset:

Proposition 3 *An investor prefers holding a small amount of the illiquid asset to holding a zero position if and only if*

$$\frac{\nu - r}{\psi} \geq \rho \frac{\mu - r}{\sigma}. \tag{15}$$

The condition for participation is identical to the Merton two-asset case and depends only on the mean-variance properties of the two securities. Somewhat surprisingly, the degree of illiquidity λ does not affect the decision to invest a small amount in the illiquid asset because of the infinite horizon of the agent: a trading opportunity will eventually arrive where the illiquid asset can be converted to liquid wealth and eventual consumption. However, even though the conditions for participation are the same as the standard Merton case, the optimal holdings of the illiquid and liquid assets are very different, as we show below.

Illiquid Asset Holdings

Illiquidity induces underinvestment in the illiquid asset relative to the Merton case. In Table 3, we present the investor's optimal rebalancing point ξ^* along with the long-run average level illiquid portfolio holdings $E[\xi]$ for different values of λ. For comparison, and in an abuse of notation, we report the consumption and portfolio policies for an investor able to continuously trade one ($\lambda = 0$) and two ($\lambda = \infty$) risky assets. The optimal holding of illiquid assets at $\lambda = 1$ upon arrival of a liquidity event is 0.37, which is lower than the optimal two-asset Merton holding at 0.60.

In addition to underinvestment in the illiquid asset, the inability to trade implies that the investor's portfolio can deviate from optimal diversification for a long time. Panel C of Figure 1 plots the stationary distribution of an investor's holding of the illiquid asset, ξ. For most of the time – the 20% to 80% range – the share of wealth allocated in illiquid securities is 0.36 to 0.45, while the 1% to 99% range is 0.30 to 0.65. Furthermore, the distribution of portfolio holdings is positively skewed, since illiquid wealth grows faster on average than liquid wealth since only the latter is used to fund consumption. As a result of this skewness, the investor chooses a rebalancing point *lower* than the mean of the steady-state distribution of portfolio holdings, that is, $\xi^* < E[\xi_t]$.

The degree of skewness is increasing in the illiquidity of the investment. When $\lambda = 1$, the mean

holding is 0.41, compared to a rebalancing value of 0.37, while the distribution of portfolio holdings has a standard deviation of 6.3% and normalized skewness coefficient of 1.9. In the case when the investor can trade once every four years on average ($\lambda = 4$), the standard deviation of the investor's illiquid holdings increases to 12% and the skewness increases to 2.3.

Liquid Asset Holdings

In addition to under-investment in the illiquid asset, illiquidity affects the investment in the liquid asset. The allocation to the liquid risk asset as a fraction of the investor's liquid holdings is equal to

$$\theta_t = \frac{\mu - r}{\sigma^2}\left(-\frac{F_W}{F_{WW}W_t}\right) + \rho\frac{\psi}{\sigma}\left(-\frac{F_{WX}X_t}{F_{WW}W_t}\right). \tag{16}$$

The allocation to the liquid asset as a function of her *total* wealth is equal to $\theta\,(1-\xi)$. There are two aspects of the optimal policy that merit attention.

First, even in the case where the liquid and illiquid asset returns are uncorrelated, $\rho = 0$, the allocation to the liquid asset differs from the Merton benchmark due to *time-varying effective risk aversion*. In Panel D of Figure 1, we compare the curvature of the investor's value function with respect to liquid wealth $-F_{WW}W/F_W$ to that of a Merton investor. For low values of allocation to illiquid assets, ξ, the two behave in a similar fashion: as the share of liquid wealth W declines in the investor's total wealth $W + X$, so does the investor's aversion to gambles in W. However, when the investor's liquid wealth becomes sufficiently low, the two lines diverge, since liquid wealth is no longer viewed as a substitute for illiquid wealth. The investor in our problem becomes much more averse to taking gambles in terms of liquid wealth than a Merton investor. Further, her effective risk aversion not only increases but it varies over time as a function of her current allocation to illiquid assets, ξ.

Second, in the case where the liquid and illiquid asset are correlated, $\rho \neq 0$, the investor hedges changes in the value of illiquid wealth. The hedging demands depend on the correlation between the liquid and illiquid asset returns, ρ, and the elasticity of substitution between liquid and illiquid wealth, $-F_{WX}X/F_{WW}W$. In Panel E of Figure 1 we plot the second component for the demand for the liquid risky asset, $-F_{WX}X/F_{WW}W$, and contrast it to the term corresponding to a Merton investor for the case of $\rho = 0.6$. For low values of X relative to total wealth the two lines are very similar, whereas they diverge dramatically as X increases relative to W. In our model, the term $-F_{WX}X/F_{WW}W$ converges to zero rather than minus infinity in the Merton case, implying *zero hedging demand at the limit*. The hedging motive disappears when illiquid securities comprise the majority of the agent's portfolio since illiquid assets are not a substitute for liquid wealth. In this

case, the investor chooses the allocation in liquid assets to smooth consumption rather than hedging fluctuations in her illiquid portfolio.

Panel F of Figure 1 plots the agent's optimal allocation to the liquid risky asset as a function of her current allocation in illiquid assets ξ. The agent partially compensates for the risk of being unable to trade the illiquid asset for a long period of time by underinvesting in the liquid risky asset. In Table 3 we summarize the average long-run holdings in the liquid risky asset for different degrees of illiquidity $1/\lambda$. Illiquidity negatively impacts the allocation to the liquid risky asset, but less so than the illiquid asset. In the case where $\lambda = 1$, the investor reduces her allocation in the liquid risky asset from 60% in the Merton benchmark to 56%, compared to a reduction from 60% to 37% for the illiquid security.

Effect of Correlation

To study the effect of correlation on portfolio policies, we focus on the interesting case when the two securities have different Sharpe ratios – for this comparison, we set the expected return of the illiquid asset to $\nu = 0.2$. In the Merton case where both assets are fully liquid: varying ρ from zero to one leads to large swings in portfolio allocations. As ρ approaches one, the investor takes large offsetting positions in the two assets that tend to plus or minus infinity. In Panel G of Figure 1 we compare the target allocation ξ^* as a function of the correlation coefficient. We see that the effect of correlation is significantly muted relative to the Merton benchmark. The investor never shorts the liquid risky asset even when the correlation approaches one. From the investor's perspective, the liquid and illiquid asset are imperfect substitutes, since only the former can be used to fund short-term consumption. This imperfect substitutability decreases the desire to use the liquid risky asset to hedge price changes in the illiquid asset.

4.5 Consumption

The investor's optimal consumption choice satisfies

$$U'(C) = F_W(W, X). \tag{17}$$

In the short run, consumption is funded by liquid assets. Hence, the investor equates the marginal utility of consumption with the marginal value of *liquid*, rather than *total* wealth. Using the form for the value function (10), the ratio of consumption to liquid wealth equals

$$c = \left((1-\gamma)H(\xi) - H'(\xi)\xi\right)^{-\frac{1}{\gamma}} (1-\xi)^{-1}. \tag{18}$$

Panel H of Figure 1 plots the agent's optimal consumption to total wealth ratio, $c(1-\xi)$, as a function of the current allocation in illiquid assets ξ. The investor always consumes a lower fraction of her total wealth than the two-asset Merton benchmark. Further, in contrast to the Merton benchmark, the consumption to wealth ratio is time-varying, since the marginal value of liquid wealth varies with the current allocation to the illiquid security.

The investor's consumption policy sheds further light on the behavior of the marginal value of liquid wealth F_W. When the share of illiquid assets in the portfolio is small, the share of total wealth consumed is insensitive to portfolio composition ξ; the investor smooths lifetime consumption by consuming a higher fraction of liquid wealth today as ξ increases. In contrast, as the share of illiquid assets ξ increases towards one, her marginal value of liquid wealth increases, leading to a lower consumption to total wealth ratio.

In Table 3 we compute the average consumption rate $E[c(1-\xi)]$ for different values of λ. As we vary the expected time until the next trading opportunity $1/\lambda$ from 1 week ($1/\lambda = 1/50$) to 10 years, the fraction of total wealth consumed per year declines from 8.8% to 5.9%. Interestingly, when the average length of the illiquidity period is sufficiently long ($\lambda \leq 1/4$), the investor consumes a lower fraction of her total wealth than an investor who is unable to trade the second asset at all. The presence of illiquidity risk constrains how an investor can fund consumption. Since consumption must be met out of liquid wealth, the greater the proportion of illiquid assets or the longer the times between liquidity events, the lower the optimal consumption.

4.6 The Cost of Illiquidity

To quantify the cost of illiquidity, we compute the fraction of initial wealth α the investor would be willing to give up at the instant of the liquidity event, in order to be fully able to trade the illiquid asset

$$K_{M2}\left((W_t + X_t)(1-\alpha)\right)^{1-\gamma} = (W_t + X_t)^{1-\gamma} H(\xi^*). \tag{19}$$

The left hand side of equation (19) is the value function of a Merton investor able to invest in two risky securities. We refer to α as the utility or certainty equivalent cost of illiquidity.

In Table 3 we compute the certainty equivalent cost of illiquidity for different values of λ. The cost of illiquidity can be substantial. Even when the investor can trade on average once a week ($\lambda = 50$), she would be willing to forego 1.8% of her total wealth in order to make the second asset fully liquid. For higher degrees of illiquidity, the cost increases substantially; an investor trading an asset with an average of 10 years between trades ($\lambda = 0.1$), such as institutional real estate, would

be willing to give up 22% of her total wealth in order to be able to continuously trade the illiquid asset. Here, we should emphasize that α is a conservative estimate of the cost of illiquidity. Since $H(\xi^*) \geq H(\xi_t)$, the investor would be willing to pay *at least* a fraction α at any point between liquidity events. Replacing $H(\xi^*)$ with its long-run average $E[H(\xi_t)]$ leads to higher costs.

5 Determinants of the Utility Cost of Illiquidity

In this section we explore the key determinants of the utlity cost of illiquidity. First, we separate the cost due to suboptimal diversification from the inability to fund consumption; we show that smoothing consumption is more important than maintaining optimal diversification. Second, we disentangle the effect of illiquidity from illiquidity risk by comparing our setup to a model with deterministic periods of illiquidity; we find that the uncertainty about the frequency of trade magnifies the utility cost of illiquidity. Third, we verify that our results are not driven by the fact that marginal utility is infinite in the extreme states where liquid wealth – and consumption – drops to zero, by allowing the investor to pay a fixed cost to transact immediately. Last, we explore the impact of preference parameters, disentangling the effect of the coefficient of risk aversion from the elasticity of intertemporal substitution; the utility cost of illiquidity is highest for investors who are willing to substitute across states but not across time.

5.1 The Effect of Consumption Smoothing

Illiquidity impedes consumption smoothing and optimal diversification. To separate these two effects, we consider an investor who only values consumption at some future stochastic terminal date τ

Problem 2 (No Intermediate Consumption) *The investor maximizes*

$$F_{nc}(W_t, X_t) = \max_{\{\theta,\, I\}} \mathrm{E}_t \left[U(C_{\hat{\tau}}) \right], \tag{20}$$

where $\hat{\tau}$ is a stochastic retirement time that is exponentially distributed according to a Poisson process with arrival rate δ subject to the budget constraints given by

$$\frac{dW_t}{W_t} = \left(r + (\mu - r)\, \theta_t \right) dt + \theta_t \sigma dZ_t^1 - \frac{dI_t}{W_t}, \tag{21}$$

and equation (6), with $C_t = W_t + X_t$. Re-balancing ($dI_t \neq 0$) occurs only when the Poisson process N_t^λ jumps.

The following proposition characterizes the solution to Problem 2

Proposition 4 (No Intermediate Consumption) *In Problem 2, the investor's value function can be written as* $F_{nc}(W,X) = (W+X)^{1-\gamma} H_{nc}(\xi)$*, where* $H_{nc}(\xi)$ *exists and is finite, continuous, and concave for* $\xi \in [0,1]$*.* $H_{nc}(\xi)$ *obtains its maximum for some* $\xi \in [0,1]$*. When a trading opportunity occurs at time* τ*, the trader selects* I_τ *so that* $\frac{X_\tau}{X_\tau + W_\tau} = \xi^*_{nc}$*, where* H^*_{nc} *and* ξ^*_{nc} *are defined as in Proposition 2. In addition,* $H_{nc}(\xi)$ *is characterized for* $t \leq \hat{\tau}$ *by*

$$0 = \max_\theta \left[\delta \left(\frac{1}{1-\gamma} - H_{nc}(\xi) \right) + \lambda \left(H^*_{nc} - H_{nc}(\xi) \right) + H_{nc}(\xi)\, A(\xi, 0, \theta) + \frac{\partial H_{nc}(\xi)}{\partial \xi} B(\xi, 0, \theta) + \frac{1}{2} \frac{\partial^2 H_{nc}(\xi)}{\partial \xi^2} C(\xi, 0, \theta) \right]. \tag{22}$$

where the functions A*,* B *and* C *are defined in* (13)*.*

The differential equation characterizing the solution to the problem without intermediate consumption is similar to our baseline model, up to a difference in the discount rate. To facilitate comparison with the baseline case, we consider the case where the investor's effective rate of impatience is equal to our baseline calibration, $\delta = \beta$.

In Table 4, we compute optimal policies and utility costs across different levels of illiquidity $1/\lambda$. Absent the motive to smooth intermediate consumption, the effects of illiquidity are quantitatively small. The illiquidity cost in terms of certainty equivalent wealth is below 40 basis points across all values of λ. Even for an average time between liquidity events of 10 years, the optimal holdings in the illiquid asset are 0.52, compared to 0.59 for the Merton case. This compares to 0.05 in Table 3 with intermediate consumption for the same average times between liquidity events. Thus, the inability to fully smooth consumption across states is the primary determinant of the cost of illiquidity, which confirms our intuition in Section 4.2.

5.2 Stochastic Versus Deterministic Trading Opportunities

To disentangle the effect of the *length* of the illiquid period from the *uncertainty* over its duration, we consider the case where the agent is allowed to re-balance her portfolio at fixed intervals, spaced T periods apart. The investor's problem is

Problem 3 (Deterministic Liquidity) *The investor maximizes* (9)*, subject to the budget constraints (5) and (6), with re-balancing* ($dI_t \neq 0$) *only at the deterministic times* $\tau = 0, T, 2T, \ldots$*.*

The following proposition characterizes the solution to Problem 3:

Proposition 5 (Deterministic Liquidity) *For Problem 3, the investor's value function can be written as* $F_T(t,W,X) = (W+X)^{1-\gamma} H_T(t,\xi)$*, where* $H_T(t,\xi)$ *exists and is finite, continuous, and*

concave in ξ *for* $\xi \in [0,1)$. $H_T(t,\xi)$ *obtains its maximum in* ξ *for some* $\xi \in [0,1)$. *We have* $H_T(\tau,\xi) = \max_\xi \lim_{\epsilon \downarrow 0} H_T(\tau+\epsilon,\xi)$, *and we define* $\xi_T^* = \arg\max_\xi \lim_{\epsilon\downarrow 0} H_T(\tau+\epsilon,\xi)$ *with* $\tau = 0, T, 2T, ...,$ *the repeated trading times. At* $t=\tau$, *the investor selects* I_τ *so that* $\frac{X_\tau}{X_\tau+W_\tau} = \xi_T^*$.[17] *The function* $H_T(t,\xi)$ *is characterized by*

$$0 = \max_{c,\theta} \left\{ \frac{1}{1-\gamma} c^{1-\gamma} (1-\xi)^{1-\gamma} - \beta H_T(t,\xi) + \frac{\partial H_T(t,\xi)}{\partial t} + H_T(t,\xi)\, A(\xi,c,\theta) \right.$$
$$\left. + \frac{\partial H_T(t,\xi)}{\partial \xi} B(\xi,c,\theta) + \frac{1}{2} \frac{\partial^2 H_T(t,\xi)}{\partial \xi^2} C(\xi,c,\theta) \right\}. \tag{23}$$

where the functions A, B *and* C *are defined in* (13).

Table 5 computes the optimal policies and the utility cost of illiquidity for different lengths of the illiquidity period. In contrast to the case with stochastic trading opportunities (see Table 3), varying the length of the deterministic illiquidity period has only a small effect on optimal policies. For example, when the time until the next trade is known in advance, varying the expected time until the next liquidity date from 1/50 to 10 years leads to a drop in the fraction of total wealth consumed per year from 8.9% to 8.4%. Similarly, the effects on welfare are small and are relative insensitive to the length of the illiquidity period; the utility cost of illiquidity varies between 1.1% (1/50 years) to 2.8% (10 years).

Comparing Table 5 to Table 3, we conclude that the uncertainty regarding the opportunity to trade is a major component of the utility cost of illiquidity. When the trading intervals are known in advance, lengthening the intervals of non-trading has a small impact on investor utility. The investor's main concern is to avoid states of the world where her liquid wealth – and therefore her consumption – drops to zero before the next opportunity to trade. If the investor can trade at deterministic intervals, this state can be avoided with probability one by investing an appropriate amount into the riskless asset and consuming a constant fraction. That is, the investor can hedge against deterministic illiquidity. Illiquidity risk represented by stochastic non-trading intervals – the unknown time until the next liquidity event – is unhedgeable and induces large portfolio choice effects.

5.3 Introducing Costly Liquidity

Our model can be interpreted as a setting where the cost of transacting is infinite, except at times when the Poisson process hits, in which case it is equal to zero. Effectively, we are assuming that sometimes there is no available counterparty with which to trade, at any price. However, this

[17] The limit statement reflects the fact that the value function is continuous in ξ but not in t at $t=\tau$: at $t=\tau+\epsilon$, ξ is a state variable, and so $H(\tau+\epsilon,\xi)$ is discretely less than $H(\tau,\xi)$, except for $\xi = \xi_T^*$.

assumption may be extreme: often, in addition to a decentralized market, there are certain trading partners that are always available, but at a cost. Here, we explore a hybrid of our baseline model and a model with transaction costs. Specifically, the investor has two options if she wants to re-balance her portfolio, she can i) wait for the arrival of the Poisson process, just like the baseline model, or ii) pay a cost and re-balance freely. We assume that this fixed fee is independent of the trading amount and scales with the investor's total wealth.

Problem 4 (Fixed Transaction Cost) *The investor maximizes* (9)*, subject to the budget constraints (5) and (6), with re-balancing ($dI_t \neq 0$) only when the Poisson process N_t^λ hits or upon the payment of a fixed fraction κ of total wealth $W + X$.*

The following proposition characterizes the solution to Problem 4:

Proposition 6 (Fixed Transaction Cost) *In Problem 4, the investor's value function can be written as $F_{fc}(W, X) = (W + X)^{1-\gamma} H_{fc}(\xi)$, where $H_{fc}(\xi)$ exists and is finite, continuous, and concave in $\xi \in \left[\underline{\xi}, \overline{\xi}\right]$. Define H_{fc}^* and ξ_{fc}^* as in Proposition 2. When ξ hits the boundaries $[\underline{\xi}, \overline{\xi}]$, the trader pays $\kappa(W + X)$ and selects I_τ so that $\frac{X_\tau}{X_\tau + W_\tau} = \xi_{fc}^*$. In the no-trade region, $\xi \in \left[\underline{\xi}, \overline{\xi}\right]$, $H_{fc}(\xi)$ solves*

$$0 = \max_{c,\theta} \left[\frac{1}{1-\gamma} c^{1-\gamma} (1-\xi)^{1-\gamma} - \beta H(\xi) + \lambda \left(H_{fc}^* - H(\xi) \right) + H(\xi) A(\xi, c, \theta) + \frac{\partial H(\xi)}{\partial \xi} B(\xi, c, \theta) \right.$$
$$\left. + \frac{1}{2} \frac{\partial^2 H(\xi)}{\partial \xi^2} C(\xi, c, \theta) \right]. \tag{24}$$

where the functions A, B and C are defined in (13) *and the boundaries of the no-trading region, $[\underline{\xi}, \overline{\xi}]$, solve $H(\underline{\xi}) = H(\overline{\xi}) = (1-\kappa)^{1-\gamma} H_{fc}^*$.*

We choose a fixed, rather than proportional, transaction cost for several reasons. First, the fixed costs of trading is the closest analogue to our baseline model, because it leads to very similar policy functions.[18] Second, the fixed cost can be interpreted as the investor paying attention to look for a counterparty in the market for the illiquid asset, in the spirit of Abel, Eberly, and Panageas (2007). Effectively, we are endowing the investor with the option to pay a fee to temporarily increase the

[18] The portfolio policies take the same form in the baseline model and the fixed transaction cost model. Intuitively, the investor must wait for a liquidity event to trade in the baseline model. Then, the investor can re-balance his portfolio freely and chooses allocation ξ^*. From Proposition 6, the fixed cost operates in a similar fashion. Once the wealth allocation process exits the set $\left[\underline{\xi}, \overline{\xi}\right]$, the investor can re-balance freely to any portfolio weight after paying the fixed cost – and therefore re-balances to the same point ξ^*. Note that $\xi^* \in \left[\underline{\xi}, \overline{\xi}\right]$, so after hitting the boundary the investor re-balances once and it will be some time before hitting the boundary again. In a proportional costs model, the investor's portfolio holdings also lie in a bounded set (see, for example, Constantinides (1986)). However, under proportional transaction costs, the investor transacts the minimal amount necessary to keep the portfolio weight within the boundary, which implies continuous trading.

arrival rate λ to infinity. Third, the numerical solution to hybrid model with proportional transacion costs is more complex and numerically challenging. We expect, however, the underlying economics between fixed and proportional transaction costs is similar: the investor always wishes to avoid high ξ states – which are states with low liquid wealth and therefore high marginal utility. The investor wishes to avoid these states by paying a cost, and whether the cost is fixed or proportional is secondary.

Table 6 computes the optimal policies and the utility cost of illiquidity for different levels of illiquidity λ and transaction cost κ. Comparing Table 6 to Table 3, we see that, the availability of the option to freely trade upon payment of the cost κ dampens the effect of illiquidity. For example, in the case where $\lambda = 1$, varying the level of the fixed cost from infinite to 1% of total wealth, the utility cost falls from 6.7% to 2.9%. In contrast to our baseline model, here the investor has the option to trade at will; hence she trades only when her illiquid allocation ξ exits the inaction region $[\underline{\xi}, \overline{\xi}]$, that is, only when the cost of illiquidity is high. Doing so allows the investor to eliminate the occurrence of states when her liquid wealth – and therefore her consumption – drops to zero. This option is exercised more frequently if the probability of otherwise finding a counterparty, λ, or the exercise cost κ, is low.

In the last row of each panel in Table 6, we compute optimal policies and the utility cost of illiquidity in the case where $\lambda = 0$. Doing so allows us to compare our baseline setup with a pure transaction costs model. Our baseline model and the pure transaction costs model lead to similar qualitative predictions as our model – the investor trades infrequently. As a comparison of the utility costs across the two models, an investor facing liquidity events arriving on average once a year ($\lambda = 1$) would be willing to pay a fraction $\kappa = 2.6\%$ of her total wealth to an intermediary each time to transact freely. This extension illustrates that the substantial utility costs are not driven by the extreme behavior of marginal utility when liquid wealth drops to zero.

We conclude that endowing that investor with the option to trade at will, though at a cost, lowers the cost of illiquidity. If investors have access to a market maker who is able to absorb the illiquid asset for a fee, this hybrid model is more appropriate. However, in many situations, liquidity is not always available, even at a cost. Adverse selection or search frictions may lead to the inability to trade immediately at any price. In the financial crisis, many markets experienced liquidity freezes where no transactions were possible at any price (see Tirole (2011)). In these cases, our baseline model is more appropriate than transaction costs notions of illiquidity.

5.4 Intertemporal Substitution and Risk Aversion

The curvature of the value function is an important determinant of the cost of illiquidity. In particular, the investor fears the probability that she reaches states with very low liquid wealth for two reasons. First, states with low liquid wealth are states with low consumption, and the investor likes to smooth consumption across *states*; the coefficient of risk aversion captures the magnitude of this preference. Second, in states where liquid wealth is low relative to total wealth, the investor faces a steeply increasing consumption profile. The agent dislikes these states because she wants to have smooth consumption paths over *time*; the elasticity of intertemporal substitution (EIS) governs this preference. A feature of time-separable preferences is that these two effects are linked. To investigate these two motives separately, we consider the case where the agent has recursive preferences.

Problem 5 (Epstein-Zin) *The investor maximizes*

$$F_{ez}(W_t, X_t) = \max_{\{\theta,\, I,\, c\}} \mathrm{E}_t \left[\int_t^\infty f\Big(C_s, F_{ez}(W_s, X_s)\Big) ds \right], \tag{25}$$

where the aggregator f is defined following Duffie and Epstein (1992) as

$$f(C, J) \equiv \frac{\beta}{1-\zeta} \left(\frac{C^{1-\zeta}}{((1-\gamma)J)^{\frac{\gamma-\zeta}{1-\gamma}}} - (1-\gamma)\, J \right), \tag{26}$$

subject to the budget constraints (5) and (6), with re-balancing ($dI_t \neq 0$) only when the Poisson process N_t^λ jumps.[19]

The following proposition characterizes the solution to Problem 5:

Proposition 7 (Epstein-Zin) *For Problem 5, the investor's value function can be written as $F_{ez}(W, X) = (W+X)^{1-\gamma} H_{ez}(\xi)$, where $H_{ez}(\xi)$ exists and is finite, continuous, and concave for $\xi \in [0, 1)$. $H_{ez}(\xi)$ obtains its maximum for some $\xi \in [0, 1)$. Define H_{ez}^* and ξ_{ez}^* as in Proposition 2. When a trading opportunity occurs at time τ, the trader selects I_τ so that $\frac{X_\tau}{X_\tau + W_\tau} = \xi_{ez}^*$. In addition, $H_{ez}(\xi)$ is characterized by*

$$\begin{aligned} 0 \;=\; \max_{c,\,\theta} \Bigg\{ & \frac{\beta}{1-\zeta} \left(c^{1-\zeta} (1-\xi)^{1-\zeta} \Big((1-\gamma) H_{ez}(\xi) \Big)^{\frac{\zeta-\gamma}{1-\gamma}} - (1-\gamma) H_{ez}(\xi) \right) + \lambda \left({H_{ez}}^* - H_{ez}(\xi) \right) \\ & + H_{ez}(\xi)\, A(\xi, c, \theta) + \frac{\partial H_{ez}(\xi)}{\partial \xi} B(\xi, c, \theta) + \frac{\partial^2 H_{ez}(\xi)}{\partial \xi^2} C(\xi, c, \theta) \Bigg\}. \end{aligned} \tag{27}$$

where the functions A, B and C are defined in (13).

[19] See Duffie and Epstein (1992) for more details. Here, β is the subjective discount rate, γ is the coefficient of risk aversion and ζ is the inverse of the elasticity of intertemporal substitution. The case of power utility corresponds to $\gamma = \zeta$.

To quantify the impact of risk aversion and the elasticity substitution on portfolio policies, we compute optimal policies and utility costs for different preference parameters and levels of illiquidity $1/\lambda$. We show the results in Table 7. In Panel A, we vary the coefficient of risk aversion γ and the likelihood of trading λ, holding the EIS fixed. As risk aversion increases, the investor's policies (ξ^*), her rate of consumption out of liquid wealth (c), and the fraction of her liquid wealth invested in the liquid asset (θ) all converge to the frictionless benchmark. For instance, in the case where $\lambda = 10$, an investor with a risk aversion of $\gamma = 3$ allocates 80% of her total wealth to the illiquid asset, compared to 118% for a Merton investor. In contrast, if the investor had a risk aversion of $\gamma = 15$, she would allocate only 22% of her total wealth to the illiquid asset, compared to 24% for a Merton investor. An investor with high risk aversion chooses to invest less in risky assets, hence for her, the cost of illiquidity is small. Indeed, as we see in Panel A.*ii*, the utility cost of illiquidity *decreases* with risk aversion.

In Panel B, we vary the elasticity of intertemporal substitution $1/\zeta$ and the likelihood of trading λ, holding risk aversion fixed. The elasticity of intertemporal substitution has a quantitatively small impact on portfolio policies. As we see in Panels B.*i* and B.*iv*, varying the EIS from 1.5 to 1/6 has essentially no impact on the allocation to the liquid asset, and a small impact on the allocation to the illiquid asset. However, as we see in Panel B.*iii*, varying the EIS has an impact on the investor's optimal consumption policy. When the investor's desire to smooth consumption across states increases ($EIS = 1/6$), her optimal fraction of wealth consumed declines faster with λ than when her elasticity of substitution is high. Hence, we find that low EIS magnifies the utility cost of illiquidity to the investor, as we see in Panel B.*ii*.

We conclude that the utility cost of illiquidity is higher for agents with low inter-temporal elasticity of substitution and *low* risk aversion. Holding portfolio allocations constant, the cost of illiquidity is higher for more risk averse agents that are also reluctant to substitute across time. However, this utility cost depends on the fraction of wealth invested in illiquid securities, and the amount of investment in the illiquid risky asset is decreasing in risk aversion.

6 Liquidity Crises

Here, we extend the model to a setting where financial markets are normally liquid, but transition to infrequent liquidity crises. These liquidity crises are temporary, and they adversely affect the liquidity of otherwise liquid securities. We consider two applications of the extended model. First, we show that the possibility of a deterioration in market liquidity leads to limited arbitrage in normal times. Second, using the investor's marginal value of wealth, we derive the price of illiquidity risk.

6.1 A Model With Systematic Liquidity Risk

The level of market liquidity depends on two states, $S_t \in \{I, L\}$. In state L, corresponding to 'normal' times, all assets are perfectly liquid, as in the Merton benchmark. In state I – the 'crisis' state – the investor needs to wait for the arrival of a trading opportunity to trade the now illiquid asset P, as in the model in Section 3. The state of market liquidity, S_t, follows a continuous-time Markov process with transition probability matrix between time t and $t + dt$ given by

$$P = \begin{pmatrix} 1 - \chi^L \, dt & \chi^L \, dt \\ \chi^I \, dt & 1 - \chi^I \, dt \end{pmatrix}. \tag{28}$$

Hence, the frequency and average duration of a liquidity crisis are χ_I and $1/\chi_L$ respectively.

The investor's problem is

Problem 6 (Liquidity Crises) *The investor maximizes (9) subject to the budget constraints (5) and (6). The state of the economy ($S_t \in \{I, L\}$) evolves as in (28). If $S_t = L$, trade in both assets is continuous; if $S_t = I$, the investor can re-balance ($dI_t \neq 0$) only when the Poisson process N_t^λ jumps.*

Our first result is that even in normal times, $S = L$, the investor will not short the *potentially* illiquid asset or have a short position in liquid wealth:

Proposition 8 *Any optimal policies in Problem 6 will have $W > 0$ and $X \geq 0$ a.s. for both $S = I$ and $S = L$.*

The possibility of a liquidity crisis affects portfolio policies in normal times. The transition from liquid to illiquid is a surprise event and occurs without the opportunity to re-balance. Consequently, the portfolio restrictions from the illiquid state – see Proposition 1 – also apply in the liquid state.

The following proposition characterizes the solution to Problem 6:

Proposition 9 (Liquidity Crises) *For Problem 6, the investor's value function can be written as $F_{LC}(W, X, S) = (W + X)^{1-\gamma} H_{LC}(\xi, S)$, where $H_{LC}(\xi)$ exists and is finite, continuous, and concave in ξ for $\xi \in [0, 1)$. $H_{LC}(\xi)$ obtains its maximum for some $\xi \in [0, 1)$. The function H_{LC} is given by*

$$H_{LC}(\xi, S) = \begin{cases} H_I(\xi), & S = I \\ H_L^*, & S = L \end{cases}, \tag{29}$$

where the function $H_I(\xi)$ satisfies the Hamilton-Jacobi-Bellman equation

$$\begin{aligned} 0 = \max_{c,\theta} \Bigg\{ & \frac{1}{1-\gamma} c^{1-\gamma} (1-\xi)^{1-\gamma} - \beta H_I(\xi) + \lambda \left(H_I^* - H_I(\xi) \right) + \chi^L \left(H_L^* - H_I(\xi) \right) \\ & + H_I(\xi)(1-\gamma) A(\xi, c, \theta) + \frac{\partial H_I(\xi)}{\partial \xi} B(\xi, c, \theta) + \frac{1}{2} \frac{\partial^2 H_I(\xi)}{\partial \xi^2} C(\xi, c, \theta) \Bigg\}, \end{aligned} \tag{30}$$

and the functions A, B and C are defined in (13)*. The constants* H_L^* *and* H_I^* *solve*

$$0 = \max_{c,\theta,\xi} \left\{ \frac{1}{1-\gamma} c^{1-\gamma} (1-\xi)^{1-\gamma} - \beta H_L^* + \chi^I (H_I(\xi) - H_L^*) + A\left(\xi, \frac{c}{1-\xi}, \frac{\theta}{1-\xi}\right) H_L^* \right\} \tag{31}$$

$$H_I^* = \max_{\xi} H_I(\xi) \tag{32}$$

The policies $\{c_I^*, \theta_I^*\}$ *and* $\{c_L^*, \theta_L^*, \xi_L^*\}$ *maximize (30) and (31) respectively. The policy* ξ_I^* *maximizes (32).*

In this case, the investor's value function depends not only on her wealth composition ξ, but also on the condition of market liquidity. In normal times, $S = L$, the investor can freely re-balance between both risky assets and thus the function $H_L(\xi)$ is a constant. During a liquidity crisis, $S = I$, the investor's problem is similar to the problem analyzed in Section 3, and, therefore, her value function H_I depends on the ratio of illiquid to total wealth ξ. When a liquidity crisis occurs, the investor is constrained to hold her current allocation to the now illiquid security until the next opportunity to trade. Hence, the investor's optimal portfolio holdings in the liquid state are affected by the possibility of a liquidity crisis.

In Figure 2, we compare portfolio policies across regimes for different values of the frequency, χ_I, average duration $1/\chi_L$ and severity $1/\lambda$ of liquidity crisis. As we see, the investor reduces her allocation in illiquid asset not only during a crisis, but also during normal times. Similarly, the investor reduces her consumption rate in both regimes. Both of these effects are increasing in χ_I, $1/\chi_L$ and $1/\lambda$. In addition, the investor holds fewer liquid risky assets, but only during the liquidity crisis; her portfolio allocation in liquid risky assets is the same as the Merton benchmark in normal times, assuming the two assets S and P are uncorrelated.

In summary, the possibility of a liquidity crisis leads to underinvestment in assets that are currently fully liquid but whose liquidity can dry up during a crisis. The same mechanism leads to limited arbitrage, which we explore below.

6.2 Limits to Arbitrage

In the absence of any trading friction, the existence of two perfectly correlated securities with different Sharpe ratios implies an arbitrage opportunity. Faced with this arbitrage, the investor should construct a zero-investment portfolio that has a positive payoff, and take an infinite position in this strategy. In our setting, the investor is reluctant to fully invest in arbitrage opportunities that involve potentially illiquid securities – even if both securities are currently fully liquid and taking advantage of this arbitrage involves no short positions in risky assets.

Corollary 10 (Limits to Arbitrage) *If* $|\rho| = 1$ *and* $\frac{\nu - r}{\psi} \neq \frac{\mu - r}{\sigma}$ *the investor's portfolio policies* θ_L^* *and* ξ_L^* *are finite and satisfy*

$$\frac{\nu - r}{\psi} - \rho \frac{\mu - r}{\sigma} = -\frac{\chi^I}{\psi} \frac{H_I'(\xi_L^*)}{H_L^*(1-\gamma)} \tag{33}$$

$$\xi_L^* \in [0, 1) \tag{34}$$

$$\theta_L^* = \frac{\mu - r}{\gamma \sigma^2} - \rho \frac{\psi}{\sigma} \xi_L^* \tag{35}$$

Limits to arbitrage arise naturally in our setting. In the event of a crisis, the investor is not able to continuously re-balance her position. Hence, even if both securities are currently fully liquid, undertaking the arbitrage exposes the investor to illiquidity risk. In equation (33), the investor will increase her holdings of the illiquid asset until the marginal welfare loss in the illiquid state – determined by $H_I'(\xi_L^*)$ – times the probability of that state occurring is proportional to the difference in the Sharpe ratios between the liquid and illiquid assets. Examining the investor's allocation to the liquid asset (35), we see two components: the first part depends on the Sharpe ratio of the liquid risky asset; the second part hedges fluctuations to her wealth due to the investment in the potentially illiquid security. Since ξ_L^* is finite, her overall portfolio is *not* riskless.

The arbitrageur's investment in this apparent arbitrage opportunity is limited by her wealth. The investor will never take a levered position in the illiquid asset during a liquidity crisis – see proposition 1 – since doing so would lead to states with zero consumption. The inability to reduce leverage at the onset of a liquidity crisis – since the illiquid asset cannot be traded immediately – dissuades her from leveraging her *potentially illiquid* investment in normal times, as we see in equation (34). Consequently, the amount of resources the agent commits to an 'arbitrage opportunity' will be bounded above by her level of wealth, leading to limited arbitrage.

To quantify the magnitude of limited arbitrage in our setting, we compute the optimal portfolio policies θ_L^* and ξ_L^* in a setting with an arbitrage opportunity. We assume the two risky assets are perfectly correlated, $|\rho| = 1$, and we set the mean return to the potentially illiquid asset to be two percentage points higher than the mean return of the liquid asset, $\nu = 0.14$. As we see in Figure 3, the leverage constraint (34) is not binding. Even when the severity of the liquidity crisis is small – the potentially illiquid asset can be traded on average once a month – the investor only allocates 95% of her wealth into the potentially illiquid security; as the severity of the crisis increases to $1/\lambda = 2$ years, the investor allocates just 60% of her wealth in the apparent arbitrage. An increase in the frequency, or the average duration of a crisis has similar qualitative results.

We conclude that the possibility of a liquidity crisis leads to limited arbitrage in normal times. Our mechanism is distinct from existing models; limits to arbitrage arise even in a state of the world

where all securities are fully liquid, there are no transaction costs, and riskless profits are possible. To the extent that there is a negative relation between the degree of aggressiveness in arbitrageurs' strategies and the degree of mispricing, our model suggests that mispricing is worse in markets where liquidity crises are more likely, last longer, and are more severe.

6.3 The Pricing of Illiquidity Risk

A model with systematic liquidity crises allows us to quantify the illiquidity *risk* premium, that is, the expected return differential between two *liquid* securities with heterogeneous return exposure to illiquidity crises. Using an approach similar to Merton's ICAPM, we derive the price of the risk of a liquidity crisis. To compute this risk premium, we examine the effect of a rise of illiquidity on the investor's marginal value of liquid wealth F_W taking the processes governing asset returns as given. During normal times, that is, in the liquid state $S = L$, the investor's marginal value of wealth process $\pi_t \equiv F_W$ evolves according to

$$\frac{d\pi_t}{\pi_t} = [\dots]\, dt + [\dots]\, dZ_t^1 + [\dots]\, dZ_t^2 + \frac{F_W(W,X,I) - F_W(W,X,L)}{F_W(W,X,L)}\Big|_{\xi=\xi_L^*} dN_t^I, \tag{36}$$

where N_t^I is a Poisson count process such that $dN_t^I = 1$ at the onset of a liquidity crisis, $S = I$. The last term in (36) corresponds to the increase in the investor's marginal utility at the onset of a financial crisis. This term determines the market price of illiquidity risk.

The marginal utility process (36) provides a theoretical justification for empirical specifications of the stochastic discount factor (SDF) that include measures of systematic illiquidity risk (e.g. Pastor and Stambaugh, 2003; Sadka, 2006; Korajczyk and Sadka, 2008). In particular, several of the illiquidity measures considered in the literature are likely to be correlated with an increase in the difficulty of finding counterparties to trade. For instance, Korajczyk and Sadka (2008) show comovement among several measures of liquidity; these measures include the average turnover in the stock market. A more direct test of the SDF (36) would include the common component of illiquidity measures across asset markets. Hu, Pan, and Wang (2012) provide some empirical evidence along these lines.

To quantify the magnitude of the illiquidity risk premium, we introduce a derivative security in zero net supply that allows the investor to hedge a deterioration in market liquidity, which we call illiquidity protection. By purchasing illiquidity protection, the investor pays an annual premium equal to $\hat{\chi}_I$ in order to receive a cash payment of \$1 dollar in the event of a liquidity crisis. The following corollary computes the cost of illiquidity insurance that would induce zero demand for the derivative security

Corollary 11 *The annual premium for liquidity protection is equal to*

$$\hat{\chi}_I = \chi_I \frac{F_W(W, X, I)}{F_W(W, X, L)}\Big|_{\xi=\xi_L^*} \tag{37}$$

Corollary 11 shows that the cost of illiquidity insurance is equal to the probability of a liquidity crisis, times the increase in the marginal value of liquid wealth on the event of a crisis. Since liquid wealth becomes more valuable during a liquidity crisis than normal times, the investor is willing to pay a higher rate that the objective probability χ_I to obtain some protection against a liquidity crisis.

The illiquidity risk premium is related to the utility cost of illiquidity. In Panel A of Figure 4 we compare the risk premium $\hat{\chi}_I - \chi_I$ of a liquidity crisis across different values of χ_I, χ_L and λ. We see that the investor is willing to pay a substantial premium over the subjective probability χ_I in order to obtain liquidity during a crisis. For example, the investor would be willing to pay an excess premium of 80 bps per year to obtain liquidity on the event of a once in a decade liquidity crisis ($\chi_I = 0.1$), with average duration of two years ($\chi_L = 0.5$), during which liquidity events arrive on average once a year ($\lambda = 1$).

In Panel B of Figure 4, we compute the utility cost of a liquidity crisis, defined as the fraction of total wealth the investor would pay *ex-ante* to eliminate the possibility of a crisis. Even though illiquidity crises are temporary, they still lead to substantial utility costs. For instance, the investor would be willing to forgo 2% of her wealth to eliminate the possibility of a once in a decade liquidity crisis ($\chi_I = 0.1$), with average duration of two years ($\chi_L = 0.5$), during which the investor can re-balance on average once a year ($\lambda = 1$). Comparing Panels A and B, we see that the illiquidity *risk* premium is related to the utility cost of illiquidity. Changes in the model specification or parameters that amplify the utility cost of illiquidity – for instance those considered in Section 5 – also lead to a higher illiquidity risk premium.

The risk premium of the security offering illiquidity protection helps us understand differences in average return among similar, liquid securities, that have different price behavior during a liquidity crisis. A classic example is the swap-treasury spread. Both securities are very liquid and have similar exposures to interest rate and credit risk; yet, swaps have been historically priced cheaper than treasuries. This difference in price could be due to their differential price reaction to changes in the level of market illiquidity. For instance, during the flight-to-quality episodes that followed the financial crisis of 2008, the swap spread increased and stayed high for some time. Similar behavior was observed for mortgage spreads, especially for the riskiest parts of the mortgage market. Our model provides a framework to understand the spread differential during a liquidity crisis, by relating

the price of illiquidity risk to the marginal utility of financial market participants.

7 Conclusion

We study the effect of illiquidity risk on portfolio choice by extending the Merton (1971) framework to allow for infrequent and stochastic trading opportunities. Relative to the Merton economy, illiquidity leads to a large reduction in the allocation to both illiquid and liquid assets, lower consumption rates and time-varying effective relative risk aversion. There are two main drivers of these results. First, consumption is financed through liquid assets. Investors care about both liquid and illiquid wealth, and as illiquid wealth becomes larger, the investor endogenously acts in a more risk-averse fashion fearing states with low liquid wealth. Second, the fact that the duration of the illiquidity period is uncertain greatly amplifies the cost of illiquidity. In contrast to models with deterministic trading dates, the investor cannot hedge against the likelihood that her liquid wealth – and therefore her intermediate consumption – drops to zero.

We study the pricing of liquidity crises by allowing the risk of illiquidity to vary over time. Motivated by the behavior of many asset markets that exhibit periodic pronounced periods where liquidity "dries up" , we extend the model to allow for multiple liquidity regimes. We consider the case where all assets are fully liquid during normal times, but there exist temporary regimes during which the illiquid asset can be traded only infrequently. Our model allows for the possibility of 'arbitrage opportunities' which occur when all assets are perfectly liquid, but agents do not take advantage of them due to the possibility that liquidity will evaporate. Our calibration implies that agents would be willing to pay an illiquidity risk premium of 2% to insure against illiquidity crises occurring once every ten years.

A Proofs and Derivations

A.1 Proof of Proposition 1

The value function is bounded below by the problem in which the illiquid asset does not exist, and it is bounded above by the problem in which the entire portfolio can be continuously rebalanced: the Merton one- and two-stock problems. Hence, there exist constants K_{M1} and K_{M2} such that

$$K_{M1}W^{1-\gamma} \leq F(W,X) \leq K_{M2}\left(W+X\right)^{1-\gamma} \leq 0. \tag{A.1}$$

Combining (A.1) and (10) yields that $H(\xi)$ exists and is finite for $\xi \in [0,1)$.

Consumption is out of liquid wealth only and the illiquid asset cannot be pledged, so $W_t \leq 0$ implies zero consumption before the next trading day, leaving the objective function (4) at $-\infty$. For $|\rho| < 1$, $X_t < 0$ implies that under any admissible investment and consumption policy, there is a positive probability that at the next trading time $W_\tau + X_\tau \leq 0$, violating limited liability, implying zero consumption, and leaving the objective function (4) at $-\infty$. For $\rho = 1$, $X_t < 0$ is ruled out by assuming that the illiquid asset has a weakly higher Sharpe ratio than the liquid asset (8). For $\rho = -1$ the investor invests positive amounts in the illiquid asset X_t and the liquid risky asset.

A.2 Proof of Proposition 2

The arguments in the proof of proposition 1 are sufficient to show that $H(1) = -\infty$. Concavity of $H(\xi)$ on $\xi \in [0,1)$ follows from Lemma 12 (below), and continuity on $\xi \in [0,1]$ from inspection. That $H(\xi)$ obtains its maximum for some $\xi \in [0,1)$ follows from concavity, continuity, and $H(1) = -\infty$.

Lemma 12 *$H(\xi)$ is concave on $\xi \in [0,1)$.*

Proof. Define $Q = X + W$ to be total wealth, and let $\{Q_0, X_0^1\}$ and $\{Q_0, X_0^3\}$ be two initial values with the associated optimal policies $\{C^1, \pi^1\}$ and $\{C^3, \pi^3\}$ where $\pi = \theta W$. For $\kappa \in (0,1)$, we consider a middle initial value $\{Q_0, X_0^2 = \kappa X_0^1 + (1-\kappa)X_0^3\}$ and the associated (possibly optimal) policies $\{C^2 = \kappa C^1 + (1-\kappa)C^3, \pi^2 = \kappa\pi^1 + (1-\kappa)\pi^3\}$, which are feasible because of the linearity of the budget constraint. From (5) and (6), we have

$$dQ_t = \left[rQ_t + (\mu - r)\pi_s + (\nu - r)X_t - C_t\right]dt + \left[\pi_s\sigma + \psi\rho X_t\right]dZ_t^1 + \psi X_t\sqrt{1-\rho^2}dZ_t^2 \tag{A.2}$$

for any time t. Thus, from the construction of our initial values and optimal policies, we have $Q_t^2 = \kappa Q_t^1 + (1-\kappa)Q_t^3$. Next, consider the objective function

$$\mathrm{E}\left[\int_0^\infty e^{-\beta t}U(C_t)\right] = \mathrm{E}\left[\int_0^\tau e^{-\beta t}U(C_t)dt + e^{-\beta\tau}Q_\tau^{1-\gamma}H^*\right] \tag{A.3}$$

Because $U(C)$ is increasing and concave, we have $U(C_t^2) > \kappa U(C_t^1) + (1-\kappa)U(C_t^3)$. From Jensen's inequality and $H^* < 0$, we have ${Q_\tau^2}^{1-\gamma}H^* > \kappa {Q_\tau^1}^{1-\gamma}H^* + (1-\kappa){Q_\tau^3}^{1-\gamma}H^*$. Thus, $\mathrm{E}^2\left[\int_0^\infty e^{-\beta t}U(C_t)\right] > \kappa\mathrm{E}^1\left[\int_0^\infty e^{-\beta t}U(C_t)\right] + (1-\kappa)\mathrm{E}^3\left[\int_0^\infty e^{-\beta t}U(C_t)\right]$, and so the value function is concave in X for fixed Q. Since $\xi = \frac{X}{Q}$, this is sufficient to show that the value function is concave in ξ for fixed Q, so H is concave. ■

To continue, we observe that the principal of optimality implies the Hamilton-Jacobi-Bellman equation between rebalancing times:

$$\begin{aligned} 0 \quad = \quad & \max_{c,\,\theta}\left[\frac{1}{1-\gamma}\left(cW\right)^{1-\gamma} - \beta F + F_W W\left(r + (\mu - r)\theta - c\right) + F_X X\nu\right. \\ & \left. +\lambda\left(F^* - F\right) + \frac{1}{2}F_{WW}W^2\theta^2\sigma^2 + \frac{1}{2}F_{XX}X^2\psi^2 + F_{WX}WX\psi\sigma\rho\theta\right] \end{aligned} \tag{A.4}$$

and substituting in (10) yields the stated ODE. A standard verification argument completes the proof.

A.3 Proof of Proposition 3

An investor prefers holding a small amount of the illiquid asset to holding a zero position if and only if $F_X(W, X = 0) \geq F_W(W, X = 0)$.

We begin by showing that $\frac{\nu - r}{\psi} - \rho\frac{\mu - r}{\sigma} \leq 0$ implies $F_X(W, X = 0) \leq F_W(W, X = 0)$.

Assume that we have $\{W_0, X_0 = \epsilon\}$, which gives rise to an optimal portfolio policy in number of shares equal to $\zeta_t = \frac{\theta_t W_t}{S_t}$ along paths for $t \in [0, \tau]$, where τ is the next trading time. $\{W_0, X_0 = \epsilon\}$ also gives rise to a consumption

policy C_t along those same paths. Then, total discounted wealth at the next trading time equals

$$\begin{aligned} e^{-r\tau}\left(W_\tau + X_\tau\right) &= W_0 + \epsilon + \int_0^\tau e^{-rt}\left[\zeta_t(\mu - r)S_t + (\nu - r)X_t - C_t\right]dt \\ &\quad + \int_0^\tau e^{-rt}\left[\zeta_t\sigma S_t + \psi\rho X_t\right]dZ_t^1 + \int_0^\tau e^{-rt}\left[\psi\sqrt{1-\rho^2}X_t\right]dZ_t^2 \end{aligned}$$

Now consider the starting point $\{\hat{W}_0 = W_0 + \epsilon, \hat{X}_0 = 0\}$ and use the previous consumption policy state-by-state (feasible because consumption is out of liquid wealth). The portfolio policy is now $\hat{\zeta}_t = \zeta_t + \frac{\psi\rho X_t}{\sigma S_t}$. Then,

$$\begin{aligned} e^{-r\tau}\left(\hat{W}_\tau + \hat{X}_\tau\right) &= W_0 + \epsilon + \int_0^\tau e^{-rt}\left[\zeta_t(\mu - r)S_t + \frac{\psi\rho X_t}{\sigma}(\mu - r) - C_t\right]dt \\ &\quad + \int_0^\tau e^{-rt}\left[\zeta_t\sigma S_t + \psi\rho X_t\right]dZ_t^1. \end{aligned}$$

The drift in the second ($\{\hat{W}_0 = W_0 + \epsilon, \hat{X}_0 = 0\}$) minus the drift in the first ($\{W_0, X_0 = \epsilon\}$) equals

$$\int_0^\tau e^{-rt}\left[\psi\rho X_t\frac{\mu - r}{\sigma} - (\nu - r)X_t\right]dt,$$

which is positive if $\psi\rho\frac{\mu-r}{\sigma} - (\nu - r) \geq 0$. Thus, the second initial condition produces higher expected wealth and lower volatility, path by path, with a possibly sub-optimal portfolio and consumption strategy. Since the value function at rebalancing (F^*) is increasing and concave, this proves that $\rho\frac{\mu-r}{\sigma} - \frac{\nu-r}{\psi} \geq 0$ implies $F(W_0 + \epsilon, 0) \geq F(W_0, \epsilon)$.

Next we will show that $\frac{\nu-r}{\psi} - \rho\frac{\mu-r}{\sigma} \geq 0$ implies $F_X(W, X = 0) \geq F_W(W, X = 0)$. Consider a deviation in which a trader starting with $\{W_0, 0\}$ is able to move an amount ϵ into X, and then withdraws it at the next trading day. This results in higher utility if

$$0 \leq -F_W(W_0, 0)\epsilon + \mathrm{E}\left[e^{-\beta\tau}F_W(W_\tau, 0)\epsilon\frac{X_\tau}{X_0}\right],$$

with W_t following the optimal portfolio and consumption policies (as a function of W_t) for $X_t = 0$. Plugging in the value function at $X = 0$, we obtain

$$1 \leq \mathrm{E}\left[e^{-\beta\tau}\left(\frac{W_\tau}{W_0}\right)^{-\gamma}\frac{X_\tau}{X_0}\right].$$

Direct calculation show that this is true if $\frac{\nu-r}{\psi} - \rho\frac{\mu-r}{\sigma} \geq 0$, and hence $F_X(W, X = 0) \geq F_W(W, X = 0)$.

A.4 Sketch of Proof for Propositions 4, 5, 7, and 9

The characterization of the value function in each of these economies closely follows the characterization in the baseline (Problem 1) economy. Each value function is bounded above and below (analogously to equation by economies in which the illiquid asset is fully liquid and in which the illiquid asset does not exist. Concavity and the other properties of the $H(\cdot)$ functions can be observed using arguments analogous to those in the proof of Proposition 2. The Hamilton-Jacobi-Bellman equation follows from the principal of optimality, and a standard verification theorem can be used to complete the characterizations. The solution to the model with fixed transaction costs follows standard arguments. See Stokey (2008) for a textbook treatment.

A.5 Proof of Proposition 8

The arguments in the proof of Proposition 1 are sufficient to show that for $|\rho| < 1$, the objective function is at $-\infty$ if either $W \leq 0$ or $X < 0$ for $S = I$. For $S = L$, we observe that the state will shift to $S = I$ without the possibility of re-balancing; as a result, if either $W \leq 0$ or $X < 0$, the objective function in the liquid state is also equal to $-\infty$. For $\rho = 1$, $X < 0$ is ruled out by (7). For $\rho = -1$ the investor invests positive amounts in the illiquid asset X_t and the liquid risky asset.

A.6 Proof of Corollary 10

In the liquid state, the investor's optimal portfolio policies θ_L^* and ξ_L^* satisfy the first order conditions

$$\begin{aligned}0 =& H_L^*(1-\gamma)(\mu - r) - H_L^*(1-\gamma)\,\gamma\,\theta_L^*\,\sigma^2 - H_L^*\gamma\,(1-\gamma)\,\rho\,\psi\,\sigma\,\xi_L^* \\ 0 =& \chi^I H_I'(\xi_L^*) + H_L^*(1-\gamma)(\nu - r) - H_L^*\gamma\,(1-\gamma)\,\psi^2\,\xi_L^* - H_L^*(1-\gamma)\gamma\rho\psi\theta_L^*\sigma\end{aligned}$$

Setting $\rho = 1$, dividing the first equation by $\sigma H_L^*(1-\gamma)$ and the second by $\psi H_L^*(1-\gamma)$, and then subtracting the first equation from the second leads to (33) and (35) with $\rho = 1$. Setting, $\rho = -1$, dividing the first equation by $\sigma H_L^*(1-\gamma)$ and the second by $\psi H_L^*(1-\gamma)$, and then adding the first equation from the second leads to (33) and (35) with $\rho = -1$.

A.7 Proof of Corollary 11

Consider a derivative security Y that pays a fixed rate of return κ when the aggregate state S switches from L to I. Denoting by dN_t^I the Poisson count process that denotes the arrival of a liquidity crisis, the price of this security evolves according to

$$\frac{dY_t}{Y_t} = (r + \mu_Y - \kappa\,\chi^I)\,dt + \kappa dN_t^I.$$

The investor is indifferent between participating in the market for security Z and her current portfolio policy as long as the excess return μ_Y is equal to

$$\begin{aligned}\mu_Y\,dt &= -cov\left(\frac{dY_t}{Y_t}, \frac{dF_W}{F_W}\right) \\ &= \kappa\,\chi^I\,\frac{H_L^*(1-\gamma) - H_I(\xi_L^*)(1-\gamma) + \xi_L^* H_I'(\xi_L^*)}{H_L^*(1-\gamma)}\,dt, \end{aligned} \tag{A.5}$$

where we have used the investor's marginal value of wealth process (36). There exists a fictitious probability measure $\mathcal{Q}$ under which the security Y has an expected excess return equal to zero,

$$E_t^{\mathcal{Q}}\left[\frac{dY_t}{Y_t}\right] = (r + \mu_Y - \kappa\,\chi^I + \kappa\hat{\chi}^I)\,dt = r\,dt$$

rearranging, using (A.5) and solving for the risk-neutral crisis probability $\hat{\chi}^I$ yields

$$\hat{\chi}_I = \chi^I\,\frac{F_W(W,X,I)}{F_W(W,X,L)} = \chi^I\,\frac{H_I(\xi_L^*)(1-\gamma) - \xi_L^* H_I'(\xi_L^*)}{H_L^*(1-\gamma)},$$

Under that measure, the present value of the expected payments p has to equal the expected payoff in the event of a liquidity crisis

$$\begin{aligned}& E_t^{\mathcal{Q}}\int_t^\tau e^{-r(s-t)}p\,ds = E_t^{\mathcal{Q}}\left[e^{-r(\tau-t)}\right] \\ \Rightarrow & \int_t^\infty e^{-(r+\hat{\chi}_I)(s-t)}p\,ds = \int_\tau^\infty e^{-(r+\hat{\chi}_I)(\tau-t)}\hat{\chi}_I\,d\tau \\ & \Rightarrow p = \hat{\chi}_I\end{aligned}$$

References

Abel, A. B., J. C. Eberly, and S. Panageas (2007). Optimal inattention to the stock market. *The American Economic Review 97*(2), pp. 244–249.

Acharya, V. V. and L. H. Pedersen (2005). Asset pricing with liquidity risk. *Journal of Financial Economics 77*(2), 375–410.

Acharya, V. V. and P. Schnabl (2010). Do global banks spread global imbalances? the case of asset-backed commercial paper during the financial crisis of 2007-09. *IMF Economic Review 58*(1), 37–73.

Anderson, R. G. and C. S. Gascon (2009). The commercial paper market, the fed, and the 2007-2009 financial crisis. *Federal Reserve Bank of St. Louis Review* (Nov), 589–612.

Ang, A., V. Bhansali, and Y. Xing (2010). Taxes on tax-exempt bonds. *Journal of Finance 65*, 565–601.

Ang, A. and N. P. Bollen (2010). Locked up by a lockup: Valuing liquidity as a real option. *Financial Management 39*(3), 1069–1096.

Ang, A. and R. C. Green (2011). Lowering borrowing costs for states and municipalities through commonmuni. Policy paper, The Hamilton Project, Brookings Institution.

Ang, A., A. Shtauber, and P. C. Tetlock (2012). Asset pricing in the dark: The cross section of otc stocks. Working paper, Columbia University.

Bao, J., J. Pan, and J. Wang (2011). The illiquidity of corporate bonds. *Journal of Finance 66*, 911–946.

Beeferman, L. W. (2008). Pension fund investment in infrastructure: A resource paper. Occasional paper no. 3, Pensions and Capital Stewardship Project, Labor and Worklife Program, Harvard Law School.

Bitsch, F., A. Buchner, and C. Kaserer (2010). Risk, return, and cash flow characteristics of infrastructure fund investments. *EIB Papers 15*, 106–136.

Borio, C. (2004). Market distress and vanishing liquidity: Anatomy and policy options. Working paper, Bank of International Settlements.

Brunnermeier, M. K. (2009). Deciphering the liquidity and credit crunch 2007-2008. *Journal of Economic Perspectives 23*(1), 77–100.

Case, K. E. and R. J. Shiller (1989). The efficiency of the market for single-family homes. *American Economic Review 79*, 125–137.

Collett, D., C. Lizieri, and C. Ward (2003). Timing and the holding periods of institutional real estate. *Real Estate Economics 31*, 205–222.

Constantinides, G. M. (1986). Capital market equilibrium with transaction costs. *Journal of Political Economy 94*(4), 842–62.

Cox, J. C. and C. Huang (1989). Optimal consumption and portfolio policies when asset prices follow a diffusion process. *Journal of Economic Theory 49*(1), 33 – 83.

Dai, M., P. Li, H. Liu, and Y. Wang (2010). Market closure and the liquidity premium puzzle. Working paper.

Daley, B. and B. Green (2012). Waiting for news in the market for lemons. *Econometrica 80*(4), 1433–1504.

Daley, B. and B. Green (2013). News and liquidity in markets with asymmetrically-informed traders. Working paper.

De Roon, F. A., J. Guo, and J. R. Ter Horst (2009). Being locked up hurts. Working paper.

Diamond, P. A. (1982). Aggregate demand management in search equilibrium. *Journal of Political Economy 90*(5), 881–94.

Driessen, J., T.-C. Lin, and L. Phalippou (2008). A new method to estimate risk and return of non-traded assets from cash flows: The case of private equity funds. NBER Working Paper 14144, National Bureau of Economic Research, Inc.

Duffie, D. and L. G. Epstein (1992). Stochastic differential utility. *Econometrica 60*(2), 353–94.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005). Over-the-counter markets. *Econometrica 73*(6), 1815–1847.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2007). Valuation in over-the-counter markets. *Review of Financial Studies 20*(6), 1865–1900.

Dwyer, G. P. and P. Tkac (2009). The financial crisis of 2008 in fixed-income markets. *Journal of International Money and Finance 28*(8), 1293–1316.

Dybvig, P. H. and C. Huang (1988). Nonnegative wealth, absence of arbitrage, and feasible consumption plans. *Review of Financial Studies 1*(4), 377–401.

Elul, R. (2008). Liquidity crises. *Business Review 2*, 13–22.

Fisher, J. D. and M. S. Young (2000). Institutional property tenure: Evidence from the ncreif database. *Journal of Real Estate Portfolio Management 6*, 327–338.

Florance, A. C., N. G. Miller, J. Spivey, and R. Peng (2010). Slicing, dicing and scoping the size of the u.s. commercial real estate market. policy paper, CoStar Group.

Frans M. Dieleman, W. A. V. C. and M. C. Deurloo (2000). The geography of residential turnover in twenty-seven large us metropolitan housing markets, 1985-95. *Urban Studies 37*, 223–245.

Gârleanu, N. and L. H. Pedersen (2011). Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies 24*(6), 2980–2022.

Gârleanu, N. B. (2009). Pricing and portfolio choice in illiquid markets. *Journal of Economic Theory 144*(2), 532–564.

Goetzmann, W. N. (1993). Accounting for taste: Art and the financial markets over three centuries. *American Economic Review 83*, 1370–1376.

Gorton, G. (2009). Information, liquidity, and the (ongoing) panic of 2007. *American Economic Review 99*(2), 567–72.

Gorton, G. and A. Metrick (2012). Securitized banking and the run on repo. *Journal of Financial Economics 104*(3), 425–451.

Gorton, G. B. (2010). Questions and answers about the financial crisis. NBER Working Paper 15787, National Bureau of Economic Research, Inc.

Gromb, D. and D. Vayanos (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics 66*(2-3), 361–407.

Grossman, S. J. and G. Laroque (1990). Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods. *Econometrica 58*(1), 25–51.

Hansen, K. A. (1998). Seasonality of moves and duration of residence. Current population report p70-66, Household Economic Studies, Census Bureau.

Heaton, J. and D. Lucas (2000). Portfolio choice and asset prices: The importance of entrepreneurial risk. *Journal of Finance 55*(3), 1163–1198.

Hu, X., J. Pan, and J. Wang (2012). Noise as information for illiquidity. Working paper, NBER.

Kahl, M., J. Liu, and F. A. Longstaff (2003). Paper millionaires: how valuable is stock to a stockholder who is restricted from selling it? *Journal of Financial Economics 67*(3), 385–410.

Kaplan, E. H. (1997). Snapshot samples. *Socio-Economic Planning Sciences 31*, 281–291.

Kaplan, G. and G. L. Violante (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics 2*(4), 53–87.

Kaplan, S. N. and A. Schoar (2005). Private equity performance: Returns, persistence, and capital flows. *Journal of Finance 60*(4), 1791–1823.

Keely, L., B. van Ark, G. Levanon, and J. Burbank (2012). The shifting nature of U.S. housing demand: The U.S. housing market is growing again – but not as we knew it. Working paper, Demand Institute.

Korajczyk, R. A. and R. Sadka (2008). Pricing the commonality across alternative measures of liquidity. *Journal of Financial Economics 87*(1), 45–72.

Koren, M. and A. Szeidl (2003). Portfolio choice with illiquid assets. CEPR Discussion Papers 3795, C.E.P.R. Discussion Papers.

Krishnamurthy, A., S. Nagel, and D. Orlov (2012). Sizing up repo. NBER Working Paper 17768, National Bureau of Economic Research, Inc.

Lagos, R. and G. Rocheteau (2009). Liquidity in asset markets with search frictions. *Econometrica 77*(2), 403–426.

Levitt, S. D. and C. Syverson (2008). Market distortions when agents are better informed: The value of information in real estate transactions. *The Review of Economics and Statistics 90*(4), 599–611.

Ling, D. C., G. Marcato, and P. McAllister (2009). Dynamics of asset prices and transaction activity in illiquid markets: The case of private commercial real estate. *Journal of Real Estate Finance and Economics 39*, 359–383.

Lo, A. W., H. Mamaysky, and J. Wang (2004). Asset prices and trading volume under fixed transactions costs. *Journal of Political Economy 112*(5), 1054–1090.

Longstaff, F. A. (2001). Optimal portfolio choice and the valuation of illiquid securities. *Review of Financial Studies 14*(2), 407–31.

Longstaff, F. A. (2009). Portfolio claustrophobia: Asset pricing in markets with illiquid assets. *American Economic Review 99*(4), 1119–44.

Lopez-de Silanes, F., L. Phalippou, and O. Gottschalg (2010). Giants at the gate: On the cross-section of private equity investment returns. MPRA Paper 28487, University Library of Munich, Germany.

McConell, John, J. and A. Saretto (2010). Auction failures and the market for auction rate securities. *Journal of Financial Economics 97*, 451–469.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory 3*(4), 373–413.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica 41*(5), 867–87.

Metrick, A. and A. Yasuda (2010). The economics of private equity funds. *Review of Financial Studies 23*, 2303–2341.

Miller, N. G., L. Peng, and M. A. Sklarz (2011). The economic impact of anticipated house price changes–evidence from home sales. *American Economic Review 39*, 345–378.

Morris, S. and H. S. Shin (2010). Liquidity black holes. *Review of Finance 8*, 1–18.

Pastor, L. and R. F. Stambaugh (2003). Liquidity risk and expected stock returns. *Journal of Political Economy 111*(3), 642–685.

Persaud, A. (2001). Liquidity black holes. Working paper, State Street Bank.

Phalippou, L. and O. Gottschalg (2009). The performance of private equity funds. *Review of Financial Studies 22*(4), 1747–1776.

Rogers, L. and O. Zane (2002). A simple model of liquidity effects. In K. Sandmann and P. Schoenbucher (Eds.), *Advances in Finance and Stochastics: Essays in Honour of Dieter Sondermann*, pp. 161–176. Springer, Berlin.

Sadka, R. (2006). Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk. *Journal of Financial Economics 80*(2), 309–349.

Santos, T. and P. Veronesi (2006). Labor income and predictable stock returns. *Review of Financial Studies 19*(1), 1–44.

Schwartz, E. S. and C. Tebaldi (2006). Illiquid assets and optimal portfolio choice. NBER Working Papers 12633, National Bureau of Economic Research, Inc.

Shleifer, A. and R. W. Vishny (1997). The limits of arbitrage. *Journal of Finance 52*(1), 35–55.

Skaterschivkov, S. (2006). *Skate's Art Investment Handbook*. KunstAM GmBH.

Stokey, N. L. (2008). *The Economics of Inaction: Stochastic Control Models with Fixed Costs*. Princeton University Press.

Taleb, N. (1997). *Dynamic Hedging*. John Wiley, New York.

Tirole, J. (2011). Illiquidity and all its friends. *Journal of Economic Literature 49*(2), 287–325.

Vayanos, D. (1998). Transaction costs and asset prices: A dynamic equilibrium model. *Review of Financial Studies 11*(1), 1–58.

Vayanos, D. and P.-O. Weill (2008). A search-based theory of the on-the-run phenomenon. *Journal of Finance 63*(3), 1361–1398.

Wermers, R. (2000). Mutual fund performance: An empirical decomposition into stock-picking talent, style, transactions costs, and expenses. *Journal of Finance 55*(4), 1655–1703.

Winchell, W. B. (2010). A brief history of the private equity secondary market. Working paper, Partner Ventures.

Table 1: Holding Periods and Turnover of Various Asset Classes

Asset Class	Typical Time between Transactions	Annualized Turnover	References
Public Equities	Within seconds	over 100%	Turnover can be computed from NYSE and NASDAQ data at nyxdata.com and nasdaqtrader.com, respectively.
OTC (Pinksheet) Equities	Within a day, but many stocks over a week	∼35%	Ang, Shtauber, and Tetlock (2012).
Corporate Bonds	Within a day	25-35%	Bao, Pan, and Wang (2011).
Municipal Bonds	∼6 months, with 5% of muni bonds trading more infrequently than once per decade	$< 10\%$	Holding period statistics computed from an updated version of the dataset from Ang, Bhansali, and Xing (2010) and summarized in Ang and Green (2011). Turnover numbers from the SEC Report on Transactions in Municipal Bond Securities, 2004.
Private Equity	Funds last for 10 years; the median investment duration is 4 years; secondary trade before exit is relatively unusual.	$< 10\%$	Private equity contracts are described by Metrick and Yasuda (2010); for duration see Lopez-de Silanes et al. (2010). For estimates of "secondaries" in private equity see http://lavca.org/2012/07/19/lp-profile-an-interview-with-maureen-downey-pantheon/ and Winchell (2010).
Residential Housing	4-5 years, but ranges from months to decades	4-6%	For median duration in residences see Hansen (1998) and Case and Shiller (1989), with Miller, Peng, and Sklarz (2011) comments on the range. Turnover numbers are computed by Frans M. Dieleman and Deurloo (2000).
Institutional Real Estate	8-11 years	∼7%	See Fisher and Young (2000) and Collett, Lizieri, and Ward (2003) for holding periods and Ling, Marcato, and McAllister (2009) for turnover.
Institutional Infrastructure	50-60 years for initial commitment, some as long as 99 years	Negligible	Beeferman (2008), Bitsch, Buchner, and Kaserer (2010).
Art	40-70 years	$< 15\%$	For holding periods see Goetzmann (1993) and Kaplan (1997). Turnover can be inferred from the size of the art market estimated by Skaterschivkov (2006) and estimated annual art sales.

Table 2: Liquid and Illiquid Asset Returns

	1981Q3 – 2010Q2			1981Q3 – 2006Q4		
	Mean	Stdev	Corr	Mean	St Dev	Corr
Equity	0.103	0.182	1.000	0.125	0.157	1.000
Illiquid Assets						
Private Equity	0.103	0.229	0.629	0.110	0.231	0.605
Buyout	0.092	0.134	0.267	0.097	0.110	0.010
Venture Capital	0.133	0.278	0.557	0.143	0.286	0.548
Illiquid Investment	0.109	0.165	0.674	0.117	0.159	0.623

The table reports summary statistics on excess returns on liquid and illiquid assets. Liquid equity returns are total returns on the S&P500. Data on private equity, buyout, and venture capital funds are obtained from Venture Economics and Cambridge Associates. We construct annual horizon log returns at the quarterly frequency. We compute log excess returns using the difference between log returns on the asset and year-on-year rollover returns on one-month T-bills expressed as a continuously compounded rate. The column "Corr" reports the correlation of excess returns with equity. The illiquid investment is a portfolio invested with equal weights in private equity, buyout, and venture capital and is rebalanced quarterly.

Table 3: The Baseline Model

Avg. Rebalancing	Optimal	Illiquidity	Average policies		
Interval $(1/\lambda)$	Rebalance (ξ^*)	Utility cost	$E[\xi]$	$E[c(1-\xi)]$	$E[\theta(1-\xi)]$
0	0.593	-	0.593	0.089	0.593
1/50	0.535	0.018	0.541	0.088	0.583
1/10	0.493	0.028	0.511	0.087	0.578
1/4	0.475	0.036	0.485	0.086	0.571
1/2	0.442	0.045	0.461	0.083	0.568
1	0.373	0.067	0.409	0.081	0.558
2	0.251	0.103	0.299	0.075	0.546
4	0.132	0.165	0.212	0.069	0.536
10	0.048	0.222	0.214	0.059	0.489
∞	-	-	-	0.070	0.593

The table displays the effect of illiquidity on portfolio choice and welfare in the baseline model. The long-run average policies are computed using a long simulation of 10,000 years. The cases $E(T) = 0$ and $E(T) = \infty$ correspond, with some abuse of notation, to the Merton one- and two-asset cases respectively. The fraction of illiquid assets to total wealth is ξ, with optimal value ξ^* at the time of re-balancing. Consumption as a fraction of total wealth is $c(1-\xi)$ and the allocation to the liquid asset as a function of total wealth is $\theta(1-\xi)$. The table is computed using the following parameter values: $\gamma = 6$, $\beta = 0.1$, $\mu = \nu = .12$, $r = .04$, $\sigma = \psi = .15$, and $\rho = 0$.

Table 4: The Model Without Intermediate Consumption

Avg. Rebalancing	Optimal	Illiquidity	Average policies	
Interval ($1/\lambda$)	Rebalance (ξ^*)	Utility cost	$E[\xi]$	$E[\theta(1-\xi)]$
0	0.593	-	0.593	0.593
1/50	0.557	0.002	0.555	0.580
1/10	0.547	0.002	0.546	0.578
1/4	0.544	0.002	0.543	0.576
1/2	0.535	0.002	0.534	0.574
1	0.527	0.003	0.526	0.571
2	0.523	0.003	0.523	0.567
4	0.520	0.003	0.518	0.563
10	0.518	0.004	0.516	0.555
∞	-	-	-	0.593

The table displays the effect of illiquidity on portfolio choice and welfare in the case without intermediate consumption, see Section 5.1 for more details. The cases $E(T) = 0$ and $E(T) = \infty$ correspond, with some abuse of notation, to the Merton one- and two-asset cases respectively. The long-run average policies are computed using a long simulation of 10,000 years. The fraction of illiquid assets to total wealth is ξ, with optimal value ξ^* at the time of re-balancing. The allocation to the liquid asset as a function of total wealth is $\theta(1-\xi)$. The table is computed using the following parameter values: $\gamma = 6$, $\beta = 0.1$, $\mu = \nu = .12$, $r = .04$, $\sigma = \psi = .15$, and $\rho = 0$.

Table 5: The Model With Deterministic Liquidity

Rebalancing	Optimal	Illiquidity	Average policies		
Interval (T)	Rebalance (ξ^*)	Utility cost	$E[\xi]$	$E[c(1-\xi)]$	$E[\theta(1-\xi)]$
0	0.593	-	0.593	0.089	0.593
1/50	0.575	0.011	0.579	0.089	0.589
1/10	0.555	0.013	0.559	0.088	0.585
1/4	0.532	0.016	0.542	0.088	0.579
1/2	0.516	0.019	0.522	0.087	0.576
1	0.494	0.022	0.512	0.087	0.575
2	0.488	0.024	0.501	0.086	0.574
4	0.455	0.025	0.481	0.085	0.573
10	0.448	0.028	0.464	0.084	0.571
∞	-	-	-	0.070	0.593

The table displays the effect of illiquidity on portfolio choice and welfare in the case where trading opportunities arrive deterministically, see Section 5.2 for more details. The cases $E(T) = 0$ and $E(T) = \infty$ correspond, with some abuse of notation, to the Merton one- and two-asset cases respectively. The long-run average policies are computed using a long simulation of 10,000 years. The fraction of illiquid assets to total wealth is ξ, with optimal value ξ^* at the time of re-balancing. Consumption as a fraction of total wealth is $c(1-\xi)$ and the allocation to the liquid asset as a function of total wealth is $\theta(1-\xi)$. Unless otherwise noted, we use the following parameter values: $\gamma = 6$, $\beta = 0.1$, $\mu = \nu = .12$, $r = .04$, $\sigma = \psi = .15$, and $\rho = 0$.

Table 6: Allowing for Costly Liquidity

	A. Frequency Cost Paid				B. Optimal Rebalance (ξ^*)				C. Illiquidity Utility Cost			
$\lambda \backslash \kappa$	1%	5%	10%	20%	1%	5%	10%	20%	1%	5%	10%	20%
50	0.000	0.000	0.000	0.000	0.535	0.535	0.535	0.535	0.017	0.017	0.017	0.017
10	0.000	0.000	0.000	0.000	0.505	0.505	0.505	0.505	0.024	0.024	0.024	0.024
4	0.001	0.000	0.000	0.000	0.496	0.496	0.496	0.496	0.025	0.025	0.025	0.025
2	0.002	0.000	0.000	0.000	0.483	0.483	0.483	0.483	0.026	0.026	0.026	0.026
1	0.007	0.001	0.001	0.001	0.460	0.454	0.451	0.446	0.029	0.034	0.037	0.038
1/2	0.034	0.006	0.004	0.003	0.435	0.409	0.397	0.38	0.034	0.040	0.043	0.048
1/4	0.063	0.023	0.014	0.009	0.418	0.367	0.341	0.308	0.039	0.053	0.061	0.073
1/10	0.101	0.046	0.033	0.021	0.405	0.332	0.291	0.242	0.043	0.065	0.081	0.103
0	0.131	0.077	0.059	0.046	0.395	0.302	0.248	0.193	0.046	0.078	0.102	0.135
	D. Average Policies											
	$E[\xi]$				$E[c(1-\xi)]$				$E[\theta(1-\xi)]$			
$\lambda \backslash \kappa$	1%	5%	10%	20%	1%	5%	10%	20%	1%	5%	10%	20%
50	0.542	0.509	0.509	0.509	0.088	0.088	0.088	0.088	0.585	0.581	0.578	0.577
12	0.518	0.509	0.509	0.508	0.088	0.088	0.088	0.087	0.581	0.577	0.576	0.575
4	0.507	0.507	0.507	0.501	0.087	0.087	0.086	0.086	0.578	0.575	0.575	0.572
2	0.502	0.502	0.501	0.491	0.087	0.087	0.086	0.085	0.574	0.574	0.573	0.570
1	0.500	0.484	0.481	0.477	0.086	0.086	0.085	0.084	0.569	0.557	0.551	0.551
1/2	0.499	0.479	0.471	0.451	0.086	0.085	0.084	0.083	0.566	0.558	0.552	0.550
1/4	0.495	0.477	0.455	0.434	0.085	0.084	0.084	0.082	0.564	0.550	0.541	0.528
1/10	0.492	0.478	0.468	0.425	0.084	0.083	0.081	0.078	0.562	0.541	0.521	0.511
0	0.490	0.476	0.475	0.471	0.084	0.082	0.079	0.073	0.550	0.536	0.509	0.489

The table reports the optimal policies and the utility cost of illiquidity for different levels of illiquidity λ and transaction cost κ, in the hybrid model with transaction costs. See Section 5.3 for more details. The long-run average policies are computed using a long simulation of 10,000 years. The fraction of illiquid assets to total wealth is ξ, with optimal value ξ^* at the time of re-balancing. Consumption as a fraction of total wealth is $c(1-\xi)$ and the allocation to the liquid asset as a function of total wealth is $\theta(1-\xi)$. The table is computed using the following parameter values: $\gamma = 6$, $\beta = 0.1$, $\mu = \nu = .12$, $r = .04$, $\sigma = \psi = .15$, and $\rho = 0$.

Table 7: The Effect of Preference Parameters on Optimal Policies and Welfare

λ	A. Risk aversion – γ															
	3	6	10	15	3	6	10	15	3	6	10	15	3	6	10	15
	i. Optimal Rebalance				ii. Utility cost				iii. $E(c(1-\xi))$				iv. $E(\theta(1-\xi))$			
∞	1.185	0.593	0.356	0.237	-	-	-	-	0.129	0.090	0.074	0.069	1.185	0.593	0.356	0.237
10	0.799	0.493	0.336	0.220	0.109	0.029	0.020	0.005	0.117	0.086	0.072	0.068	1.122	0.572	0.349	0.255
4	0.666	0.475	0.291	0.201	0.114	0.037	0.031	0.021	0.117	0.086	0.064	0.065	1.113	0.571	0.349	0.238
2	0.530	0.442	0.285	0.199	0.132	0.041	0.035	0.030	0.114	0.082	0.068	0.065	1.097	0.561	0.343	0.234
1	0.409	0.373	0.280	0.197	0.172	0.067	0.041	0.035	0.110	0.081	0.072	0.065	1.076	0.552	0.335	0.233
1/2	0.289	0.246	0.192	0.149	0.256	0.107	0.051	0.041	0.105	0.080	0.069	0.063	1.069	0.548	0.331	0.224
1/4	0.178	0.132	0.101	0.075	0.295	0.165	0.106	0.073	0.093	0.075	0.065	0.060	1.045	0.544	0.332	0.225
1/10	0.061	0.047	0.035	0.025	0.325	0.222	0.158	0.116	0.088	0.068	0.061	0.057	1.020	0.539	0.328	0.222
0	-	-	-	-	-	-	-	-	0.090	0.070	0.062	0.058	1.185	0.593	0.356	0.237
λ	B. Inverse Elasticity of intertemporal substitution – ζ															
	2/3	1.5	3	6	2/3	1.5	3	6								
	i. Optimal Rebalance				ii. Utility cost				iii. $E(c(1-\xi))$				iv. $E(\theta(1-\xi))$			
∞	0.593	0.593	0.593	0.593	-	-	-	-	0.109	0.096	0.092	0.090	0.593	0.593	0.593	0.593
10	0.501	0.500	0.499	0.493	0.010	0.023	0.026	0.029	0.108	0.095	0.089	0.087	0.573	0.572	0.572	0.572
4	0.495	0.490	0.489	0.475	0.019	0.027	0.029	0.037	0.108	0.095	0.089	0.087	0.572	0.571	0.571	0.571
2	0.433	0.439	0.441	0.442	0.036	0.039	0.040	0.041	0.108	0.094	0.089	0.086	0.569	0.569	0.568	0.568
1	0.391	0.398	0.383	0.373	0.040	0.047	0.052	0.067	0.107	0.093	0.088	0.085	0.561	0.558	0.558	0.557
1/2	0.244	0.245	0.248	0.251	0.077	0.093	0.100	0.103	0.103	0.089	0.083	0.080	0.542	0.544	0.545	0.545
1/4	0.114	0.121	0.129	0.131	0.139	0.154	0.161	0.165	0.103	0.084	0.077	0.074	0.534	0.534	0.534	0.534
1/10	0.026	0.038	0.045	0.048	0.177	0.203	0.216	0.222	0.099	0.078	0.069	0.065	0.508	0.510	0.511	0.511
0	-	-	-	-	-	-	-	-	0.118	0.088	0.076	0.070	0.593	0.593	0.593	0.593

The table shows the effect on welfare and portfolio choice of preference parameters – the coefficient of relative risk aversion γ and the inverse of the elasticity of intertemporal substitution ζ. See Section 5.4 for more details. The long-run average policies are computed using a long simulation of 10,000 years. The fraction of illiquid assets to total wealth is ξ, with optimal value ξ^* at the time of re-balancing. Consumption as a fraction of total wealth is $c(1-\xi)$ and the allocation to the liquid asset as a function of total wealth is $\theta(1-\xi)$. Unless noted otherwise, the table is computed using the following parameter values: $\zeta = 6$, $\gamma = 6$, $\beta = 0.1$, $\mu = \nu = .12$, $r = .04$, $\sigma = \psi = .15$, and $\rho = 0$.

Figure 1: Model Solution

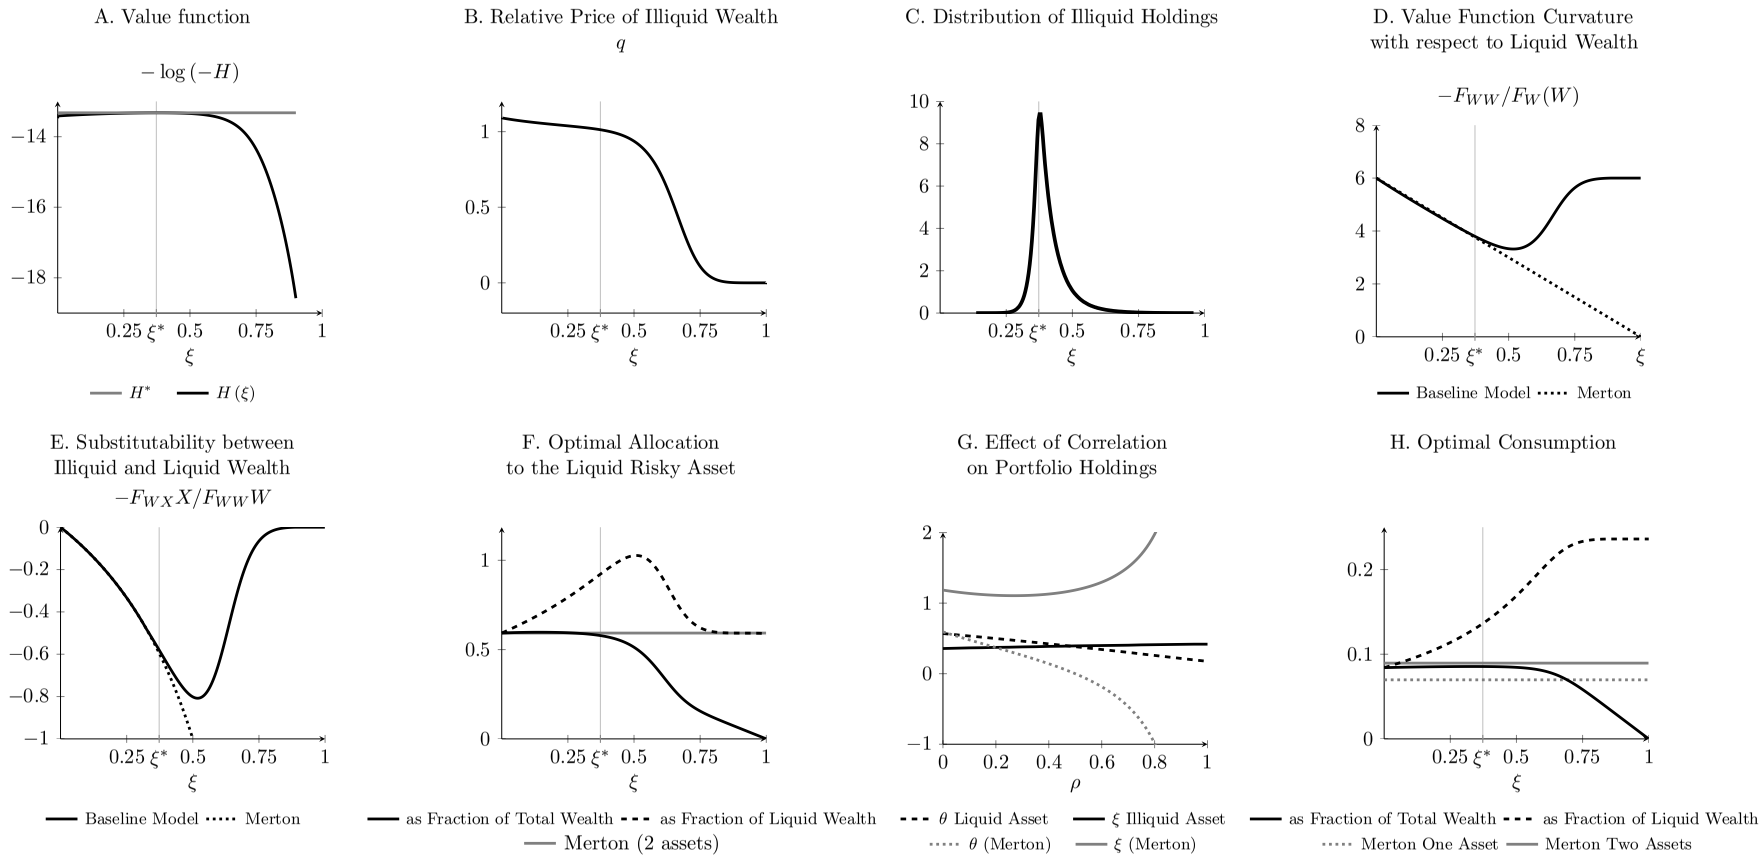

Panel A plots the value function. The vertical solid gray line corresponds to the value of the optimal rebalancing point ξ^*. Panel B plots the relative price of illiquid wealth q. Panel C plots the stationary distribution of allocation to the illiquid asset as a fraction of total wealth, $\xi = \frac{X}{X+W}$. We obtain the stationary distribution using a long simulation of 10,000 years. Panels D and E plot the curvature of the value function with respect to liquid wealth $-F_{WW}W/F_W$, and the elasticity of substitution in the value function between liquid and illiquid wealth, $F_{WX}X/F_{WW}W$, respectively. The solid lines represents the case where $\lambda = 1$. The dotted lines correspond to the Merton case. In Panel E, we assume $\rho = 0.6$. Panel F displays the optimal allocation to the liquid assets. Panel G examines the effect of correlation on risky asset holdings. We plot the optimal allocations to the liquid risky asset θ and the illiquid risky asset ξ as a fraction of total wealth at the rebalancing time, both as a function of ρ. The remainder is allocated to the riskless asset. For this panel we use $\nu = 0.2$. Panel H graphs consumption policy. In Panels F and H, the gray horizontal line corresponds to the allocation to the risky asset in the one- and/or two-asset Merton economy. Unless noted otherwise, all figures use $\gamma = 6$, $\beta = 0.1$, $\mu = \nu = .12$, $r = .04$, $\lambda = 1$, $\sigma = \psi = .15$, and $\rho = 0$.

Figure 2: Optimal Policies Across Liquidity Regimes

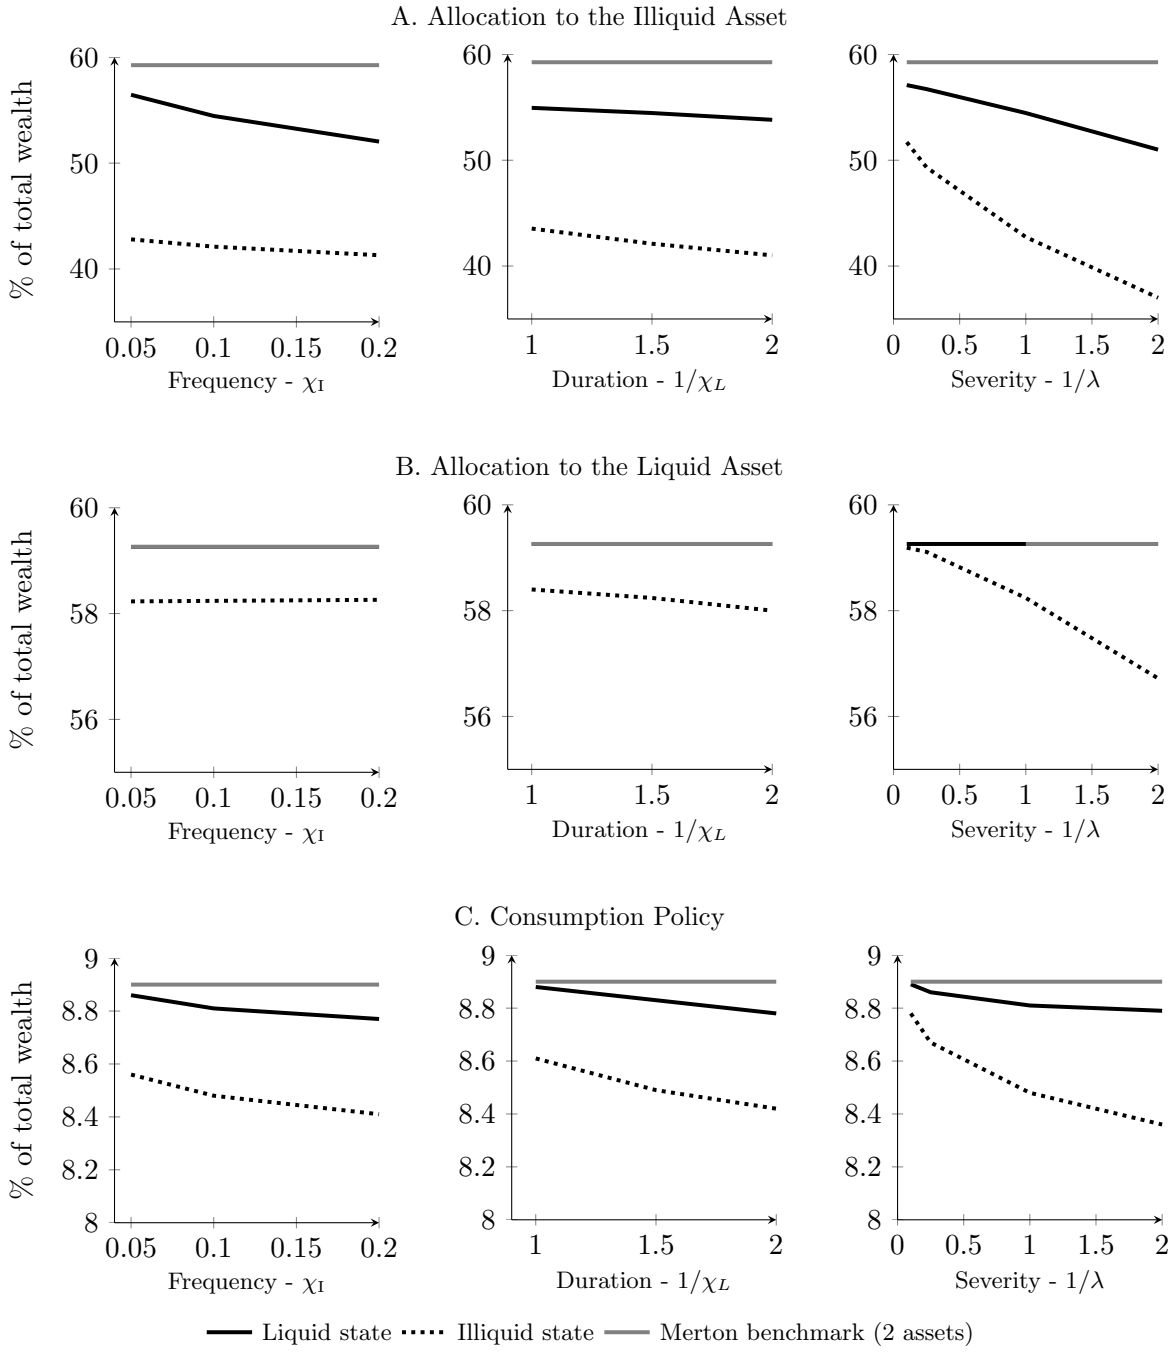

The figure compares portfolio (Panels A and B) and consumption (Panel C) policies across the liquid (solid line) and illiquid (dotted line) regimes for different frequency (χ_I), average duration ($1/\chi_L$), and severity ($1/\lambda$) of liquidity crises. Unless noted otherwise, the curves are plotted with the following parameter values: $\gamma = 6$, $\beta = 0.1$, $\mu = \nu = .12$, $r = .04$, $\sigma = \psi = .15$, $\rho = 0$, $\chi_I = 0.1$, $\chi_L = 1/1.5$ and $\lambda = 1$

Figure 3: Limits to Arbitrage

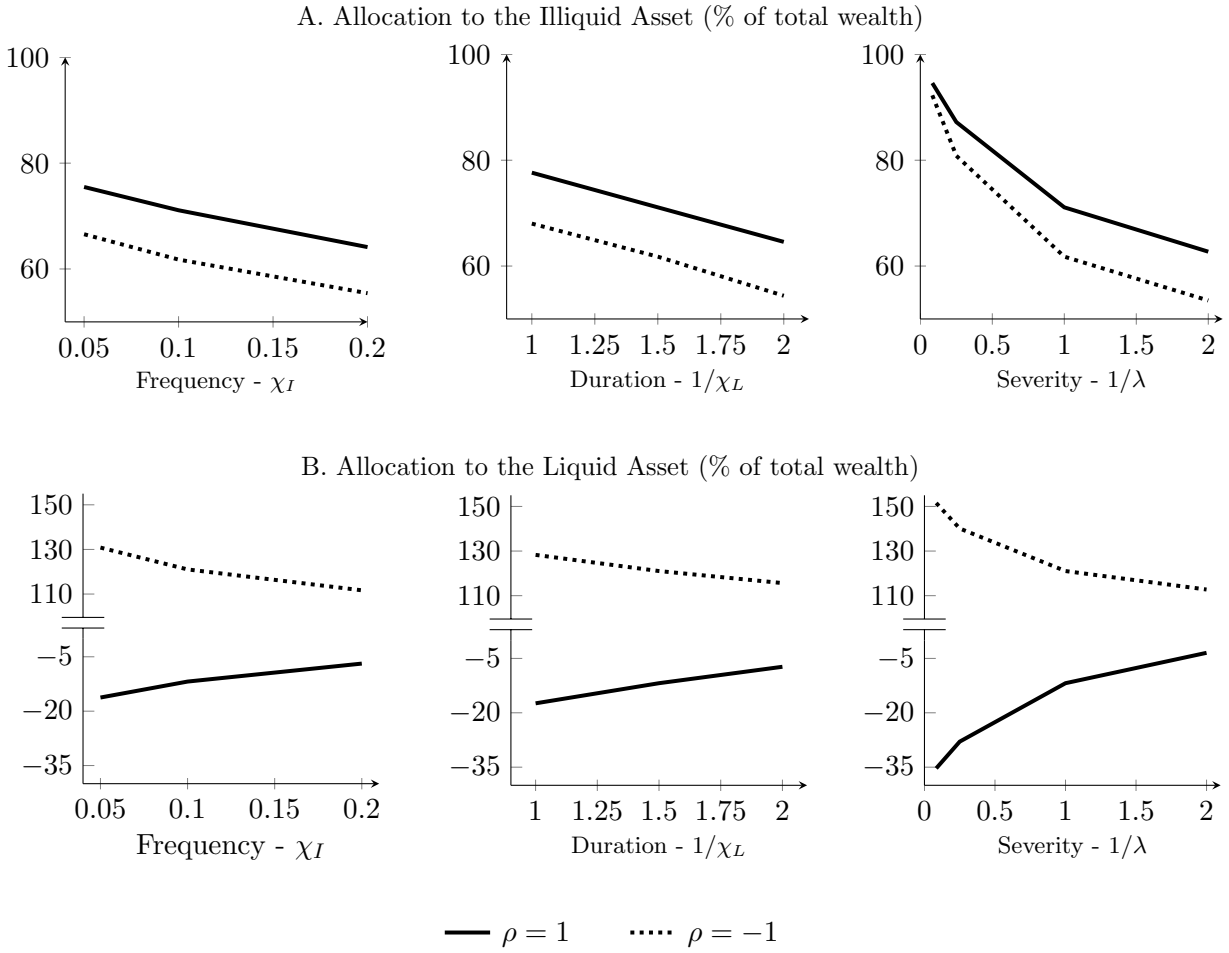

The figure compares portfolio policies of liquid assets (Panel A) and illiquid assets (Panel B) in the liquid regime in the case of an apparent arbitrage opportunity. We vary the frequency (χ_I), average duration ($1/\chi_L$) and severity ($1/\lambda$) of liquidity crises. Unless noted otherwise, the curves are plotted with the following parameter values: $\gamma = 6$, $\beta = 0.1$, $\mu = .12$, $\nu = .14$, $r = .04$, $\sigma = \psi = .15$, $\chi_I = 0.1$, $\chi_L = 1/1.5$ and $\lambda = 1$

Figure 4: Liquidity Risk Premium and the Welfare Cost of Illiquidity

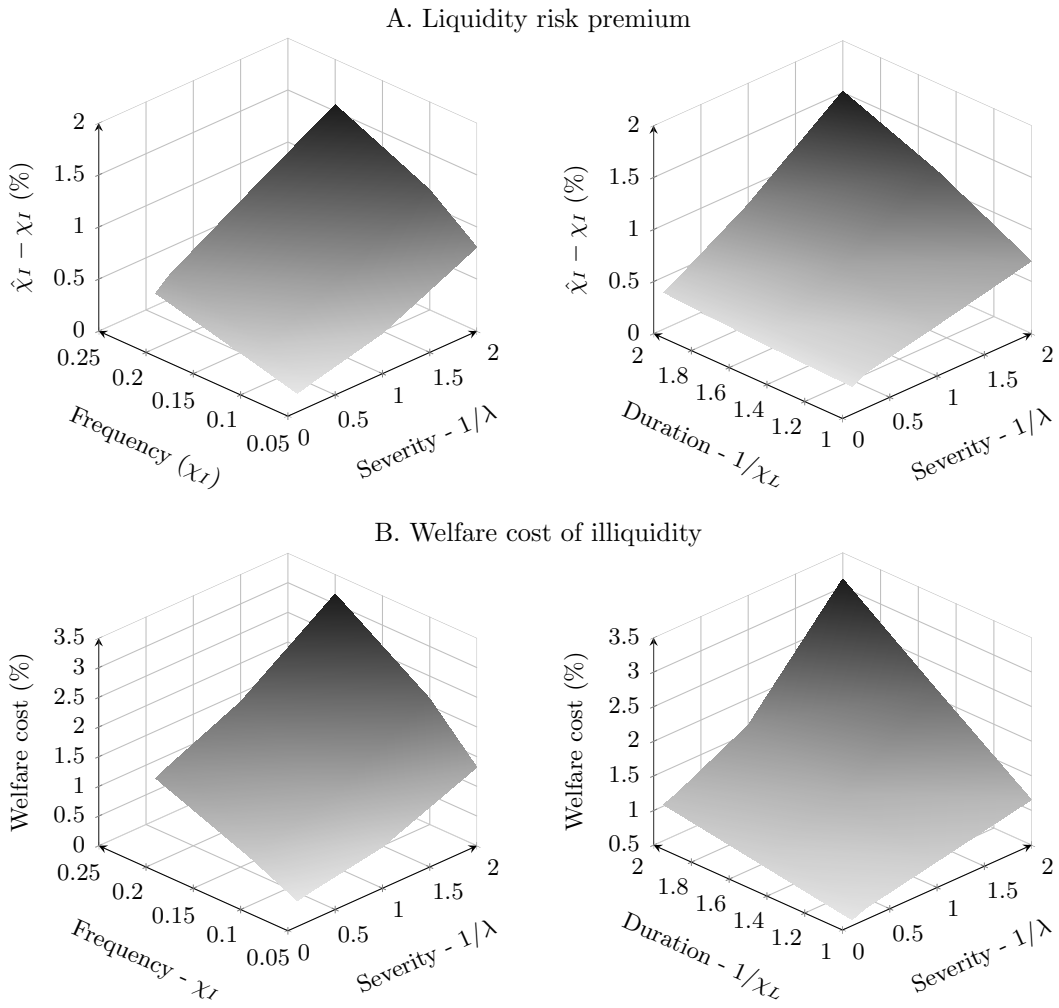

The figure shows the risk premium associated with liquidity insurance $\hat{\chi}_I - \chi_I$ (Panel A) and the welfare cost of illiquidity (Panel B) for different frequency (χ_I), average duration ($1/\chi_L$) and severity ($1/\lambda$) of liquidity crises. Unless noted otherwise, the curves are plotted with the following parameter values: $\gamma = 6$, $\beta = 0.1$, $\mu = \nu = .12$, $r = .04$, $\sigma = \psi = .15$, $\rho = 0$, $\chi_I = 0.1$, $\chi_L = 1/1.5$