

Explaining the Magnitude of Liquidity Premia: The Roles of Return Predictability, Wealth Shocks and State-dependent Transaction Costs*

Anthony W. Lynch[†]
New York University and NBER

Sinan Tan[‡]
New York University

First Version: 9 December 2002

This Version: 8 December 2004

Work in progress.
Comments welcome.

*The authors would like to thank Viral Acharya, George Constantinides, Joel Hasbrouck, John Heaton, Lasse Pedersen, Hans Stoll and the participants of the NBER Spring Asset Pricing Group Meeting, NYU Monday Finance Seminar, the NYU Macro-Finance Reading Group, the NHH Finance Seminar, the UCLA Finance Seminar and the Vanderbilt Finance Seminar for helpful comments and suggestions. All remaining errors are of course the authors' responsibility.

[†]Stern School of Business, New York University, 44 West Fourth Street, Suite 9-190, New York, NY 10012-1126, alynch@stern.nyu.edu, (212) 998-0350.

[‡]Stern School of Business, New York University, 44 West Fourth Street, Suite 9-190, New York, NY 10012-1126, stan@stern.nyu.edu, (212) 998-0321.

Explaining the Magnitude of Liquidity Premia: The Roles of Return Predictability, Wealth Shocks and State-dependent Transaction Costs

Abstract

The seminal work of Constantinides (1986) documents how, when the risky return is calibrated to the U.S. market return, the impact of transaction costs on per-annum liquidity premia is an order of magnitude smaller than the cost rate itself. A number of recent papers have formed portfolios sorted on liquidity measures and found a spread in expected per-annum return that is definitely not an order of magnitude smaller than the transaction cost spread: the expected per-annum return spread is found to be around 6-7% per annum. Our paper bridges the gap between Constantinides' theoretical result and the empirical magnitude of the liquidity premium by examining dynamic portfolio choice with transaction costs in a variety of more elaborate settings that move the problem closer to the one solved by real-world investors. In particular, we allow returns to be predictable and transaction costs to be stochastic, and we introduce wealth shocks, both stationary multiplicative and labor income. With predictable returns, we also allow the wealth shocks and transaction costs to be state dependent. We find that adding these real world complications to the canonical problem can cause transactions costs to produce per-annum liquidity premia that are no longer an order of magnitude smaller than the rate, but are instead the same order of magnitude. For example, predictable returns and i.i.d. labor income growth causes the liquidity premium for an agent with a wealth to monthly labor income ratio of 0 or 10 to be 1.68% and 1.20% respectively; these are 21-fold and 15-fold increases, respectively, relative to that in the standard i.i.d. return case. We conclude that the effect of proportional transaction costs on the standard consumption and portfolio allocation problem with i.i.d. returns can be materially altered by reasonable perturbations that bring the problem closer to the one investors are actually solving.

1 Introduction

A number of recent papers have found a difference in expected return across portfolios sorted on liquidity measures. While the finding of a difference is not surprising, the magnitude is, with expected return differences on the order of 6-7% per annum.¹ The magnitude of the difference seems too large to be explained by realistic transactions costs. In particular, the seminal work of Constantinides (1986) documents how investor utility is largely insensitive to transaction costs when the investor solves the canonical problem of i.i.d. returns calibrated to U.S. data, no non-financial income, and a constant proportional cost rate. For realistic proportional costs, Constantinides shows that the per-annum liquidity premium that must be offered to induce an CRRA investor to hold the transaction-cost asset instead of an otherwise identical no-transaction-cost asset is an order of magnitude smaller than the transaction cost rate itself. Constantinides also provides the intuition for this result. Investors respond to transaction costs by turning over the portfolio much less frequently than annually, and since the value function is insensitive to quite large deviations from the optimal no-transaction-cost portfolio allocation, the overall impact of transaction costs on utility is small. Our paper bridges the gap between this theoretical result and the empirical magnitude of the liquidity premium by examining dynamic portfolio choice with transaction costs in a variety of more elaborate settings that move the problem closer to the one solved by real-world investors. In particular, we allow returns to be predictable and transaction costs to be stochastic, and we introduce wealth shocks, both stationary multiplicative and labor income. With predictable returns, we also allow the wealth shocks and transaction costs to be state dependent.

We find that adding these real world complications to the canonical problem can cause transactions costs to produce per-annum liquidity premia that are no longer an order of magnitude smaller than the rate, but are instead the same order of magnitude. In particular, return predictability and i.i.d labor income calibrated to U.S. data are sufficient to obtain per-annum liquidity premia of the same order of magnitude as the cost rate, for realistic wealth income ratios. For this reason, our results provide an important new insight into the effect of transaction costs on investor behavior. In particular, the effect of proportional transaction costs on the standard consumption and portfolio allocation problem with i.i.d. returns can be materially altered by reasonable perturbations that bring the problem closer to the one investors are actually solving. Moreover, if the agent we consider can be regarded as the marginal investor, then the implication is that equilibrium liquidity premia can be the same order of magnitude as the cost rate, once the theoretical economy is complicated

¹See, for example, Brennan and Subrahmanyam (1996) and Pastor and Stambaugh (2003).

along the dimensions we consider. Further, so long as the agent we consider is one who holds low liquidity stocks (which we recognize is a big if), then the liquidity premium we report can be viewed as a lower bound on the equilibrium liquidity premium for the economy that the agent inhabits.

Our canonical agent has power utility with a relative risk aversion coefficient of 6 and has access to a risky asset whose return is calibrated to that of a portfolio of low liquidity stocks. This portfolio is formed from the 13 least liquid of 25 liquidity sorted U.S stock portfolios (see Acharya and Pedersen, 2002). Proportional one-way transactions costs on this portfolio are set at 2%, as a lower bound on the difference between the cost rate for low and high liquidity portfolios formed by bifurcating on the basis of liquidity. We rely on recent empirical work by Lesmond, Ogden and Trzcinka (1999) who quantify the transaction costs associated with trading individual stocks and find a 3% cost for the 5 smallest size deciles and a 1% cost for the 5 largest.

We find that return predictability calibrated to that in the data increases the liquidity premium on the low liquidity portfolio by a factor of 5. An important wealth shock is labor income which we calibrate with a permanent and a temporary component, as in Carroll (1996, 1997) and Viceira (1997), using parameter estimates obtained from U.S. PSID data by Gakidis (1997). In the i.i.d. return case with a fixed transaction cost rate, we find that liquidity premia are unaffected by excluding the temporary component. Consequently, we ignore the temporary component when examining the effects of labor income on liquidity premia. With i.i.d. returns and a fixed transaction cost rate, the inclusion of labor income causes the liquidity premium on the low liquidity portfolio to become 1.42% for an agent with no financial wealth, an almost 18-fold increase relative to the canonical i.i.d. case. Even if the agent's wealth to permanent monthly income ratio is 100, the increase is still 3.4-fold. Once returns are allowed to be predictable (the transaction cost rate remains a constant), i.i.d. labor income causes the liquidity premium to be 1.68% and 1.20% for an agent with a wealth to monthly labor income ratio of 0 and 10 respectively. These represent 21-fold and 15-fold increases, respectively, relative to the standard i.i.d. case. The premium does drop to 0.61% per annum once the wealth to permanent monthly income ratio hits 100 but it can be made as high as 0.87% by allowing labor income growth to be state-dependent.

We also consider multiplicative wealth shocks, for which comparative static analysis is easier because of the lighter computational burden. The liquidity premium also increases by a factor of 5 when i.i.d. wealth shocks to portfolio wealth are introduced with an annualized volatility of 17% and a mean of zero. The simultaneous presence of return predictability and wealth shocks can generate liquidity premia up to 42 times larger than in the standard i.i.d. return case. When

returns are predictable, the unconditional distribution of the wealth shocks is always the same as that for the i.i.d. shocks. The liquidity premium is largest when the wealth shock has negative unconditional correlation with the risky asset's expected return in the current period. The intuition is as follows. The implied hedging demand with respect to future labor income is negative when this correlation is negative. The resulting lower average holding means a higher premium for the same amount of trading. The liquidity premium is also large when the wealth shock has positive unconditional correlation with the risky asset's expected return in the current period: a 31-fold increase relative to standard case. The intuition is as follows. Since unexpected windfalls and unexpected outflows most likely affect the investor's cash balance, the wealth shock is modelled as affecting the investor's riskless asset holding. Hence, a positive wealth shock causes the inherited allocation to the risky asset to decrease. So a positive unconditional relation between the wealth shock and conditional expected risky-asset return means that the inherited allocation to the risky asset is low because of the wealth shock precisely when the investor wants to hold the risky asset because of its high conditional expected return. The result is a particularly large liquidity premium.

We also allow the proportional transaction cost rate to be state dependent. The unconditional mean of this rate is always fixed at 2%. Intuition suggests that wealth shocks are especially painful if negative wealth shocks occur when the transaction cost rate is high. Our results confirm this intuition with the liquidity premium always decreasing in the unconditional correlation between wealth shocks and the transaction cost rate, holding all else equal. Similar intuition suggests that transaction costs are especially painful if the rate is high when current and future conditional expected returns are low since this is when the investor is poor in utility terms. A second channel recognizes that extreme expected return states provoke more rebalancing and so a strong relation between the transaction cost rate in either direction means the investor is rebalancing most when the transaction cost rate is high. A final channel recognizes that if the cost rate is usually high when the agent is at either no short-selling boundary, then the liquidity premia will be low. We find that the first and last channels are most important when the wealth shock covaries positively with expected return, the second dominates when these two are uncorrelated, and the third is the most important when the two are negatively correlated. Since the sign of the hedging demand with respect to future labor income is the same as the sign of this correlation, negative correlation means the zero risky-asset weight constraint binds when dividend yield is low. So a high transaction cost rate when the yield is low means a low liquidity premium, and this is exactly what we find when the correlation between the dividend yield and the the wealth shock is negative, with the premia

declining as the correlation between the cost rate and dividend yield declines.

When the multiplicative wealth shock covaries negatively with expected return and the transaction cost rate, and these latter two covary together positively, the liquidity premium is higher relative to the standard i.i.d. return case by a factor of 57: a liquidity premium of 0.079% per annum in the standard i.i.d. return case becomes 4.55%, which is a number that is in the same ballpark as the numbers reported in empirical work. Smaller liquidity premia are obtained for the labor income problem with these correlations, most likely because the zero risky-asset allocation constraint is binding in many states early in life. However, we still get a liquidity premium of 1.14% for an agent with no financial wealth and one of 0.80% for an agent whose wealth to permanent monthly income ratio is 100. These represent 14-fold and 10-fold increases, respectively, relative to the standard i.i.d. case. Interestingly, when we estimate these correlations using U.S. data, we find exactly this pattern of correlations between the three variables. In the multiplicative wealth shock model, the liquidity premium becomes 1.56% when the correlations are set to the data point estimates, which are admittedly very noisy estimates of the true correlations, especially since aggregate labor income data is being used.

One concern is that the large liquidity premia and utility costs that we document are not robust to the presence of a second low-transaction-cost asset that is positively correlated with the high-transaction-cost asset. To address this concern, we allow the investor access to a high liquidity portfolio in addition to the low liquidity one. The joint return process for the two assets is again calibrated using U.S. stock portfolios formed by sorting on liquidity (see Acharya and Pedersen, 2002). In practice, investors face transaction costs on both portfolios but intuition suggests that the spread in transaction costs across the two portfolios is what is critical for the spread in expected return across the two. Consequently, we set the transaction cost rate on the second asset to zero and keep the rate on the low-liquidity portfolio at 2%. As expected, the presence of the second asset ameliorates the liquidity premium on the low-liquidity asset, but for many of the scenarios considered, the reduction is quite modest. For example, in the case in which the multiplicative wealth shock covaries negatively with expected return and the transaction cost rate, and these latter two covary together positively, the liquidity premium on the low-liquidity asset is still 4.13% per annum in the presence of a second risky asset. This number is similar to the 4.55% per annum obtained when only the low-liquidity asset is available. Similar results are obtained when the other risky asset is the market portfolio rather than the high liquidity portfolio. Thus, the large liquidity premia that we report appear to be robust to the introduction of a second risky asset

whose correlation with the low-liquidity portfolio mirrors that of the high-liquidity portfolio in the data.

Since we are concerned with the cross-section, it is useful to generalize the concept of liquidity premium for an asset having a given cost rate rather than a lower cost rate, where the lower rate may be larger than zero. The premium is defined to be the required reduction in the mean return for an otherwise identical lower-cost asset so that an agent with access to the given-cost asset would prefer not to switch to the lower-cost asset if it was offered instead. As we explained above, we use a cost rate of 2% because it represents the difference in transactions cost rates across large and small firms and then we calculate the premium for the low liquidity portfolio comparing this 2% rate to a rate of 0%. However, this premium may not be the same as the one obtained when the low liquidity portfolio with a cost rate of 3% is compared to the low liquidity portfolio with a cost rate of 1%. When we check, we find that the premia in all cases are lower when the 3% to 1% comparison is performed, but the ratios of the premia relative to the premium in the standard case are similar. We are still able to generate premia as high as 1.97% with a constant cost rate and as high as 2.56% when we allow the rate to be stochastic.

We also examine the utility cost of transaction costs which we define to be the fraction of her wealth that the investor would be willing to give up to be able to trade the low-liquidity asset without incurring the transaction costs. We find the pattern of these utility costs largely mirrors the pattern for the liquidity premia. The only exception is the effect of adding a second risky asset which attenuates this utility cost to a much larger extent than it does the liquidity premium.

Another variable of interest is turnover. Using simulations, we find that annual turnover goes from about 4% in the standard case to 163% when the correlation between expected return and the multiplicative wealth shock is negative and the cost rate is constant. Once the cost rate is allowed to be stochastic, the annual turnover can be as high as 144%, which occurs when the multiplicative wealth shock covaries negatively with expected return and the transaction cost rate, and these latter two covary together positively. Recall that these two cases with negative correlation between expected return and the wealth shock are also cases with large liquidity premia. When the correlations are set to the data point estimates, the annual turnover becomes 73%.

The turnover number can be multiplied by the average cost to get the direct effect of transaction costs on expected return. The extent to which the liquidity premium exceeds this direct cost can be attributed to some combination of risk premium for trading more when the agent is poor in utility terms, and in the case of a stochastic cost rate, to the agent trading more when the cost rate is high.

Not surprisingly, the direct cost is close to the liquidity premium when there are no wealth shocks and the cost rate is constant. The direct cost as a fraction of the liquidity premium is highest when the expected return and the multiplicative wealth shock are negatively correlated, taking a value as high as 0.97 when the cost rate is a constant. The fraction is smallest when the transaction cost rate covaries negatively with expected return and the multiplicative wealth shock, and these latter two covary together positively; it takes a value of 0.22. The effect of the wealth shock-expected return correlation on this ratio is not surprising since the sign of this correlation determines the sign of the hedging demand with respect to future labor income. When the correlation is negative, the zero risky-asset weight constraint binds in low dividend yield states, which means the agent tends to trade when she is wealthy in utility terms because of high future expected returns and high future labor income growth rates. The converse is true when the correlation is positive.

Increasing the investor’s risk aversion from 6 to 8 increases the liquidity premia for the low-liquidity portfolio in all cases. The liquidity premium gets as high as 5.66% in the case in which the multiplicative wealth shock covaries negatively with expected return and the transaction cost rate, and these latter two covary together positively.²

Clearly, our paper is but a first step toward bridging the gap between the theoretical literature and the empirical work finding large spreads in expected returns for portfolios formed on the basis of liquidity. One important limitation of our analysis is that it is partial equilibrium. Hence, it says nothing about how transaction costs affect equilibrium prices by limiting the ability of agents to risk share.

The paper is organized as follows. Section 2 discusses related literature while section 3 describes the investor’s dynamic optimization problem with predictable returns and transaction costs. Section 4 calibrates the state variables, the asset returns and the parameters to the U.S. economy. Section 5 discusses the liquidity premia and utility cost results while Section 6 discusses the allocation results. Section 7 concludes.

2 Related Literature

A number of recent empirical papers have examined how expected returns vary with measures of liquidity. Brennan and Subrahmanyam (1996) form 25 portfolios forming quintiles on size and then

²We also consider the case in which the investor has access to a single risky asset whose returns and transaction cost rate are calibrated to the U.S. market portfolio. In unreported results, we again find that state-dependent expected returns, wealth shocks and transaction cost rates together produce per-annum liquidity premia of same order of magnitude as the cost rate.

within each size quintile, forming quintiles on the Kyle (1985) inverse measure of market depth, λ , estimated as in Glosten and Harris (1988). They find a 6.6% per annum spread in average abnormal return from the Fama-French (1993) three-factor model between the low- λ and the high- λ quintiles. Rather than sorting on a measure of stock illiquidity, Pastor and Stambaugh (2002) form deciles based on covariance of return with a measure of market liquidity and find a spread in abnormal return between the two decile extremes of 7.5% per annum with respect to a four factor model that accounts for sensitivities to the market, size and book-to-market factors of Fama-French (1993) and a momentum factor. Easley, Hvidkjaer and O'Hara (2002) examine how information-based trading affects asset returns and reports that a difference of 10 percentage points in the probability of information-based trading between two stocks leads to a difference in expected returns of 2.5% per annum. Other papers to examine how expected returns vary with measures of liquidity include Brennan, Chordia and Subrahmanyam (1998), Amihud (2002) and Hasbrouck (2003).

Several theoretical and numerical papers have considered how illiquidity and transaction costs affect asset prices. Early work by Stoll (1978) and Ho and Stoll (1981) examines how a dealer sets the spread given that she faces inventory carrying costs. In a single period setting, Amihud and Mendelson (1986) show how transaction costs can affect expected returns on stocks. In an economy with two classes of agents, Heaton and Lucas (1996) examine numerically how idiosyncratic and uninsurable labor income affects equilibrium expected returns both with and without transaction costs on the riskless and equity assets. In an overlapping generations economy, Vayanos (1998) shows how prices are affected by the presence of transaction costs. In his model, agents have a lifecycle motive for trading and trading behavior is predetermined. Huang (2002) studies an equilibrium model in which agents receive unexpected liquidity shocks and can invest in liquid and illiquid riskless assets. Other papers to examine theoretically how illiquidity and transaction costs affect assets prices include Lo, Mamaysky and Wang (2001), Aiyagari and Gertler (1991) and Acharya and Pedersen (2002). Recent papers also examine multiperiod portfolio choice in the presence of labor income, but none incorporate portfolio rebalancing costs (see Viceira, 1997 and 2001, Cocco, Gomes and Maenhout, 2002, and Gomes and Michaelides, 2002)

3 The investor's portfolio allocation problem with transaction costs and predictable returns

This section lays out the preferences of and constraints faced by the investor. We characterize the optimization problem for a dynamic investor who faces either i.i.d. or predictable returns. We also

describe how the problem is complicated by labor income as the wealth shock. We also describe the solution technique for numerically solving the investor's problem.

3.1 Constraints and preferences

We consider the portfolio allocation between N risky assets and a riskless asset. In all the cases we consider, the investor has access to either one or two risky assets. The investor faces transaction costs that are proportional to wealth. The law of motion of the investor's wealth, W , is given by

$$W_{t+1} = (W_t - c_t)(1 + L_t)(1 - f_t) \left[\alpha_t'(\mathbf{R}_{t+1} - R_t^f \mathbf{i}_N) + R_t^f \right],$$

for $t = 1, \dots, T - 1,$ (1)

where c is consumption, α is an $N \times 1$ vector of portfolio weights in the N risky assets, \mathbf{R} is an $N \times 1$ vector of returns on the N risky assets, R^f is the risk-free rate, L is the wealth shock expressed as the percentage change in wealth as result of the shock, and f is the transactions cost per dollar of portfolio value. The wealth shock L is exogenous and assumed to follow a stationary process. The dollar wealth shock is $(W_t - c_t)(1 + L_t)$. At the terminal date T , $c_T = W_T$ so the investor does not receive a wealth shock at the terminal date. Dollar transaction costs at t are $(W_t - c_t)(1 + L_t)f_t$, and are paid by costlessly liquidating the i th risky and the riskless assets in the proportions α_t^i and $(1 - \alpha_t' \mathbf{i}_N)$.

Let $\hat{\alpha}_t^i$ be the allocation to the i th risky asset inherited from the previous period. Then

$$\hat{\alpha}_t^i \equiv \frac{\alpha_{t-1}^i (W_{t-1} - c_{t-1})(1 + L_{t-1})(1 - f_{t-1}) R_t^i}{W_t} = \frac{\alpha_{t-1}^i R_t^i}{\alpha_{t-1}'(\mathbf{R}_t - R_{t-1}^f \mathbf{i}_N) + R_{t-1}^f}. \quad (2)$$

where $\hat{\alpha}_t$ be the $N \times 1$ vector of these inherited portfolio weights. We assume that consumption at time t is obtained by liquidating costlessly the i th risky asset and the riskless asset in the proportions $\hat{\alpha}_t^i$ and $(1 - \hat{\alpha}_t' \mathbf{i}_N)$. This assumption is not so onerous given the availability of money-market bank accounts and given that equities pay dividends. To the extent that the sum of the risky assets' dividends exceeds the consumption out of the risky asset, c , a dividend reinvestment plan can be used to costlessly reinvest the excess dividend in the risky asset.

We allow returns to be predictable and assume that there exists a "predictive" variable, D which affects the conditional mean of the risky assets' return. We assume D follows a first-order Markov process. For simplicity, the riskless rate is assumed to be constant, and so $R_t^f = R^f$ for every t .

The wealth shock L_t is assumed to affect the riskless asset holding. Consequently the vector of inherited risky asset holdings becomes $\frac{\hat{\alpha}}{1+L_t}$ after the shock. When L_t is positive, the shock

is like a cash inflow and when L_t is negative, the shock is like a loan repayment obligation. The investor sees the wealth shock at t after the consumption decision at t but before the allocation decision. When returns are predictable, we allow the distribution of the wealth shock L_t to be state dependent and thus depend on D_t . However, the wealth shock at t contains no information about future returns $(\mathbf{R}_{t+1}, \mathbf{R}_{t+2}, \dots)$ or about future D values $(D_{t+1}, D_{t+2}, \dots)$.

The transaction cost function f depends on the chosen portfolio weights α , the $N \times 1$ vector of portfolio weights inherited from the previous period, $\hat{\alpha}_t$ and the wealth shock L :

$$f_t = \mathbf{\Phi}_t' \left| \alpha_t - \frac{\hat{\alpha}_t}{1 + L_t} \right|. \quad (3)$$

This specification accommodates transaction costs on an asset that are proportional to the change in the value of the portfolio holding of that asset, as in Constantinides (1986). In general, the $N \times 1$ vector $\mathbf{\Phi}$ has i th element Φ^i which gives the proportional cost rate associated with trading the i th risky asset. However, in all applications considered here, only one asset has a non-zero transaction cost rate and so Φ is used to denote that rate. It is straightforward to modify the transaction cost function to accommodate more elaborate transaction cost functions, like, for example, a cost that is the same fraction of portfolio value irrespective of how much of the asset is traded (see Lynch and Tan, 2003).

We also allow the cost parameters $\mathbf{\Phi}_t$ to be random, with distributions that can be state dependent and thus depend on D_t when returns are predictable. The investor sees the cost parameter realizations at t , $\mathbf{\Phi}_t$, after the consumption decision at t but before the allocation decision at t . Again, the cost parameter realizations for t contains no information about future returns $(\mathbf{R}_{t+1}, \mathbf{R}_{t+2}, \dots)$ or about future D values $(D_{t+1}, D_{t+2}, \dots)$. The advantage of the portfolio allocation decisions being made with knowledge of the realized wealth shock and cost parameters is the rebalancing rule has the same form as in the standard case; i.e., a no-trade region with rebalancing to the boundary (see Lynch and Tan, 2003).

We consider the optimal portfolio problem of a investor with a finite life of T periods and utility over intermediate consumption. Preferences are time separable and exhibit constant relative risk aversion (CRRA):

$$E \left[\sum_{t=1}^T \delta^t \frac{c_t^{1-\gamma}}{1-\gamma} \mid D_1, \hat{\alpha}_1 \right], \quad (4)$$

where γ is the relative-risk-aversion coefficient and δ is the time-discount parameter. Note that the expected lifetime utility depends on the state of the economy at time 1. Further, the inherited

portfolio weight for the i th risky asset $\hat{\alpha}_1^i$ is a state variable whenever the i th element of Φ is greater than zero, since the value of this inherited portfolio weight determines the transaction costs to be paid at time 1. These preferences have been extensively used in empirical work by Grossman and Shiller (1981), Hansen and Singleton (1982), and many others.

3.2 Optimization problem

We define κ as the fraction of wealth consumed and R_W as the rate of return on wealth, after the wealth shock and net of the transaction costs incurred. Given our parametric assumptions, the Bellman equation faced by the investor is given by

$$\frac{a(D_t, \hat{\alpha}_t, t)W_t^{1-\gamma}}{1-\gamma} = \max_{\kappa(D_t, \hat{\alpha}_t, t)} \left\{ \frac{\kappa_t^{1-\gamma}W_t^{1-\gamma}}{1-\gamma} + \delta \frac{(1-\kappa_t)^{1-\gamma}W_t^{1-\gamma}}{1-\gamma} \right. \\ \left. E \left[\max_{\alpha(D_t, \hat{\alpha}_t, L_t, \Phi_t, t)} \left\{ E \left[a(D_{t+1}, \hat{\alpha}_{t+1}, t+1)R_{W,t+1}^{1-\gamma} | D_t, \hat{\alpha}_t, L_t, \Phi_t \right] \right\} | D_t, \hat{\alpha}_t \right] \right\} \\ \text{for } t = 1, \dots, T-1, \quad (5)$$

where $E[.|D, \hat{\alpha}]$ denotes the expectation taken using the conditional distribution given D and the fraction of portfolio value allocated to the risky asset at time t , and $\alpha_t \equiv \alpha(D_t, \hat{\alpha}_t, L_t, \Phi_t, t)$ and $\kappa_t \equiv \kappa(D_t, \hat{\alpha}_t, t)$, both time dependent since the time horizon T is finite. Notice that portfolio choice α_t is allowed to depend on the wealth shock and transaction cost realizations at t but not the consumption choice κ_t . This form of the value function derives from the CRRA utility specification in eqs. (4), and from the linearity in W of the budget constraint, eq.(1). The Bellman eq. (5) is solved by *backward iteration*, starting with $t = T - 1$ and $a(D, \hat{\alpha}, T) = 1$.

If returns are not predictable, D_t is no longer a state variable for the problem and so the a is no longer a function of D_t . Moreover, the expectation in eq. (5) does not depend on D_t and neither does the optimal allocation. So $\alpha(\hat{\alpha}_t, L_t, \Phi_t, t)$ can denote α_t .

3.3 Labor income problem

The wealth shock introduced above is a multiplicative shock to financial wealth, where the shock itself is a stationary variable. This structure makes the wealth shock easy to handle, since its presence has no impact on the number of state variables for the agent's problem. Unfortunately, many wealth shocks are largely unrelated to the magnitude of current financial wealth, making this structure inapplicable one important example is labor income. An interesting question is whether the liquidity premia for realistic labor income processes are similar in magnitude to those for the multiplicative wealth shock processes we consider.

Following Carroll (1996) and (1997), labor income is specified to have both permanent and temporary components:

$$y_t = y_t^P + \epsilon_t, \quad (6)$$

$$g_t \equiv y_t^P - y_{t-1}^P = \bar{g} + b_g d_t + u_{t+1}, \quad (7)$$

where y_t is log labor income received at t , y_t^P is log permanent labor income at t , ϵ_t is log temporary labor income at t , $d_t \equiv \ln(D_t)$, and ϵ_t and u_{t+1} are uncorrelated i.i.d. processes. Consistent with our specification for the multiplicative wealth shock, ϵ_t and u_{t+1} contain no information about future returns ($\mathbf{R}_{t+1}, \mathbf{R}_{t+2}, \dots$) or about future D values (D_{t+1}, D_{t+2}, \dots).

With labor income, the law of motion for the investor's wealth, W , is given by

$$W_{t+1} = (W_t + Y_t - c_t)(1 - f_t) \left[\alpha_t(\mathbf{R}_{t+1} - R_t^f \mathbf{i}_N) + R_t^f \right] \quad \text{for } t = 1, \dots, T-1, \quad (8)$$

where Y_t is labor income received at time- t . At the terminal date T , $c_T = W_T$ so the investor does not receive labor income at the terminal date. Labor income is assumed to be received as the riskless asset, and so transaction cost function f_t becomes:

$$f_t = \Phi_t' | \alpha_t - \frac{\hat{\alpha}_t \Gamma_t}{\Gamma_t + \exp\{g_t + \epsilon_t\}}|. \quad (9)$$

where $Y_{t-1}^P = \exp\{y_{t-1}^P\}$ is permanent labor income at $t-1$ and Γ_t is defined to be $\frac{W_t}{Y_{t-1}^P}$. For a given inherited risky-asset allocation at t $\hat{\alpha}_t$, the post-labor income inherited risky-asset allocation is decreasing in Y_t .

Given this specification of the agent's problem with labor income, the value function at t is homogenous in Y_{t-1}^P and has an additional state variable: the ratio of financial wealth at t to lagged permanent labor income Γ_t . The law of motion for the investor's wealth, W , can be rewritten as

$$\Gamma_{t+1} = (\Gamma_t - \hat{\kappa}_t + \exp\{g_t + \epsilon_t\})(1 - f_t) \exp\{-g_t\} \left[\alpha_t(\mathbf{R}_{t+1} - R_t^f \mathbf{i}_N) + R_t^f \right] \quad \text{for } t = 1, \dots, T-1. \quad (10)$$

where $\hat{\kappa}_t \equiv \frac{c_t}{Y_{t-1}^P}$. Eq.(10) is also the evolution equation for the state variable Γ_t while the evolution equation for other state variable $\hat{\alpha}_t$ remains eq.(2).

Given our parametric assumptions, the Bellman equation faced by the investor is given by

$$\begin{aligned} \frac{a(\Gamma_t, D_t, \hat{\alpha}_t, t)(Y_{t-1}^P)^{1-\gamma}}{1-\gamma} = E \left[\max_{\hat{\kappa}(\Gamma_t, D_t, \hat{\alpha}_t, g_t, \epsilon_t, \Phi_t, t), \alpha(\Gamma_t, D_t, \hat{\alpha}_t, g_t, \epsilon_t, \Phi_t, t)} \left\{ \frac{\hat{\kappa}_t^{1-\gamma} (Y_{t-1}^P)^{1-\gamma}}{1-\gamma} \right. \right. \\ \left. \left. + \delta \frac{(Y_{t-1}^P)^{1-\gamma}}{1-\gamma} E [a(\Gamma_{t+1}, D_{t+1}, \hat{\alpha}_{t+1}, t+1)(\exp\{g_t\})^{1-\gamma} | \Gamma_t, D_t, \hat{\alpha}_t, g_t, \epsilon_t, \Phi_t] \right\} | \Gamma_t, D_t, \hat{\alpha}_t \right], \end{aligned} \quad (11)$$

for $t = 1, \dots, T-1$,

where $\alpha_t \equiv \alpha(\Gamma_t, D_t, \hat{\alpha}_t, g_t, \epsilon_t, \Phi_t, t)$ and $\hat{\kappa}_t \equiv \hat{\kappa}(\Gamma_t, D_t, \hat{\alpha}_t, g_t, \epsilon_t, \Phi_t, t)$. There is one difference between the formulation of this problem and the multiplicative wealth shock problem. In that problem, consumption is not allowed to depend on the current date's wealth shock or transaction cost shock, whereas here it is allowed to depend on the current date's permanent and temporary labor shocks as well as the transaction cost shock.³ As in the multiplicative wealth shock case, the Bellman eq. (11) is solved by *backward iteration*, starting with $t = T-1$ and $a(\Gamma, D, \hat{\alpha}, T) = \Gamma^{1-\gamma}$.

3.4 Liquidity premia and utility comparisons

Each of the investor problems described above imply a policy function that, in turn, yields a particular level of expected lifetime utility. Specifically, for the stationary multiplicative wealth shock problem, the policy functions $\{\alpha(D_t, \hat{\alpha}_t, L_t, \Phi_t, t)\}_{t=1}^{T-1}$ and $\{\kappa(D_t, \hat{\alpha}_t, t)\}_{t=1}^{T-1}$ can be substituted into the actual law of motion for investor's wealth eq. (1) to obtain the consumption sequence $\{c_t = \kappa(D_t, \hat{\alpha}_t, t)W_t\}_{t=1}^T$. This consumption sequence is then substituted into eq. (4) to obtain the investor's expected lifetime utility. Analogous substitutions can be performed for the labor income problem to obtain the investor's expected lifetime utility.

Constantinides (1986) found that for a CRRA investor with access to a risky asset whose return is i.i.d., proportional transaction costs produce per-annum liquidity premia that are an order of magnitude smaller than the cost rate. We are interested in determining whether this result is robust to the introduction of the real-world complications like return predictability, state-dependent wealth shocks and state dependent transaction costs. Consequently, our definition of liquidity premium is in line with that adopted by Constantinides (1986). With only one risky asset, the liquidity premium is defined to be the decrease in the unconditional mean return on this asset that the investor requires to be indifferent between having access to the risky asset without the transaction costs rather than with them. The mean is decreased by subtracting a constant from every state.

³We solved the multiplicative shock problem allowing consumption to depend on the current dates's permanent and temporary labor shocks. The liquidity premia were not materially affected.

We also perform utility comparisons. In particular, we are interested in the fraction of wealth that an investor paying the transaction costs would be prepared to give up to be allowed to trade the risky asset without paying the costs.

As mentioned above, the expected lifetime utility depends on the initial value of the inherited portfolio allocation, $\hat{\alpha}_1$, and the initial value of the vector characterizing the state of the economy, D_1 . For simplicity, we always take the inherited allocation to be the optimal allocation for the analogous no-transaction-cost problem. For those problems with wealth shocks, both stationary multiplicative and labor income, the inherited allocation for a given state is taken to be the average optimal allocation in that state for the analogous no-transaction-cost problem.

3.5 *Solution technique*

The dynamic programming problems are solved by backward recursion. With one or two risky assets, the state variable $\hat{\alpha}_1$ is discretized and the value function is linearly interpolated between $\hat{\alpha}_1$ points. This technique yields an approximate solution that converges to the actual solution as the $\hat{\alpha}_1$ grid becomes finer. In all the optimizations, the holdings of both the risky and the riskless assets are constrained to be non-negative. In the two-risky asset case, this restricts the action space with respect to allocation choice to the triangular region characterized by $\alpha_1 + \alpha_2 \leq 1$, $\alpha_1 \geq 0$ and $\alpha_2 \geq 0$. When actions are restricted to this set, implied inherited allocations for any return realization on the assets are again in the same region. We use a finer discretization for the action space of allocations than the state space of inherited allocations. Allocation choices on each asset available to the investor always include the discrete grid $\{0.000, 0.001, \dots, 0.999, 1.000\}$.

For the labor income problem, the presence of an additional state variable, the wealth to lagged permanent income ratio, considerably complicates the methodology needed to obtain accurate solutions in a manageable time-frame. Appendix A details the methodology employed.

3.6 *Turnover and Direct Trading Cost*

Another variable of interest is turnover. Turnover is calculated for the multiplicative wealth shock problem by simulating 100,000 paths and applying the optimal policies to each path. Per-annum turnover is defined to be:

$$\frac{12}{240} \sum_{t=1}^{240} \frac{\text{av}(|A_t - \hat{A}_t|)}{\text{av}(\max(A_t, \hat{A}_t/(1 + L_t)))}, \quad (12)$$

where $\text{av}(\cdot)$ is obtained by taking the average across the simulation paths, $A_t = \alpha_t W_t (1 - \kappa_t)(1 + L_t)$ and $\hat{A}_t = \hat{\alpha}_t W_t (1 - \kappa_t)$. So A_t is the dollar chosen risky-asset holdings at time t and \hat{A}_t is the dollar

effective inherited risky-asset allocation at time t . The fraction being summed in (12) represents turnover for a given month of life. The denominator must be some measure of the dollar risky-asset holdings for the month. We take the maximum of the dollar chosen risky-asset allocation and the dollar effective inherited risky-asset allocation for the month.⁴ This measure of turnover equally weighs the 240 months of the agent’s life.⁵

The turnover number can be multiplied by the average transaction cost rate to get the direct effect of transaction costs on expected return. The extent to which the liquidity premium exceeds this direct cost can be attributed to some combination of risk premium for trading more when the agent is poor in utility terms, and in the case of a stochastic cost rate, to the agent trading more when the cost rate is high.

4 Calibration

This section describes how the return processes are calibrated to data, the choice of predictive variable, and the parameter value choices.

4.1 *Return calibration*

We use the one-month Treasury-bill rate to obtain a proxy for the risk-free rate, we use the 12-month dividend yield on the value-weighted NYSE index as a proxy for the predictive variable D , and we use a risky-return dataset that consists of three portfolios, a high and a low liquidity portfolio and the market portfolio. The Acharya and Pedersen (2002) data set provides 25 value-weighted portfolios of NYSE and AMEX stocks sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The high liquidity portfolio is the value-weighted portfolio of the most liquid 12 portfolios and the low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. The market portfolio is the monthly rate of return on the value-weighted NYSE index. The data period is from February 1964 to December 1996. The stock returns and interest rate series are deflated using monthly CPI inflation. The continuously compounded riskfree rate is estimated to be the mean of the continuously compounded real one-month Treasury-bill rate over this period, which gives values for R^f of 0.110 percent.

We estimate a VAR for two sets of portfolio returns. For each, assume that \mathbf{R} is an $N \times 1$

⁴We also tried using a measure of holding based on the average of the chosen risky-asset allocation and the effective inherited allocation. The results were virtually identical.

⁵We tried other weighting schemes including one that weighs months according to the average dollar value of holdings over the months. The results were qualitatively similar to the ones we report.

return vector and let $\mathbf{r} \equiv \ln(\mathbf{1} + \mathbf{R})$ and $d \equiv \ln(1 + D)$. We assume that $[\mathbf{r}'d]'$ follows the vector autoregressive model (VAR):

$$\mathbf{r}_{t+1} = \mathbf{a}_r + \mathbf{b}_r d_t + \mathbf{e}_{t+1}, \quad (13)$$

$$d_{t+1} = a_d + b_d d_t + v_{t+1}, \quad (14)$$

where \mathbf{a}_r , $N \times 1$, and a_d are intercepts, \mathbf{b}_r , $N \times 1$, and b_d are coefficients and $[\mathbf{e}'v]'$ is an i.i.d. vector of mean-zero, multivariate normal disturbances, with *constant* covariance matrix $\Sigma_{\mathbf{e}v, \mathbf{e}v}$; the covariance matrix of \mathbf{e} is $\Sigma_{\mathbf{e}, \mathbf{e}}$ and the variance of v is σ_v^2 . Similarly, the unconditional covariance matrix $[\mathbf{r}'d]'$ is $\Sigma_{\mathbf{r}d, \mathbf{r}d}$; the unconditional variance matrices for \mathbf{r} and d are $\Sigma_{r, r}$ and σ_d^2 respectively. Without loss of generality, we normalize the mean of d , μ_d , to be zero and its variance, σ_d^2 , to be 1. Specification (13)-(14) assumes that d_t is the only state variable needed to forecast r_{t+1} which is in line with other papers on optimal portfolio selection (e.g., Barberis (2000) and Campbell and Viceira (1999)).

The data VAR is estimated using ordinary least squares (OLS) and discretized using a variation of Tauchen and Hussey's (1991) Gaussian quadrature method; the variation is designed to ensure that d is the only state variable (see Balduzzi and Lynch (1999) for details). However, following Lynch (2000), this study implements the discretization in a manner that produces exact matches for important moments for portfolio choice.⁶ We choose 19 quadrature points for the dividend yield and 3 points for the stock-return innovations since Balduzzi and Lynch (1999) find that the resulting approximation is able to capture important dimensions of the return predictability in the data.

Table 1 present data and quadrature VAR parameter values for the high and low liquidity returns while Table 2 presents the same information for the market and low liquidity returns. In each table, Panel A reports the slope coefficients \mathbf{b}_r and b_d as well as unconditional means for \mathbf{r} and d . Panel B reports the unconditional covariance matrix for $[\mathbf{r}'d]'$ and the cross-correlations. Panel C reports the unconditional covariance matrix for $[\mathbf{e}'v]'$ and the cross-correlations. Throughout both tables, the quadrature values almost always replicate the data values, which suggests that the discretization is capturing the important features of the data.

⁶In particular, the procedure matches both the conditional mean vector and the covariance matrix for log returns at all grid points of the predictive variables, as well as the unconditional volatilities of the predictive variables and the correlations of log returns with the predictive variables.

4.2 *Parameter choices*

The investor's risk aversion parameter, γ , is set to six almost always, while a comparative static analysis increases its value to eight. These γ choices are motivated by the Mehra and Prescott (1985) argument that the existing evidence from macro and micro studies constitutes an *a priori* justification for restricting the value of γ to be less than ten. The horizon T of the young investor is 240 periods or 20 years, since the return processes are calibrated to monthly returns. A 20-year horizon is a realistic investment horizon for an investor who retires at time 1. The time preference parameter, δ is set equal to the inverse of the riskfree return.

When calibrating the transaction cost rates, it is the cost of trading the individual stocks and not the portfolio itself that is relevant. The transaction cost rate on the low liquidity portfolio is calibrated to the transaction cost spread between stocks in a low transaction cost portfolio and stocks in a high transaction cost portfolio. Lesmond, Ogden and Trzcinka (1999) form size deciles and then report the average round-trip transaction cost for the individual stocks in each decile. According to Table 3 of Lesmond, Ogden and Trzcinka (1999), the average round-trip cost of trading a stock in the 5 largest portfolios less the average to trade one in the 5 smallest equals 4.01%, and so we take 2% to be the one-way transaction cost rate for the low liquidity portfolio. This number is likely to be a ballpark figure for the transaction cost spread between the low and high liquidity portfolios and a lower bound for the cost of trading the low liquidity portfolio. Turning to the value-weighted market, Lesmond, Ogden and Trzcinka (1999) report an average round-trip cost of 2.26% to trade a stock in the 5 largest deciles, so we take 1% to be the one-way transaction cost rate to trade the market portfolio. This rate likely represents a lower bound on the average cost to trade a stock in the value-weighted market, though we are equal weighting across the 5 portfolios. Throughout, the high liquidity portfolio is only available when the low liquidity portfolio is also available and moreover, the cost rate on the high liquidity portfolio is always taken to be zero. In practice, investors face transaction costs on both portfolios, but intuition suggests that the spread in transaction costs across the two portfolios is what is critical for the spread in expected return across the two. Keeping the transaction cost rate on one of the risky assets equal to zero keeps the inherited allocation state space one-dimensional, which keeps computation time manageable.

When we allow the cost rate, Φ , to be stochastic, we always keep its unconditional mean μ_Φ equal to its value in the analogous case when it's a constant. We apply Gaussian quadrature rules to $\phi \equiv \ln(1 + \Phi)$ in such a way that the volatility of ϕ , σ_ϕ , and the mean of Φ match chosen values. In particular, for a given μ_Φ , the unconditional volatility of σ_ϕ is always chosen such that

$\frac{\mu_\Phi}{\sigma_\phi}$ equals 2.63. The three values taken by the ϕ shock are always the same and 2.63 is chosen to ensure a large spread in Φ but without any negative values. When the risky return is predictable, the conditional expectation of ϕ is linear in the d value.

The wealth shock, L , is always mean zero unconditionally and in most cases, the unconditional volatility of $l = \ln(1 + L)$, σ_l , is calibrated to be 0.05. Gaussian quadrature is applied to l and ϕ to obtain their joint distribution, with the conditional expectation of l also linear in the d value. Three grid points are used for the l shock. Fixing the unconditional volatility of the wealth shock (if present) and the unconditional mean and volatility of the transaction cost rate (if stochastic), the environment facing the investor depends on: 1) Presence or not of wealth shocks. 2) Transaction cost rate stochastic or not. 3) Contemporaneous correlation between the wealth shock and the transaction cost rate, $\rho_{l,\phi} \equiv \rho[l_t, \phi_t]$. 4) Contemporaneous correlation between the wealth shock and dividend yield, $\rho_{l,d} \equiv \rho[l_t, d_t]$. 5) Contemporaneous correlation between the transaction cost rate and dividend yield, $\rho_{d,\phi} \equiv \rho[d_t, \phi_t]$.

A monthly volatility of 5% for the wealth shock translates into an annual volatility of about 17%, which someone could just as easily think is small or large. Consequently, for some cases, we will report liquidity premia as a function of this volatility, to provide some information about how sensitive the results are to the choice of σ_l value. Further, it is not easy to calibrate the wealth shock and cost rate correlations to data, though we make an attempt. Because of this difficulty, the investor's problem is evaluated over a full grid of values for these parameters for each $\{\mu_\Phi, \sigma_\phi, \sigma_l\}$ triplet considered.

Turning to the labor income process, parameter values are chosen to be the baseline values in Viceira (1997, 2001) who describes these values as consistent with those obtained by Gakidis (1997) based on PSID data for professionals and managers not self-employed under age 45. Liquidity premia given these values are of interest since this is a population group that holds stocks. Viceira's baseline value for the standard deviation of the change in log permanent labor income is 15% per year, for the standard deviation of the log temporary shock is 10% per year and for the mean growth of permanent labor income is 3% per annum. These are used to back out values for the volatility of ϵ , σ_ϵ , the unconditional volatility of g , σ_g , and the unconditional mean of g_t , \bar{g} . As in Viceira, the multiplicative temporary shock $\exp\{\epsilon_t\}$ has a mean of 1. A number of papers (see, for example, Chamberlain and Hirano, 1997 and Carroll and Samwick, 1995) have estimated labor income parameters and a range of values are reported across these studies. However, the Gakidis values seem to lie within this range, which makes them reasonable to use.

The temporary shock ϵ , when present, is assumed to be orthogonal to all other shocks and its distribution is obtained using three Gaussian quadrature grid points. Gaussian quadrature is used to deliver the joint distribution for (g_t, ϕ_t, d_{t+1}) in exactly the same way as it is used to obtain a joint distribution for (l_t, ϕ_t, d_{t+1}) in the multiplicative wealth shock case. Unconditional moments for ϕ are chosen as before. Taking the presence of labor income as given and fixing the unconditional mean and volatility of permanent labor income growth and the transaction cost rate (if stochastic), the environment facing the investor depends on: 1) Presence or not of a temporary labor income shock. 2) Transaction cost rate stochastic or not. 3) Contemporaneous correlation between permanent labor income growth and the transaction cost rate, $\rho_{g,\phi} \equiv \rho[g_t, \phi_t]$. 4) Contemporaneous correlation between permanent labor income growth and dividend yield, $\rho_{g,d} \equiv \rho[g_t, d_t]$. 5) Contemporaneous correlation between the transaction cost rate and dividend yield, $\rho_{d,\phi} \equiv \rho[d_t, \phi_t]$.

4.3 U.S. Data: Correlations and Transaction Cost Rate Volatility

This section presents some results for U.S. data regarding the correlations between dividend yield, the transaction cost rate and the growth rate of labor income, as well as the volatility of the cost rate. The point estimates we obtain are quite noisy but they can be used to frame the discussion of the liquidity premia results in the next section.

The dividend yield series that we use is the one described above. But we need proxies for the log growth in permanent labor income, g , and the transaction cost rate ϕ . For correlations, it is enough to have a proxy that is a linear function of the variable of interest. The log aggregate income series constructed by Lettau and Ludvigson (2001) is first-differenced to obtain a proxy for the log growth in permanent labor income, g . While log growth in aggregate labor income is a noisy proxy for log growth in individual permanent labor income, the hope is that the idiosyncratic component of g correlates with d and ϕ in a manner that is similar to way the systematic component correlates with these variables. Likely the sign of the data correlation provides information about the sign of the correlation for individual permanent labor income growth, but its magnitude probably represents a loose upper bound on the magnitude for individual permanent labor income growth since the idiosyncratic component of g is likely much less correlated with the dividend yield and the transaction cost rate than the systematic component. The data is quarterly.

Two proxies for ϕ are used. The first is only available at an annual frequency but it can be used to calculate volatility. Hasbrouck (2003) employs a Bayesian approach to estimating half-spreads using daily data for all ordinary common equity issues on the CRSP daily database from 1962 to

2002 and averages the daily numbers to obtain half-spread estimates at an annual frequency for each stock. Fixing year, we sort the stocks on half-spread and equally weight the top and the bottom 50 % to obtain annual time series estimates for two hypothetical portfolios (liquid portfolio and illiquid portfolio respectively). If the half-spread is greater than 15% for an observation, we set that observation to 15%, analogous to the treatment of outliers in Amihud’s ILLIQ measure by Acharya and Pedersen (2002). We construct a third series by subtracting the liquid portfolio half-spread from the illiquid portfolio half-spread, which gives us our proxy for Φ . The second proxy is available at a monthly frequency but it is a less direct measure of the per-trade cost than Hasbrouck’s measure so we only use it to calculate correlations. The Acharya and Pedersen (2002) data set, in addition to monthly returns on 25 value-weighted portfolios sorted on Amihud’s ILLIQ measure, also provides a corresponding normalized ILLIQ for each portfolio, which they argue can be interpreted as a one-way transactions cost rate. We value-weight to obtain monthly series of one-way proportional transactions cost rates for the low and the high liquidity portfolios. Again, we construct a third series by subtracting the high liquidity portfolio cost rate from the low liquidity portfolio cost rate. We implement the transformation $\phi = \log(1 + \Phi)$ on all series.

Table 3 presents means, standard deviations and correlations for a data period from February 1964 to December 1996. Panel A reports the unconditional means and the standard deviations (in percent) of the three Hasbrouck series. The volatility of the half-spread difference is 7.08% which is much larger than the 0.76% used in the liquidity premia calculations. This result suggests that we are using a conservative volatility number for ϕ when we allow ϕ to be stochastic. Interestingly, the mean of 5.44% is also larger than the 2% used in the liquidity premia calculations.

Panel B reports the correlations, taking d to be the dividend yield at the start of the quarter and ϕ to be the average of the three ϕ s over the quarter. Again we focus on the cost rate differential across the low and high liquidity portfolios. We find that the correlation of g with d is negative, consistent with the intuition that labor income is lower during recessions. We also find that ϕ is positively correlated with d , which is consistent with the idea that stocks are more expensive to trade in recessions. The correlation between ϕ and g is negative. We realize that the series we are using to proxy g and ϕ are quite noisy and so the reported correlations are likely to be noisy estimates of the true correlations. But these numbers should at least be informative as to the direction of the relations and may even contain some information as to the magnitudes.

5 Liquidity premia and utility cost results

5.1 *One risky asset*

Table 4 reports annual liquidity premia for a power utility agent with a relative risk aversion coefficient of 6 and who only has access to the low liquidity portfolio. The proportional one-way transactions cost rate on this portfolio has an unconditional mean of 2%. Panel A tabulates constant transactions costs cases while Panel B tabulates random transactions costs cases. The first subpanel of each panel considers the case with no wealth shocks while the second subpanel introduces wealth shocks that always have an unconditional mean of zero and an unconditional volatility of 5% per month.

The top left entry in Panel A reports a per annum liquidity premium of 0.079% for the canonical allocation problem with i.i.d. returns, a constant transaction cost rate and no wealth shocks. Since the transaction cost rate is 2%, this premium is an order of magnitude smaller than the rate, consistent results in Constantinides (1986). The top right entry in Panel A reports the liquidity premium when return predictability calibrated to that in the data is introduced. The reported premium is 0.463% per annum, which is a 5-fold increase relative to the canonical case. As the left entry in the second subpanel of Panel A indicates, the liquidity premium also increases by a factor of 5 when i.i.d. wealth shocks to portfolio wealth are introduced.

Turning to the right entries in the second subpanel of Panel A, the simultaneous presence of return predictability and wealth shocks can generate liquidity premia up to 42 times larger than in the standard i.i.d. return case. The unconditional distribution of the wealth shocks is always the same as that for the i.i.d. shocks. Taking a value of 3.37%, the liquidity premium is largest when the wealth shock has negative unconditional correlation with the risky asset's expected return in the current period. The intuition is as follows. The implied hedging demand with respect to future labor income is negative when this correlation is negative. The resulting lower average holding means a higher premium for the same amount of trading. The liquidity premium is also large when the wealth shock has positive unconditional correlation with the risky asset's expected return in the current period: the 2.51% represents a 31-fold increase relative to standard case. The intuition is as follows. Since unexpected windfalls and unexpected outflows most likely affect the investor's cash balance, the wealth shock is modelled as affecting the investor's riskless asset holding. Hence, a positive wealth shock causes the inherited allocation to the risky asset to decrease. So a positive unconditional relation between the wealth shock and conditional expected risky-asset return means that the inherited allocation to the risky asset is low because of the wealth shock precisely when

the investor wants to hold the risky asset because of its high conditional expected return. The result is a particularly large liquidity premium.

Since the liquidity premium is higher for $\rho_{l,d}$ equal to -0.5 or 0.5 than for $\rho_{l,d}$ equal to zero, it is of interest to ask what $\rho_{l,d}$ value minimizes the liquidity premium. This is especially important since the data point estimate of -0.21 suggests a negative value in the data that is smaller than -0.5. Figure 1 plots the relation between the liquidity premia and $\rho_{l,d}$ and we see that the premium is minimized at a small negative $\rho_{l,d}$ value. The premium is about 1.5% at the data point estimate for $\rho_{l,d}$ and its minimum value is still more than 0.9%.

We also allow the proportional transaction cost rate to be state dependent and the results are reported in Panel B. The first subpanel reports premia for cases with no wealth shocks and we see that randomness in the cost rate alone at best produces very small increases in the premia and often produces decreases. Once wealth shocks are allowed, intuition suggests that they are especially painful if negative shocks occur when the transaction cost rate is high. Our results in the second subpanel confirm this intuition with the liquidity premium always decreasing in the unconditional correlation between wealth shocks and the transaction cost rate, holding all else equal. This is done by moving from left to right across a row.⁷

Similar intuition suggests that transaction costs are especially painful if the rate is high when current and future conditional expected returns are low since this is when the investor is poor in utility terms. A second channel recognizes that extreme expected return states provoke more rebalancing and so a strong relation between the transaction cost rate in either direction means the investor is rebalancing most when the transaction cost rate is high. A final channel recognizes that if the cost rate is usually high when the agent is at either no short-selling boundary, then the liquidity premia will be low. The results in the last 5 rows of the second subpanel of Panel B (with $\rho_{l,d} = 0.5$) show that the first and last channels are more important when the wealth shock covaries positively with expected return, while the second channel dominates when these two are uncorrelated, as the middle five rows (with $\rho_{l,d} = 0$) show.

The third channel is the most important when the two are negatively correlated, as shown by the first 5 rows (with $\rho_{l,d} = -0.5$). Since the sign of the hedging demand with respect to future labor income is the same as the sign of this correlation, negative correlation means the zero risky-asset weight constraint binds when dividend yield is low. So a high transaction cost rate when the

⁷This effect is related to pricing results in the Acharya and Pedersen (2002) model. In their economy, investors hold the market each period and liquidate their portfolios at the end of each period. In equilibrium, there is a negative relation between an asset's expected return and the covariance of its transaction costs with the market return.

yield is low means a low liquidity premium, and this is exactly what we find when the correlation between the dividend yield and the the wealth shock is negative, with the premia declining as the correlation between the cost rate and dividend yield declines.

It is worth looking at the case that delivers the highest premium to get some idea as to how far one can go with the three complications that we consider. When the multiplicative wealth shock covaries negatively with expected return and the transaction cost rate, and these latter two covary together positively, the liquidity premium is higher relative to the standard i.i.d. return case by a factor of 57: a liquidity premium of 0.079% per annum in the standard i.i.d. return case becomes 4.55%, which is a number that is in the same ballpark as the numbers reported in empirical work. Interestingly, this is exactly the pattern of correlations between the three variables that we reported in Table 3 using U.S. data. When the correlations are set to the data point estimates, which are admittedly very noisy estimates of the true correlations, the liquidity premium is 1.56% (see Panel B of Table 9), which is the same order of magnitude as the empirical numbers.

Table 5 is organized in exactly the same way as Table 4 and considers the exact same set of cases. The only difference is that Table 4 reports liquidity premia while Table 5 reports utility costs. Interestingly, we find the pattern of these utility costs largely mirrors the pattern for the liquidity premia. We are also interested in the magnitude of the fraction of her wealth that the investor would be willing to give up to be able to trade the low-liquidity asset without incurring the transaction costs. While this fraction is only 0.354% in the canonical case, it gets as high as 34.59% in the case with $\rho_{\phi,l} = -0.5$, $\rho_{l,d} = 0.5$ and $\rho_{\phi,d} = -1$.

Table 6 reports liquidity premia for the low liquidity asset in environments with labor income. Panel A reports results for a fixed transaction cost rate of 2%. The first two columns of Panel A report premia for i.i.d. returns, while the last four report premia when returns are predictable. As the ratio of wealth to lagged permanent labor income, Γ , goes to infinity, the agent's problem converges to the analogous problem with no labor income. This explains why the $\Gamma = \infty$ row contains the same values as reported in Table 4 for cases with a fixed transaction cost and no wealth shock: 0.079% in the i.i.d. return case and 0.463% in the predictable return case.

To assess whether the temporary component of labor income has an effect on liquidity premia, we calculate liquidity premia in the i.i.d. case, with and without the temporary component set to zero. The results are in the two i.i.d columns, with the second setting the temporary component to zero. The liquidity premia are indistinguishable across the two columns. This result suggests that presence or not of the temporary component is irrelevant and for this reason we set it to zero for

all other cases considered. The i.i.d. columns also show that the inclusion of labor income, when returns are i.i.d. and the transaction cost rate is fixed, causes the liquidity premium on the low liquidity portfolio to become 1.42% for an agent with no financial wealth, an almost 18-fold increase relative to the canonical i.i.d. case. Even if the agent's wealth to permanent monthly income ratio is 100, the increase is still 3.4-fold, which is comparable to the 5-fold increase we reported above when we include the stationary multiplicative shock.

Turning to the predictable return cases, the correlation between the growth in permanent labor income and expected return is -0.5 in the fourth last column of panel A, -0.21, the data value, in the third last column, zero in the second last column and 0.5 in the last. For Γ s of 10 and above, the premium is higher when the correlation is 0.5 than when it is zero, and for Γ s of 100 and above, the premium is higher when the correlation is -0.5 or -0.21 than when it is zero. These findings are similar to that reported in Table 4 for the case with a stationary multiplicative wealth shock, and the intuition is similar as well. On the other hand, for low Γ s, the liquidity premium is lower when the correlation is either 0.5, -0.21 or -0.5 rather than 0, which is a result that is much more difficult to explain.

Shifting the focus to the magnitude of the liquidity premia when returns are allowed to be predictable, the four columns show that the liquidity premium can be as high as 1.68% for an agent with no financial wealth and as high as 0.87% for an agent whose wealth to permanent monthly income ratio is 100. These represent 21-fold and 11-fold increases, respectively, relative to the standard i.i.d. case. Again these increases are comparable (though smaller) than those obtained by introducing return predictability and stationary multiplicative wealth shocks. When the correlation between the expected return and the log growth in permanent labor income is -0.21, the per-annum liquidity premium is 1.02% for $\Gamma = 0$ and only drops to 0.62% for $\Gamma = 100$, which are still roughly the same order of magnitude as the transaction cost rate.

Panel B reports results for cases with predictable returns and a transaction cost rate that is state dependant. The cases reported are the two that generates the two largest premia when the wealth shock is stationary multiplicative plus the data correlations: $\rho_{g,d} = -0.5, \rho_{\phi,d} = 1, \rho_{\psi,d} = -0.5$; $\rho_{g,d} = 0.5, \rho_{\phi,d} = -1, \rho_{\psi,d} = -0.5$; and the data correlations $\rho_{g,d} = -0.21, \rho_{\phi,d} = 0.11, \rho_{\psi,d} = -0.18$. For the case with $\rho_{g,d} = -0.5, \rho_{\phi,d} = 1, \rho_{\psi,d} = -0.5$, which is the same correlation pattern as the data, the per-annum liquidity premia is 0.80% when the agent's Γ equals 100, and it climbs to 1.14% when Γ equals 0. Notice that this premium is much smaller in magnitude than the 4.55% obtained in the analogous multiplicative wealth shock case. For the case with data correlations $\rho_{g,d} =$

$-0.21, \rho_{\phi,d} = 0.11, \rho_{\phi,d} = -0.18$, the per-annum liquidity premia is 0.64% when the agent's Γ equals 100, and it climbs to 1.21% when Γ equals 0. The premia are slightly higher for the third case ($\rho_{g,d} = 0.5, \rho_{\phi,d} = -1, \rho_{\phi,d} = -0.5$) with a per-annum liquidity premia of 1.03% when the agent's Γ equals 100, and a value of 1.69% when Γ equals 0. Generally, the premia for the labor income cases, when expressed as fractions of the premia for the analogous multiplicative wealth shock cases, are much smaller for the predictability cases than the i.i.d case: 0.3 (fixed $\phi, \rho_{g,d} = -0.5$), 1.8 (fixed $\phi, \rho_{g,d} = 0$), 0.6 (fixed $\phi, \rho_{g,d} = 0.5$), 0.25 ($\rho_{g,d} = -0.5, \rho_{\phi,d} = 1, \rho_{\phi,d} = -0.5$) and 0.5 ($\rho_{g,d} = 0.5, \rho_{\phi,d} = -1, \rho_{\phi,d} = -0.5$) versus 3.3 (i.i.d. case).

These results give rise to two questions. First, why are the $\Gamma = 0$ liquidity premia (measured relative to multiplicative shock premia) smaller for the predictable return cases than the i.i.d. case? And second, why does the $\Gamma = 0$ liquidity premia go down as $\rho_{g,d}$ goes from either 0 to 0.5 or 0 to -0.5 contrary to intuition? Examining the effect of the no short-selling restrictions on the agent's behavior can answer both these questions. With predictable returns, being at either the 100%-stock boundary or the 0%-stock boundary inhibits the agent's ability to risky-asset time, so trading is reduced, transactions costs are less important, and liquidity premia are reduced. With i.i.d. returns, no risky-asset timing occurs, so being at the 100%-stock boundary has less of an effect on trading and thus on liquidity premia. When $\rho_{g,d} \geq 0$, the agent with a low Γ early in life likes stocks, and except for very low states, holds the maximum possible 100% in stock. When $\rho_{g,d} = -0.5$, even the agent with a low Γ early in life holds zero in stocks in all but the highest dividend states because of the negative hedging demand with respect to future labor income. This explains why the $\Gamma = 0$ liquidity premia (measured relative to multiplicative shock premia) are smaller for the predictable return cases than the i.i.d. case.

When $\rho_{g,d} \geq 0$, the agent with a low Γ early in life likes stocks more and is at the 100% stocks boundary in even more states when $\rho_{g,d}$ is 0.5 rather than 0. This explains why the $\Gamma = 0$ liquidity premia goes down as $\rho_{g,d}$ goes from 0 to 0.5. When $\rho_{g,d} = -0.5$, the agent with a low Γ early in life holds zero in stocks in a sufficiently large number of low states that the $\Gamma = 0$ liquidity premia goes down as $\rho_{g,d}$ goes from 0 to -0.5.

5.2 *Two risky assets*

One concern is that the large liquidity premia and utility costs that we document are not robust to the presence of a second low-transaction-cost asset that is positively correlated with the high-transaction-cost asset. To address this concern, we allow the investor access to a high liquidity

portfolio in addition to the low liquidity one. As described above, the joint return process for the two assets is again calibrated using U.S stock portfolios formed by sorting on liquidity (see Acharya and Pedersen, 2002). Throughout, we set the transaction cost rate on the second asset to zero and keep the average rate on the low-liquidity portfolio at 2%.

The liquidity premia results are in Table 7. As expected, the presence of the second asset ameliorates the liquidity premium on the low-liquidity asset, but for many of the scenarios considered, the reduction is quite modest. For example, in the case in which the multiplicative wealth shock covaries negatively with expected return and the transaction cost rate, and these latter two covary together positively, the liquidity premium on the low-liquidity asset is still 4.13% per annum in the presence of a second risky asset. This number is similar to the 4.55% per annum obtained when only the low-liquidity asset is available. Thus, the large liquidity premia that we report appear to be robust to the introduction of a second risky asset whose correlation with the low-liquidity portfolio mirrors that of the high-liquidity portfolio in the data.

Table 8 is a companion table to Table 6, reporting utility costs rather than liquidity premia. Again, we find the pattern of these utility costs largely mirrors the pattern for the liquidity premia. However, with respect to the utility cost magnitudes, the effect of adding a second risky asset attenuates utility costs to a much larger extent than it does the liquidity premium.

We also allow the second risky asset to be the market portfolio instead of the low liquidity portfolio. The results are contained in the second-last column of Panel B in Table 9 for a selected subset of the cases. Again, the reduction in liquidity premia is quite modest.

5.3 Low liquidity portfolio: 3% versus 1% transaction cost rate comparison

Since we are concerned with the cross-section, it is useful to generalize the concept of liquidity premium for an asset having a given cost rate rather than a lower cost rate, where the lower rate may be larger than zero. The premium is defined to be the required reduction in the mean return for an otherwise identical lower-cost asset so that an agent with access to the given-cost asset would prefer not to switch to the lower-cost asset if it was offered instead. As we explained above, we use a cost rate of 2% because it represents the difference in transactions cost rates across large and small firms and then we calculate the premium for the low liquidity portfolio comparing this 2% rate to a rate of 0%. However, this premium may not be the same as the one obtained when the low liquidity portfolio with a cost rate of 3% is compared to the low liquidity portfolio with a cost rate of 1%.

The third column of Panel B in Table 9 reports liquidity premia for the 3% versus 1% comparison for a subset of the cases, while the first column of Panel B reproduces liquidity for premia for the initial 2% versus 0% comparison for the same subset. We find that the premia in all cases are lower for the 3% to 1% comparison, but the ratios of the premia relative to the premium in the standard case are similar. We are still able to generate premia as high as 1.97% with a constant cost rate and as high as 2.56% when we allow the rate to be stochastic.

5.4 *Turnover and direct trading costs*

Turnover results are reported in Panel A of Table 9 for four problems, one column for each problem: low liquidity 2%; low liquidity 3%; high liquidity plus low liquidity 2%; market plus low liquidity 2%. Focusing on the problem in which the agent has access to low liquidity portfolio at a 2% transaction cost rate, which is reported in the first column, we find that annual turnover goes from about 4% in the standard case to 163% when the correlation between expected return and the multiplicative wealth shock is negative and the cost rate is constant. Once the cost rate is allowed to be stochastic, the annual turnover can be as high as 145%, which occurs when the multiplicative wealth shock covaries negatively with expected return and the transaction cost rate, and these latter two covary together positively. Recall that these two cases with negative correlation between expected return and the wealth shock are also cases with large liquidity premia. When the correlations are set to the data point estimates, the annual turnover becomes 73%. Examining annual turnover when the transaction cost rate is increased to 3%, we see that turnover decreases in all cases but is still higher than 100% when the multiplicative wealth shock covaries negatively with expected return and the transaction cost rate, and these latter two covary together positively. It is still more than 60% when the correlations are set to the data point estimates.

The direct effect of transaction costs on expected return is reported in Panel B of Table 9 for four comparisons, two columns for each comparison: low liquidity 2% versus low liquidity 0%; low liquidity 3% versus low liquidity 1%; high liquidity plus low liquidity 2% versus high liquidity plus low liquidity 0%; market plus low liquidity 2% versus market plus low liquidity 0%. For each comparison, the first column reports liquidity premia while the second reports the direct effect of transaction costs on expected return. Not surprisingly, for all four comparisons, the direct cost is close to the liquidity premium when there are no wealth shocks and the cost rate is constant. The direct cost as a fraction of the liquidity premium is highest when the expected return and the multiplicative wealth shock are negatively correlated. For the first comparison (low liquidity 2%

versus 0%), it can take a value as high as 0.97 and this occurs when the cost rate is a constant. The fraction is smallest when the expected return and the multiplicative wealth shock are positively correlated. For the first comparison, it can take a value as low as 0.22 and this occurs when the transaction cost rate covaries negatively with expected return and the multiplicative wealth shock, and these latter two covary together positively.

Recall that the extent to which the liquidity premium exceeds this direct cost can be attributed to a risk premium for trading more when the agent is poor in utility terms, and in the case of a stochastic cost rate, to the agent trading more when the cost rate is high. The effect of the wealth shock-expected return correlation on this ratio is not surprising since the sign of this correlation determines the sign of the hedging demand with respect to future labor income. When the correlation is negative, the zero risky-asset weight constraint binds in low dividend yield states, which means the agent tends to trade when she is wealthy in utility terms because of high future expected returns and high future labor income growth rates. The result is a high ratio. The converse is true when the correlation is positive.

5.5 *Comparative statics*

We are interested in assessing how the liquidity premia vary with changes in key parameters of the problem. Table 10 reports liquidity premia for the same set of cases as Table 4 but for an investor with risk aversion of 8 rather than 6. The table shows that increasing the investor's risk aversion from 6 to 8 increases the liquidity premia on the low-liquidity portfolio in all cases. The liquidity premium gets as high as 5.66% in the case in which the multiplicative wealth shock covaries negatively with expected return and the transaction cost rate, and these latter two covary together positively.

We are also interested in how the liquidity premium is affected by changes in the volatility of the multiplicative wealth shock. This comparative static is of particular interest since it is difficult to calibrate this volatility to data. We focus on the case with the correlations set equal to the data point estimates: $\rho_{l,d} = -0.21$, $\rho_{\phi,d} = 0.18$ and $\rho_{\phi,l} = -0.11$. Figure 2 plots the liquidity premium as a function of the shock volatility for this data correlations case. Plots for risk aversion of 6 and 8 are contained in the graph. We see that the liquidity premia is an increasing convex function of the wealth shock volatility. As the wealth shock volatility goes to zero, the liquidity premium converges to that for the case without wealth shocks and the specified correlation between the transaction cost rate and expected return.

6 Stock Allocations

Figure 3 plots average stock allocation as a function of age for a variety of wealth to permanent income ratios ranging from 0 up to 1000 for the case in which the correlations are set to the data point estimates. When the wealth-income ratio is 1000, which means the agent is financially wealthy relative to her income stream, the average allocation is a concave, decreasing function of age, as would be expected given that the allocation is close to that of an agent who does not receive any labor income. However, once the wealth income ratio is less than 100, the relation becomes positive with a hump near the end of the agent's life. If the wealth-income ratio is less than 10, then the agent's average allocation in the first month is always at least 30% lower than the 70% chosen by the agent with a wealth-income ratio of 1000. The lower average allocations lead to higher liquidity premia holding the quantity of trading fixed. Lynch and Tan (2004) explore in detail the effect of business cycle variation in the first two moments of permanent labor income growth on allocation.

7 Conclusions

The seminal work of Constantinides (1986) documents how, when the risky return is calibrated to the U.S. market return, the impact of transaction costs on per-annum liquidity premia is an order of magnitude smaller than the cost rate itself. A number of recent papers have formed portfolios sorted on liquidity measures and found a spread in expected per-annum return that is definitely not an order of magnitude smaller than the transaction cost spread: the expected per-annum return spread is found to be around 6-7% per annum. Our paper bridges the gap between Constantinides' theoretical result and the empirical magnitude of the liquidity premium by examining dynamic portfolio choice with transaction costs in a variety of more elaborate settings that move the problem closer to the one solved by real-world investors. In particular, we allow returns to be predictable and transaction costs to be stochastic, and we introduce wealth shocks, both stationary multiplicative and labor income. With predictable returns, we also allow the wealth shocks and transaction costs to be state dependent.

With labor income calibrated to U.S. PSID data as the wealth shock and its correlation with dividend yield assumed to be negative, the liquidity premium can be as high as 1.14% for an agent with no financial wealth and as high as 0.80% for an agent whose wealth to permanent monthly income ratio is 100. When the multiplicative wealth shock covaries negatively with expected return and the transaction cost rate, and these latter two covary together positively, the liquidity premium

is higher relative to the standard i.i.d. return case by a factor of 57: a liquidity premium of 0.079% per annum in the standard i.i.d. return case becomes 4.55%, which is a number that is in the same ballpark as the numbers reported in empirical work. Interestingly, when we estimate these correlations using U.S. data, we find exactly this pattern of correlations between the three variables. The liquidity premium becomes 1.56% when the correlations are set to the data point estimates, which are admittedly very noisy estimates of the true correlations. We conclude that the effect of proportional transaction costs on the standard consumption and portfolio allocation problem with i.i.d. returns can be materially altered by reasonable perturbations that bring the problem closer to the one investors are actually solving.

Clearly, our paper is but a first step toward bridging the gap between the theoretical literature and the empirical work finding large spreads in expected returns for portfolios formed on the basis of liquidity. One important limitation of our analysis is that it is partial equilibrium. Hence, it says nothing about how transaction costs affect equilibrium prices by limiting the ability of agents to risk share. More work is needed to understand how transaction costs effect prices and returns in a general equilibrium setting.

Appendix A: Solution Technique for the Labor Income Problem

This appendix sketches the numerical procedures associated with computing liquidity premia in a dynamic savings and portfolio choice problem with predictable or i.i.d. returns and labor income.

There are three key elements to the implementation. The first is to endogenize the discrete state representation of the value function to bound error propagation at each iteration. The second is to resort to extrapolation only when the problem on hand is economically sufficiently close (in a sense to be made clear below) to a problem for which the functional form for the value function is known. The third is to exploit a natural sense in which the algorithm can be parallelized across computational units to obtain linear reductions in run-time.

The concern that gives rise to the above elements is that the wealth to lagged permanent labor income (wealth-income, henceforth) ratio state is unbounded on the non-negative side of the real line. To represent the value function on this dimension, this range is partitioned into three disjoint, non-degenerate intervals. A different algorithm is applied to obtain an approximation on each interval. At each iteration, the lower intermediate boundary point is chosen to be the smallest value of wealth-income ratio such that an agent without the labor income for all periods to terminal date requires no more than 10% extra wealth to be equally happy as an otherwise identical agent with the labor income at that wealth-income ratio. The upper intermediate boundary point is chosen to satisfy the same definition at 1%. Over the lower-end interval, the value function is approximated as a piecewise linear form or piecewise shape-preserving monotone cubic hermite interpolant of Fritsch and Carlson (1980) for the wealth-income ratio state and as a piecewise linear form for the inherited allocation state. Over the higher-end interval, the value function is taken to be that of the otherwise identical problem without the labor income. Over the middle interval, the value function is approximated as a function of the form, $V(W) = a \times (W - b)^c + d$, where a, b, c, d are constants in \mathbb{R} , to match the function and the first derivative values at the upper and the lower intermediate boundary points. Further, at any given iteration and at any given grid node of the discretized state space, the objective function has to be jointly solved for consumption and portfolio policies subject to the short sales constraints on the T-bill and the stock. A recursive golden section algorithm is used to optimize the consumption policy, defined as the fraction of last period's permanent labor income consumed, accurate to the fourth decimal digit and the portfolio policy, defined as the fraction of wealth invested in the risky asset, accurate to the third decimal digit.

We suggest a dynamic gridding algorithm to bound errors on policy functions at each iteration.

This algorithm takes the value function representation for the previous iteration as given, computes the intermediate boundary points for the iteration at hand and continues to add points in the lower-end interval in a particular way until policies at no point on a representative grid for the next iteration differ by more than prespecified magnitudes across increasingly denser grids for the iteration on hand.

The steps of this algorithm in detail are:

- 1) take value function representation for the previous iteration, $V_{t+1}(\cdot)$, as given
- 2) initialize a set of grid nodes, $X_t = \{\Gamma_{t,1}, \Gamma_{t,2}, \Gamma_{t,3}, \dots, \Gamma_{t,n}\}$, in increasing order, where $\Gamma_{t,1} = 0$, $\Gamma_{t,n} = LP_{t,10\%}$, n is a constant, and $LP_{t,10\%}$ is the lower intermediate boundary point for iteration t with the defining parameter set to 10%
- 3) define $Y_t = \{(\Gamma_{t,1} + \Gamma_{t,2})/2, (\Gamma_{t,2} + \Gamma_{t,3})/2, \dots, (\Gamma_{t,n-1} + \Gamma_{t,n})/2\}$, in increasing order
- 4) evaluate $V_t(X_t)$ using the value function representation, $V_{t+1}(\cdot)$
- 5) evaluate $V_t(Y_t)$ using the value function representation, $V_{t+1}(\cdot)$
- 6) evaluate the lower intermediate point, $LP_{t-1,10\%}$ for iteration $t - 1$ using V_t as represented by $X_t \cup Y_t$
- 7) define $A_{t-1} = \{(i - 1) \times (LP_{t-1,10\%})/q\}_{i=1}^{q+1}$, where $q \geq 2$ is a constant.
- 8) evaluate $V_{t-1}(A_{t-1})$ using V_t as represented by X_t
- 9) evaluate $V_{t-1}(A_{t-1})$ using V_t as represented by $X_t \cup Y_t$
- 10) define $Z_{t-1} \subseteq A_{t-1} = \{x \in A_{t-1} \mid \text{and policy functions at } x \text{ using representations in (8) and (9) differ more than prespecified magnitudes}\}$
- 11) if $Z_{t-1} = \emptyset$, accept $X_t \cup Y_t$ as sufficient representation for V_t , exit subroutine.
- 12) compute the range of wealth-income, Γ_t , range the system can possibly assume by starting at any point in Z under any realization of the return and the labor income shocks and under any allowed policy. This range is generically of the form $[0 \ h_t]$, for some $h_t > 0$.
- 13) update X_t as $X_t \cup Y_t$, in increasing order,
- 14) update Y_t as $\{(X_{t,1} + X_{t,2})/2, (X_{t,2} + X_{t,3})/2, \dots, (X_{t,k-1} + X_{t,k})/2\}$; where k is the smallest positive integer satisfying $X_{t,k} \geq \min(h_t, LP_t)$.
- 15) return to step 4 and repeat.

Step (10) ensures that the maximal absolute scaled deviation in consumption policy (defined as the fraction of last period's permanent component of labor income consumed), $|(\hat{\kappa}_{coarse} - \hat{\kappa}_{dense})/\hat{\kappa}_{dense}|$ is bounded from above by 10^{-3} and the maximal absolute deviation in portfolio policy (defined as the fraction of wealth invested in the risky asset), $|\alpha_{coarse} - \alpha_{dense}|$, is bounded

from above by 10^{-2} . q and n are set to 50. The initialization in step 2 is designed to take advantage of the previous iteration's (V_{t+2}) representation.

This algorithm, given the choice of the interpolant produces a set of grid nodes which then, can represent the value function at hand sufficiently well before one goes on to the next iteration. For the i.i.d. case without the temporary shocks to labor income under the parametrization given in table 6, choice of a piecewise linear interpolant leads, on average (over the life cycle), to approximately 400 grid nodes, and the choice of the piecewise shape-preserving monotone cubic hermite interpolant leads, on average (over the life cycle), to approximately 130 grid nodes. For the predictability cases, we always use the shape-preserving monotone cubic hermite interpolant. For the case with $\rho_{g,d} = 0$ and other parameters as given in table 6, this algorithm leads, on average (over the life cycle), to approximately 180 grid nodes.

For any distinct parametric specification and return generating process, this dynamic gridding scheme is run without the transactions costs and each node on the resultant grid for the wealth-income ratio state is augmented with the 51-node uniform inherited allocation grid on $[0, 1]$ to get the state discretization for the transactions cost problem. This procedure implies that the joint state for the case with predictable returns and transactions costs, with $\rho(g, d) = 0$, is on average (over the life cycle), represented with $19 \times 180 \times 51 \cong 170,000$ grid nodes.

To determine the reduction in unconditional mean returns required to offset elimination of the transactions costs and keep the agent's expected utility the same, a standard bisection algorithm is used. This algorithm is set to produce liquidity premia accurate to the fourth decimal digit.

While reducing the unconditional mean return in the no transactions cost problem to calculate liquidity premia, the gridding scheme obtained for the otherwise similar case with the returns calibrated to the data is used. For the i.i.d. case without the temporary shocks in table 6 and for the no financial wealth case, we rerun the gridding scheme under the return processes with the unconditional mean reduced exactly in the magnitude of the liquidity premium for this case, and recompute liquidity premium. The liquidity premium is unchanged to reported precision.

To ensure that results are robust to the parameters of the solution algorithm, for the i.i.d. case without the temporary shocks in table 6, we change the defining parameter for the lower boundary point from 10% to 9%, 8%, 7%, and 6%. The maximal $|(\hat{\kappa}_{10\%} - \hat{\kappa}_{x\%})/\hat{\kappa}_{x\%}|$, for an equally spaced 100-node grid on $[0, LP_{t,10\%}]$ is less than 10^{-3} for all $t = 1, \dots, 240$ and $x = 6, 7, 8, 9$. Similarly, the maximal $|\alpha_{10\%} - \alpha_{x\%}|$ is less than 10^{-2} . Fixing the defining parameter for the lower boundary point at 10%, we change the precision on the consumption policy, $\hat{\kappa}$, from 10^{-4} to 10^{-5} and 10^{-6} .

We again note that the maximal absolute scaled consumption policies is less than than 10^{-3} across the life cycle.

Finally, to keep run-time manageable, state space at each iteration is decomposed into disjoint subsets. Each subset is optimized by a particular peripheral computational unit, the value function values are aggregated by a central aggregator unit which in turn passes the interpolant representations back to the peripheral units. This parallelization scheme leads to linear reductions in run-time.

References

- Acharya, V.V., Pedersen, L. H., 2002. Asset pricing with liquidity risk. Unpublished working paper. New York University, NY.
- Aiyagari, S.R., Gertler, M., 1991. Asset returns with transactions costs and uninsured individual risk: a stage III exercise. *Journal of Monetary Economics* 27, 309-31.
- Amihud, Y., 2002. Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets* 5, 31-56.
- Amihud, Y., Mendelson, H., 1986. Asset pricing and the bid-ask spread. *Journal of Financial Economics* 17, 223-49.
- Balduzzi, P., Lynch, A., 1999. Transaction costs and predictability: some utility cost calculations. *Journal of Financial Economics* 52, 47-78.
- Barberis, N., 2000. Investing for the long run when returns are predictable. *Journal of Finance* 55, 225-64.
- Brennan, M., Chordia, T. and Subrahmanyam, A., 1998. Alternative factor specifications, security characteristics and the cross-section of expected returns, *Journal of Financial Economics* 49, 345-374.
- Brennan, M., Subrahmanyam, A., 1996. Market microstructure and asset pricing: on the compensation for illiquidity in stock returns. *Journal of Financial Economics* 41, 441-64.
- Campbell, J., Viceira, L., 1999. Consumption and portfolio decisions when expected returns are time varying. *Quarterly Journal of Economics* 114, 433-495.
- Carroll, C., 1996. Buffer-stock saving: Some theory, Working paper, John Hopkins University.
- Carroll, C., 1997. Buffer-stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics* 112, 1-55.
- Carroll, C., Sanwick, A., 1997. The nature of precautionary wealth. *Journal of Monetary Economics* 40, 41-71.

- Cocco, J., Gomes, F., Maenhout, P., 2002. Consumption and portfolio choice over the lifecycle. Working paper, London Business School.
- Chamberlain, G., Hirano, K., 1997. Predictive distributions based on longitudinal earnings data. Working paper, Harvard University.
- Constantinides, G., 1986. Capital market equilibrium with transaction costs. *Journal of Political Economy* 94, 842-62.
- Easley, D., Hvidkjaer, S., O'Hara, M., 2002. Is information risk a determinant of asset prices?. Forthcoming *Journal of Finance*.
- Fama, E., French, K., 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3-56.
- Fritsch, F. N., Carlson, R. E., 1980. Monotone Piecewise Cubic Interpolation. *SIAM Journal of Numerical Analysis* 17, 238-246
- Gakidis, H., 1997. Earnings uncertainty and life-cycle portfolio choice. Working paper, MIT.
- Gomes, F., Michaelides, A., 2002. Portfolio choice with habit formation: A life-cycle model with uninsurable labor income risk. Working paper, London School of Economics.
- Glosten, L., Harris, L., 1988. Estimating the components of the bid-ask spread. *Journal of Financial Economics* 21, 123-42.
- Grossman, S., Shiller, R., 1981. The determinants of the variability of stock market prices. *American Economic Review* 71, 222-227.
- Hansen, L., Singleton, K., 1982. Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica* 50, 1269-1289.
- Hasbrouck, J., 2003. Trading costs and returns for U.S. equities: evidence from daily data. Unpublished working paper. New York University, NY.
- Ho, T. and Stoll, H., 1981, Optimal Dealer Pricing Under Transactions and Return Uncertainty, *Journal of Financial Economics* 9, 47-73.

- Heaton, J., Lucas, D., 1996. Evaluating the effects of incomplete markets on risk sharing and asset pricing. *Journal of Political Economy* 104, 443-487.
- Huang, M., 2002. Liquidity shocks and equilibrium liquidity premia. *Journal of Economic Theory*, forthcoming.
- Kyle, A., 1985. Continuous auctions and insider trading. *Econometrica* 53, 1315-35.
- Lesmond, D., Ogden, J., Trzcinka, T., 1999. A new measure of total transaction costs. *Review of Financial Studies* 12, 1113-41.
- Lettau, M., Ludvigson, S., 2001. Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying. *Journal of Political Economy*, 109(6), 1238-1287.
- Lo, A., Mamaysky, H., Wang, J., 2001. Asset prices and trading volume under fixed transactions costs. Unpublished working paper. MIT and Yale University.
- Lynch A. W., 2000. Portfolio choice and equity characteristics: characterizing the hedging demands induced by return predictability. *Journal of Financial Economics* 62, 67-130.
- Lynch A.W., Tan,S., 2003. Multiple risky assets, transaction costs and return predictability: implications for portfolio choice. Unpublished working paper. New York University, NY.
- Lynch A.W., Tan,S., 2004. Labor Income Dynamics at Business-cycle Frequencies: Implications for Portfolio Choice. Unpublished working paper. New York University, NY.
- Mehra, R., Prescott, E., 1985. The equity premium: a puzzle. *Journal of Monetary Economics* 15, 145-161.
- Pástor, L., Stambaugh, R., 2003. Liquidity risk and expected stock returns. *Journal of Political Economy* 111, 642-685.
- Stoll, H., 1978, The Supply of Dealer Services in Securities Markets, *Journal of Finance* 33, 1133-1151.
- Tauchen, G., Hussey, R., 1991. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica* 59, 317-96.

Vayanos, D., 1998. Transaction costs and asset prices: a dynamic equilibrium model. *Review of Financial Studies* 11, 1-58.

Viceira, L., 1997. Optimal portfolio choice for long-horizon investors with non-tradable labor income. Working paper, Harvard University.

Viceira, L., 2001. Optimal portfolio choice for long-horizon investors with non-tradable labor income. *Journal of Finance* 56, 433-470.

Table 1. SAMPLE STATISTICS, VAR COEFFICIENTS AND QUADRATURE APPROXIMATION: LOW AND HIGH LIQUIDITY PORTFOLIOS. The table reports moments and parameters for the high and low liquidity portfolios estimated from U.S. data and calculated for the quadrature approximation based on the VAR that uses log dividend yield as the only state variable. The VAR is described in section 3.1. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The high liquidity portfolio is the value-weighted portfolio of the most liquid 12 portfolios and the low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. The data period is from February 1964 to December 1996. Panel A reports unconditional means, VAR slopes and VAR R^2 's for the data and the quadrature approximation. b is the vector of VAR slopes and R^2 denotes the regression R^2 . Panel B reports the unconditional covariance matrix for the data and for the quadrature approximation. Panels C reports the conditional covariance matrices for the data VAR and the quadrature VAR. All results are for continuously compounded returns. Returns are expressed per month and in percent.

Data				Quadrature			
<i>Panel A: Unconditional sample moments and VAR coefficients</i>							
Asset/Variable	Uncond. Mean	b	R^2	Uncond. Mean	b	R^2	
High Liquidity	0.45	0.37	0.73	0.45	0.37	0.73	
Low Liquidity	0.84	0.58	1.19	0.84	0.58	1.19	
Dividend Yield	0.00	0.98	95.67	0.00	0.96	92.58	
Asset/Variable	High Liquidity	Low Liquidity	Dividend Yield	High Liquidity	Low Liquidity	Dividend Yield	
<i>Panel B: Unconditional standard deviations, covariances (above diagonal), and correlations (below)</i>							
High Liquidity	4.34	20.42	-0.49	4.34	20.42	-0.75	
Low Liquidity	0.88	5.33	-0.36	0.88	5.33	-0.65	
Dividend Yield	-0.11	-0.07	1.00	-0.17	-0.12	1.00	
<i>Panel C: Conditional standard deviations, covariances (above diagonal), and correlations (below) for the VAR with Dividend Yield as Predictor</i>							
High Liquidity	4.32	20.20	-0.85	4.32	20.20	-1.11	
Low Liquidity	0.88	5.30	-0.92	0.88	5.30	-1.21	
Dividend Yield	-0.94	-0.84	0.21	-0.94	-0.84	0.27	

Table 2. SAMPLE STATISTICS, VAR COEFFICIENTS AND QUADRATURE APPROXIMATION: LOW LIQUIDITY AND MARKET PORTFOLIOS. The table reports moments and parameters for the low liquidity and market portfolios estimated from U.S. data and calculated for the quadrature approximation based on the VAR that uses log dividend yield as the only state variable. The VAR is described in section 3.1. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The Low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. Market portfolio is the CRSP value weighted return on the NYSE index. The data period is from February 1964 to December 1996. Panel A reports unconditional means, VAR slopes and VAR R^2 s for the data and the quadrature approximation: b is the vector of VAR slopes and R^2 denotes the regression R^2 . Panel B reports the unconditional covariance matrix for the data and for the quadrature approximation. Panels C reports the conditional covariance matrices for the data VAR and the quadrature VAR. All results are for continuously compounded returns. Returns are expressed per month and in percent.

Data				Quadrature			
<i>Panel A: Unconditional sample moments and VAR coefficients</i>							
Asset/Variable	Uncond. Mean	b	R^2	Uncond. Mean	b	R^2	
Low Liquidity	0.84	0.58	1.19	0.84	0.58	1.19	
Market	0.47	0.38	0.76	0.47	0.38	0.76	
Dividend Yield	0.00	0.98	95.67	0.00	0.96	92.58	
Asset/Variable	Low Liquidity	Market	Dividend Yield	Low Liquidity	Market	Dividend Yield	
<i>Panel B: Unconditional standard deviations, covariances (above diagonal), and correlations (below)</i>							
Low Liquidity	5.33	20.56	-0.36	5.33	20.56	-0.65	
Market	0.89	4.31	-0.48	0.89	4.31	-0.74	
Dividend Yield	-0.07	-0.11	1.00	-0.12	-0.17	1.00	
<i>Panel C: Conditional standard deviations, covariances (above diagonal), and correlations (below) for the VAR with Dividend Yield as Predictor</i>							
Low Liquidity	5.30	20.34	-0.92	5.30	20.34	-1.21	
Market Liquidity	0.89	4.30	-0.85	0.89	4.30	-1.11	
Dividend Yield	-0.84	-0.95	0.21	-0.84	-0.95	0.27	

Table 3. EMPIRICAL ESTIMATES OF MEANS, STANDARD DEVIATIONS, AND CORRELATIONS OF WEALTH SHOCKS, TRANSACTIONS COSTS AND DIVIDEND YIELD. The table presents empirical means, standard deviations and correlations of wealth shocks, proportional transactions costs and dividend yield, where the wealth shock is taken to be the change in the permanent income. Hasbrouck (2003), employs a bayesian approach to estimating half-spreads using daily data from all ordinary common equity issues on the CRSP daily database from 1962 to 2002 and averages the daily numbers to obtain half-spread estimates at an annual frequency for each stock. Fixing year, we sort the stocks on half-spread and equally weight the top and the bottom 50 % to obtain annual time series estimates for two hypothetical portfolios, *liquid portfolio* and *illiquid portfolio*, respectively. If half spread is greater than 30% for an observation, we set that observation to 30% analogous to a treatment in Acharya and Pedersen (2002). We construct a third time series, *Illiquid - Liquid*, by subtracting the liquid portfolio half spread from the illiquid portfolio half spread. In Panel A, we report the unconditional means and the standard deviations (in percent) of these series. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. The high liquidity portfolio is the value-weighted portfolio of the most liquid 12 portfolios. Acharya & Pedersen (2002) data set also provides the corresponding normalized ILLIQ which can be interpreted as a round trip transactions cost rate at the same monthly frequency. We halve this normalized variable and value weight to obtain one way proportional transactions cost rates for the low and the high liquidity portfolios. We also construct a third time series, *Low - High*, by subtracting the transactions cost rate of the high liquidity portfolio from that of the low liquidity portfolio. The data period is from February 1964 to December 1996. We correlate these time series, with the dividend yield series, d , (12 month dividend yield on the value weighted NYSE index) on monthly data and with the change in permanent labor income, g , where we take the y of Lettau and Ludvigson (2001), as the permanent income, on quarterly data. We average monthly transactions cost rates for three months to get quarterly rates. Φ denotes the proportional transactions cost percentage, ρ denotes unconditional correlation, μ denotes unconditional mean and σ denotes unconditional standard deviation. Further, we have that $\phi = \log(1 + \Phi)$.

<i>Panel A: Means and Standard Deviations</i>		
<u>Φ Series</u>	<u>μ_Φ</u>	<u>σ_Φ</u>
Liquid	0.39	0.40
Illiquid	5.84	7.24
Illiquid-Liquid	5.44	7.08

<i>Panel B: Correlations</i>		
$\rho_{g,d}$	-0.21	
<u>Φ Series</u>	<u>$\rho_{\phi,d}$</u>	<u>$\rho_{\phi,g}$</u>
High Liquidity	0.13	0.06
Low Liquidity	0.11	-0.17
Low-High	0.11	-0.18

Table 4. LIQUIDITY PREMIA ON THE LOW LIQUIDITY PORTFOLIO: NO OTHER RISKY ASSET IS AVAILABLE, RISK AVERSION IS 6 AND THE MEAN TRANSACTIONS COST PERCENTAGE IS 2. The table reports annual liquidity premia (in percent) on the low liquidity portfolio as defined in Constantinides (1986) when the only risky asset available to the investor is the low liquidity portfolio. The liquidity premium is the increase in the low liquidity portfolio's annual expected log return going from the no-transactions-cost-case to the otherwise identical transactions-cost-case that makes the investor indifferent between the two cases. Panel A tabulates constant transactions costs cases while Panel B tabulates random transactions costs cases and the subpanels vary the wealth shock dynamics. Returns can be predictable or i.i.d. and in either case are calibrated to U.S. data using a quadrature approximation. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. The data period is from February 1964 to December 1996. Φ denotes the proportional transactions cost percentage on the low liquidity portfolio, L denotes the monthly percentage wealth shock, d denotes log dividend yield, ρ denotes unconditional correlation, μ denotes unconditional mean and σ denotes unconditional standard deviation. Further, we have that $\phi = \log(1 + \Phi)$ and $l = \log(1 + L)$. The risk aversion parameter, γ , is set to 6.

<i>Panel A:</i>		$\Phi = 2\%$	
<u>$L = 0$</u>			
		<i>i.i.d.</i>	<i>Predictable</i>
		0.08	0.46
<u>$\mu_L = 0 \quad \sigma_l = 0.05$</u>			
$\rho_{l,d}$		<i>i.i.d.</i>	<i>Predictable</i>
-0.5			3.37
0		0.42	0.90
0.5			2.51
<i>Panel B:</i>		$\mu_\Phi = 2\% \quad \sigma_\Phi = 0.76\%$	
<u>$L = 0$</u>			
$\rho_{\phi,d}$		<i>i.i.d.</i>	<i>Predictable</i>
-1			0.51
-0.5			0.44
0		0.05	0.39
0.5			0.36
1			0.39
<u>$\mu_L = 0 \quad \sigma_l = 0.05$</u>			
		<u>$\rho_{\phi,l} = -0.5$</u>	
$\rho_{l,d}$	$\rho_{\phi,d}$	<i>i.i.d.</i>	<i>Predictable</i>
	-1		—
	-0.5		2.64
-0.5	0		3.17
	0.5		3.79
	1		4.55
	-1		—
	-0.5		0.93
0	0	0.34	0.93
	0.5		0.96
	1		—
	-1		—
	-0.5		3.55
0.5	0		3.00
	0.5		2.53
	1		2.10
	-1		—
	-0.5		3.00
	0		2.53
	0.5		2.10
	1		—
		<u>$\rho_{\phi,l} = 0$</u>	
		<i>i.i.d.</i>	<i>Predictable</i>
			0.85
			0.81
		0.28	0.81
			0.84
			0.91
		<u>$\rho_{\phi,l} = 0.5$</u>	
		<i>i.i.d.</i>	<i>Predictable</i>
			1.82
			2.17
			2.75
			3.40
			—
			—
			0.73
		0.23	0.73
			0.75
			—
			—
			2.55
			2.09
			1.65
			1.29

Table 5. UTILITY COSTS OF TRANSACTIONS COSTS ON THE LOW LIQUIDITY PORTFOLIO WHEN NO OTHER RISKY ASSET IS AVAILABLE, RISK AVERSION IS 6 AND THE MEAN TRANSACTIONS COST PERCENTAGE IS 2. When returns are i.i.d., the table reports the percentage of wealth a consumer would be prepared to give up to get access to the analogous no-transactions-costs-case, when returns are predictable it reports the means of such percentage. Panel A tabulates constant transactions costs cases while Panel B tabulates random transactions costs cases and subpanels vary the wealth shock dynamics. Returns can be predictable or i.i.d. and in either case are calibrated to the U.S. data by a quadrature approximation. The only risky asset available to the agent is a low liquidity portfolio. In each case the dynamic consumer program is solved and the initial value of the inherited portfolio allocation is set equal to beginning-of-horizon portfolio allocation for the analogous dynamic program with no transactions costs. Means are calculated using the unconditional distribution for the dividend yield. All values are reported in percent. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. The data period is from February 1964 to December 1996. Φ denotes the proportional transactions cost percentage on the low liquidity portfolio, L denotes the monthly percentage wealth shock, d denotes log dividend yield, ρ denotes unconditional correlation, μ denotes unconditional mean and σ denotes unconditional standard deviation. Further, we have that $\phi = \log(1 + \Phi)$ and $l = \log(1 + L)$. The risk aversion parameter, γ , is set to 6.

<i>Panel A:</i>		$\Phi = 2\%$					
		$L = 0$					
		<i>i.i.d.</i>	<i>Predictable</i>				
		0.35	2.82				
		$\mu_L = 0 \quad \sigma_l = 0.05$					
$\rho_{l,d}$		<i>i.i.d.</i>	<i>Predictable</i>				
-0.5			23.87				
0		2.38	7.08				
0.5			26.39				
<i>Panel B:</i>		$\mu_\Phi = 2\% \quad \sigma_\Phi = 0.76\%$					
		$L = 0$					
		<i>i.i.d.</i>	<i>Predictable</i>				
	$\rho_{\phi,d}$						
	-1		3.11				
	-0.5		2.67				
	0	0.22	2.37				
	0.5		2.22				
	1		2.36				
		$\mu_L = 0 \quad \sigma_l = 0.05$					
		$\rho_{\phi,l} = -0.5$					
		$\rho_{\phi,l} = 0$					
		$\rho_{\phi,l} = 0.5$					
$\rho_{l,d}$	$\rho_{\phi,d}$	<i>i.i.d.</i>	<i>Predictable</i>	<i>i.i.d.</i>	<i>Predictable</i>	<i>i.i.d.</i>	<i>Predictable</i>
	-1		—		—		14.57
	-0.5		19.83		18.44		16.94
-0.5	0		22.82		21.67		20.45
	0.5		25.98		25.02		24.01
	1		29.31		—		—
	-1		—		6.69		—
	-0.5		7.34		6.44		5.79
0	0	1.94	7.31	1.57	6.43	1.31	5.78
	0.5		7.51		6.62		5.98
	1		—		7.15		—
	-1		34.59		—		—
	-0.5		30.40		28.57		26.71
0.5	0		26.56		24.60		22.65
	0.5		22.81		20.76		18.51
	1		—		—		14.89

Table 6. LIQUIDITY PREMIA ON THE LOW LIQUIDITY PORTFOLIO IN THE PRESENCE OF LABOR INCOME: NO OTHER RISKY ASSET IS AVAILABLE, RISK AVERSION IS 6 AND THE MEAN TRANSACTIONS COST PERCENTAGE IS 2. The table reports annual liquidity premia (in percent) on the low liquidity portfolio as defined in Constantinides (1986) when the only risky asset available to the investor is the low liquidity portfolio. The liquidity premium is the increase in the low liquidity portfolio's annual expected log return going from the no-transactions-cost-case to the otherwise identical transactions-cost-case that makes the investor indifferent between the two cases. Panel A tabulates constant transactions costs cases while panel B tabulates random transactions costs cases. Returns can be predictable or i.i.d. and in either case are calibrated to U.S. data using a quadrature approximation. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. The data period is from February 1964 to December 1996. Φ denotes the proportional transactions cost percentage on the low liquidity portfolio, ϵ denotes the log monthly temporary shock to labor income, g denotes monthly change in log permanent labor income, d denotes log dividend yield, ρ denotes unconditional correlation, μ denotes unconditional mean and σ denotes unconditional standard deviation. Further, we have that $\phi = \log(1 + \Phi)$. \bar{g} and σ_g are chosen to match 3% and 15% respectively, for the mean annual growth in permanent labor income and standard deviation of the change in annual permanent labor income. When labor income has a temporary component, μ_ϵ and σ_ϵ are chosen to match 1.00 and 10% respectively, for the mean annual multiplicative temporary shock and standard deviation of the log annual temporary shock. Data sets $\rho_{g,d} = -0.21, \rho_{\phi,d} = 0.11, \rho_{\phi,g} = -0.18$. The risk aversion parameter, γ , is set to 6.

<i>Panel A : $\Phi = 2\%$</i>						
$\epsilon \neq 0$	$\epsilon = 0$		$\epsilon = 0$			
	<i>i.i.d.</i>		<i>Predictable</i>			
Γ			$\rho_{g,d} = -0.5$	$\rho_{g,d} = -0.21$	$\rho_{g,d} = 0$	$\rho_{g,d} = 0.5$
0	1.42	1.42	0.84	1.02	1.68	1.43
1	1.38	1.38	0.83	0.99	1.63	1.40
10	0.94	0.94	0.79	0.89	1.20	1.22
100	0.27	0.27	0.67	0.62	0.61	0.87
1000	0.11	0.11	0.56	0.50	0.53	0.68
∞	0.08	0.08	0.46	0.46	0.46	0.46
<i>Panel B : $\mu_\Phi = 2\%, \sigma_\Phi = 0.76\%$</i>						
$\epsilon = 0$						
<i>Predictable</i>						
Γ	$\rho_{g,d} = -0.5, \rho_{\phi,d} = 1, \rho_{\phi,g} = -0.5$		<i>Data</i>	$\rho_{g,d} = 0.5, \rho_{\phi,d} = -1, \rho_{\phi,g} = -0.5$		
0	1.14		1.21	1.69		
1	1.13		1.12	1.66		
10	1.05		0.99	1.42		
100	0.80		0.64	1.03		
1000	0.57		0.51	0.80		
∞	0.39		0.37	0.51		

Table 7. LIQUIDITY PREMIA ON THE LOW LIQUIDITY PORTFOLIO WHEN THE HIGH LIQUIDITY PORTFOLIO IS ALSO AVAILABLE, RISK AVERSION IS 6 AND MEAN TRANSACTIONS COST PERCENTAGE IS 2. The table reports annual liquidity premia (in percent) on the low liquidity portfolio as defined in Constantinides (1986) when a portfolio of high liquidity stocks (high liquidity portfolio) is also available. The liquidity premium is the increase in the low liquidity portfolio's annual expected log return going from the no-transactions-cost-case to the otherwise identical transactions-cost-case that makes the investor indifferent between the two cases. Panel A tabulates constant transactions costs cases while Panel B tabulates random transactions costs cases and subpanels vary the wealth shock dynamics. Returns can be predictable or i.i.d. and in either case are calibrated to U.S. data using a quadrature approximation. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. The high liquidity portfolio is the value-weighted portfolio of the most liquid 13 portfolios. The data period is from February 1964 to December 1996. Φ denotes the proportional transactions cost percentage on the low liquidity portfolio, L denotes the monthly percentage wealth shock, d denotes log dividend yield, ρ denotes unconditional correlation, μ denotes unconditional mean and σ denotes unconditional standard deviation. Further, we have that $\phi = \log(1 + \Phi)$ and $l = \log(1 + L)$. The risk aversion parameter, γ , is set to 6.

<i>Panel A:</i>		$\Phi = 2\%$	
<u>$L = 0$</u>			
		<i>i.i.d.</i>	<i>Predictable</i>
		0.08	0.43
<u>$\mu_L = 0 \sigma_l = 0.05$</u>			
$\rho_{l,d}$		<i>i.i.d.</i>	<i>Predictable</i>
-0.5			3.19
0		0.41	0.76
0.5			2.02
<i>Panel B:</i>		$\mu_\Phi = 2\% \sigma_\Phi = 0.76\%$	
<u>$L = 0$</u>			
		<i>i.i.d.</i>	<i>Predictable</i>
$\rho_{\phi,d}$			0.46
-1			0.38
-0.5		0.05	0.34
0			0.33
0.5			0.37
1			
<u>$\mu_L = 0 \sigma_l = 0.05$</u>			
		<u>$\rho_{\phi,l} = -0.5$</u>	
		<u>$\rho_{\phi,l} = 0$</u>	
		<u>$\rho_{\phi,l} = 0.5$</u>	
$\rho_{l,d}$	$\rho_{\phi,d}$	<i>i.i.d.</i>	<i>Predictable</i>
	-1		—
	-0.5		1.81
	0		2.12
-0.5	0		2.34
	0.5		2.81
	1		2.63
			3.19
			—
			—
	-1		0.71
	-0.5		0.67
0	0	0.32	0.77
	0.5		0.67
	1		0.69
			0.77
			—
	-1		—
	-0.5		3.32
	0		2.12
0.5	0		1.95
	0.5		1.37
	1		1.22
			1.53
			1.19
			1.03
			0.88
			—
			0.86

Table 8. UTILITY COSTS OF TRANSACTIONS COSTS ON THE LOW LIQUIDITY PORTFOLIO WHEN THE HIGH LIQUIDITY PORTFOLIO IS ALSO AVAILABLE, RISK AVERSION IS 6 AND MEAN TRANSACTIONS COST PERCENTAGE IS 2. When returns are i.i.d., the table reports the percentage of wealth a consumer would be prepared to give up to get access to the analogous no-transactions-costs-case, when returns are predictable it reports the means of such percentage. Panel A tabulates constant transactions costs cases while Panel B tabulates random transactions costs cases and subpanels vary the wealth shock dynamics. Returns are calibrated to the U.S. data using a quadrature approximation. The risky assets available to the agent are a low liquidity portfolio and a high liquidity portfolio. In each case the dynamic consumer program is solved and the initial value of the inherited portfolio allocation is set equal to beginning-of-horizon portfolio allocation for the analogous dynamic program with no transactions costs. Means are calculated using the unconditional distribution for the dividend yield. All values are reported in percent. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. The high liquidity portfolio is the value-weighted portfolio of the most liquid 13 portfolios. The data period is from February 1964 to December 1996. Φ denotes the proportional transactions cost percentage, L denotes the monthly percentage wealth shock, d denotes log dividend yield, ρ denotes unconditional correlation, μ denotes unconditional mean and σ denotes unconditional standard deviation. Further, we have that $\phi = \log(1 + \Phi)$ and $l = \log(1 + L)$. The risk aversion parameter, γ , is set to 6.

<i>Panel A:</i>		$\Phi = 2\%$	
		$L = 0$	
		<i>i.i.d.</i>	<i>Predictable</i>
		0.35	2.60
		$\mu_L = 0 \sigma_l = 0.05$	
$\rho_{l,d}$		<i>i.i.d.</i>	<i>Predictable</i>
-0.5			22.94
0		2.36	6.04
0.5			11.01
<i>Panel B:</i>		$\mu_\Phi = 2\% \sigma_\Phi = 0.76\%$	
		$L = 0$	
		<i>i.i.d.</i>	<i>Predictable</i>
$\rho_{\phi,d}$			2.79
-1			2.34
-0.5		0.21	2.07
0			1.99
0.5			2.26
1			
		$\mu_L = 0 \sigma_l = 0.05$	
		$\rho_{\phi,l} = -0.5$	
$\rho_{l,d}$	$\rho_{\phi,d}$	<i>i.i.d.</i>	<i>Predictable</i>
	-1		—
	-0.5		19.26
-0.5	0		21.87
	0.5		24.63
	1		27.53
	-1		—
	-0.5		6.12
0	0	1.84	6.07
	0.5		6.30
	1		—
	-1		14.20
	-0.5		11.36
0.5	0		9.19
	0.5		7.66
	1		—
		$\rho_{\phi,l} = 0$	
$\rho_{l,d}$	$\rho_{\phi,d}$	<i>i.i.d.</i>	<i>Predictable</i>
	-1		—
	-0.5		17.98
-0.5	0		20.85
	0.5		23.79
	1		—
	-1		5.70
	-0.5		5.36
0	0	1.40	5.32
	0.5		5.54
	1		6.12
	-1		—
	-0.5		10.78
0.5	0		8.48
	0.5		6.86
	1		—
		$\rho_{\phi,l} = 0.5$	
$\rho_{l,d}$	$\rho_{\phi,d}$	<i>i.i.d.</i>	<i>Predictable</i>
	-1		14.53
	-0.5		16.56
-0.5	0		19.76
	0.5		22.92
	1		—
	-1		—
	-0.5		4.73
0	0	1.20	4.70
	0.5		4.91
	1		—
	-1		—
	-0.5		10.24
0.5	0		7.78
	0.5		6.03
	1		5.93

Table 9. TURNOVER, LIQUIDITY PREMIA AND DIRECT TRADING COSTS ON THE LOW LIQUIDITY PORTFOLIO. The table presents per annum turnover in Panel A and per annum direct trading costs (columns labelled *Direct Cost*) (in percent) next to the per annum liquidity premia (in percent) as defined in Constantinides (1986) (columns labelled *Liquidity Premia*) on the low liquidity portfolio in Panel B. Row titles indicated the return and the wealth shock and transactions cost dynamics. *Data* sets $\rho_{\phi,d} = 0.11, \rho_{\phi,l} = -0.18, \rho_{l,d} = -0.21$ as in Table 3 calibration. The per annum turnover is defined to be $12/240 \times \sum_{t=1}^{240} \text{av}(|\alpha_t - \hat{\alpha}_t|) / \text{av}(\max(\alpha_t, \hat{\alpha}_t) / (1 + L_t))$, where $\text{av}(\cdot)$ is taking the average across the 100,000 simulation paths under the optimal policies. Direct trading cost is defined to be $\mu_{\Phi} \times \text{turnover}$. The liquidity premia as defined in Constantinides (1986) is the reduction in the unconditional mean of the otherwise identical lower-cost-benchmark case to equate utility across the two cases. The associated μ_{Φ} 's for the two cases are given as column titles. The stochastic transactions costs cases always set μ_{Φ}/σ_{Φ} to 2.63. The market and the high liquidity portfolios trade costlessly when available. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. The high liquidity portfolio is the value-weighted portfolio of the most liquid 12 portfolios. Market portfolio is the CRSP value weighted return on the NYSE index. Φ denotes the proportional transactions cost percentage on the low liquidity portfolio, L denotes the monthly percentage wealth shock, d denotes log dividend yield, ρ denotes unconditional correlation, μ denotes unconditional mean and σ denotes unconditional standard deviation. Further, we have that $\phi = \log(1 + \Phi)$ and $l = \log(1 + L)$. We have $\mu_L = 0$ and $\sigma_L = 0.05$. The risk aversion parameter, γ , is set to 6.

Return Process	Transactions Costs and Wealth Shock		Alone		Plus High Liquidity		Plus Market	
			2% vs. 0%	3% vs. 1%	2% vs. 0%	2% vs. 0%	Liquidity Premia	Direct Cost
<i>i. i. d.</i>	constant Φ		3.89	3.42	3.90	3.89		
<i>i. i. d.</i>	constant $\Phi, \rho_{l,d} = 0$		10.44	9.38	9.07	9.09		
<i>Predictable</i>	constant Φ		20.54	15.21	19.49	19.33		
<i>Predictable</i>	constant $\Phi, \rho_{l,d} = -0.5$		163.66	147.17	156.04	151.96		
<i>Predictable</i>	constant $\Phi, \rho_{l,d} = 0$		28.72	22.70	22.70	22.37		
<i>Predictable</i>	constant $\Phi, \rho_{l,d} = 0.5$		41.83	38.56	27.87	28.16		
<i>Predictable</i>	$\rho_{\phi,d} = 1, \rho_{\phi,l} = -0.5, \rho_{l,d} = -0.5$		144.94	123.17	133.15	131.74		
<i>Predictable</i>	$\rho_{\phi,d} = -1, \rho_{\phi,l} = -0.5, \rho_{l,d} = 0.5$		40.17	37.59	34.84	35.78		
<i>Predictable</i>	<i>Data</i>		73.50	62.23	70.12	69.23		

Panel A: Turnover

	Liquidity Premia and Direct Trading Costs		Liquidity Premia		Direct Cost		Liquidity Premia		Direct Cost	
	Liquidity Premia	Direct Cost	Liquidity Premia	Direct Cost	Liquidity Premia	Direct Cost	Liquidity Premia	Direct Cost	Liquidity Premia	Direct Cost
<i>i. i. d.</i>	0.08	0.08	0.05	0.05	0.08	0.08	0.08	0.08	0.08	0.08
<i>i. i. d.</i>	0.42	0.21	0.40	0.16	0.41	0.18	0.42	0.18	0.42	0.18
<i>Predictable</i>	0.46	0.41	0.26	0.14	0.43	0.39	0.42	0.39	0.42	0.39
<i>Predictable</i>	3.37	3.27	1.97	1.83	3.19	3.12	3.16	3.03	3.16	3.03
<i>Predictable</i>	0.90	0.57	0.62	0.27	0.76	0.45	0.73	0.45	0.73	0.45
<i>Predictable</i>	2.51	0.84	1.52	0.67	2.02	0.56	2.12	0.56	2.12	0.56
<i>Predictable</i>	4.55	2.90	2.56	1.76	4.13	2.66	4.06	2.63	4.06	2.63
<i>Predictable</i>	3.55	0.80	2.25	0.68	3.32	0.70	3.46	0.72	3.46	0.72
<i>Predictable</i>	1.56	1.47	1.00	0.93	1.49	1.40	1.48	1.38	1.48	1.38

Panel B: Liquidity Premia and Direct Trading Costs

Table 10. LIQUIDITY PREMIA ON THE LOW LIQUIDITY PORTFOLIO: NO OTHER RISKY ASSET IS AVAILABLE, RISK AVERSION IS 8 AND THE MEAN TRANSACTIONS COST PERCENTAGE IS 2. The table reports annual liquidity premia (in percent) on the low liquidity portfolio as defined in Constantinides (1986) when the only risky asset available to the investor is the low liquidity portfolio. The liquidity premium is the increase in the low liquidity portfolio's annual expected log return going from the no-transactions-cost-case to the otherwise identical transactions-cost-case that makes the investor indifferent between the two cases. Panel A tabulates constant transactions costs cases while Panel B tabulates random transactions costs cases and the subpanels vary the wealth shock dynamics. Returns can be predictable or i.i.d. and in either case are calibrated to U.S. data using a quadrature approximation. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. The data period is from February 1964 to December 1996. Φ denotes the proportional transactions cost percentage on the low liquidity portfolio, L denotes the monthly percentage wealth shock, d denotes log dividend yield, ρ denotes unconditional correlation, μ denotes unconditional mean and σ denotes unconditional standard deviation. Further, we have that $\phi = \log(1 + \Phi)$ and $l = \log(1 + L)$. The risk aversion parameter, γ , is set to 8.

<i>Panel A:</i>		$\Phi = 2\%$	
<u>$L = 0$</u>			
		<i>i.i.d.</i>	<i>Predictable</i>
		0.11	0.50
<u>$\mu_L = 0 \quad \sigma_l = 0.05$</u>			
$\rho_{l,d}$		<i>i.i.d.</i>	<i>Predictable</i>
-0.5			4.23
0		0.52	0.99
0.5			3.00
<i>Panel B:</i>		$\mu_\Phi = 2\% \quad \sigma_\Phi = 0.76\%$	
<u>$L = 0$</u>			
		<i>i.i.d.</i>	<i>Predictable</i>
$\rho_{\phi,d}$			0.52
-1			0.45
-0.5			0.42
0		0.07	0.41
0.5			0.46
1			
<u>$\mu_L = 0 \quad \sigma_l = 0.05$</u>			
		<u>$\rho_{\phi,l} = -0.5$</u>	
$\rho_{l,d}$	$\rho_{\phi,d}$	<i>i.i.d.</i>	<i>Predictable</i>
	-1		—
	-0.5		3.35
-0.5	0		4.03
	0.5		4.79
	1		5.66
		<u>$\rho_{\phi,l} = 0$</u>	
	-1		—
	-0.5		0.91
0	0	0.42	0.97
	0.5		1.03
	1		—
		<u>$\rho_{\phi,l} = 0.5$</u>	
	-1		—
	-0.5		3.04
0.5	0	0.26	3.02
	0.5		2.50
	1		—
		<u>$\rho_{\phi,l} = 0$</u>	
	-1		—
	-0.5		3.31
0.5	0		2.74
	0.5		2.22
	1		—
		<u>$\rho_{\phi,l} = 0.5$</u>	
	-1		—
	-0.5		3.04
0.5	0		2.47
	0.5		1.92
	1		1.44




Figure 1. LIQUIDITY PREMIA AND CORRELATION OF WEALTH SHOCKS WITH DIVIDEND YIELD WITH CONSTANT 2% TRANSACTIONS COSTS. The graph reports annual liquidity premia in percent on the low liquidity portfolio when this portfolio is the only risky asset available to the investor against the unconditional correlation of the monthly log wealth shock with the dividend yield, $\rho_{l,d}$, when $\Phi = 2\%$. The liquidity premium is the increase in the low liquidity portfolio's expected return going from the no-transactions-cost-case to the current case that makes the investor indifferent between the two cases as defined in Constantinides (1986). Returns are predictable and are calibrated to data by a quadrature approximation. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. Data period is from February 1964 to December 1996. Let Φ denote the proportional transactions cost rate, L denote the monthly percentage wealth shock, d denote the log dividend yield, ρ denote unconditional correlation, μ denote unconditional mean and σ denote unconditional standard deviation. Further, define $\phi = \log(1+\Phi)$ and $l = \log(1+L)$. We set $\mu_L = 0$ and $\sigma_l = 0.05$. The Risk aversion parameter, γ , is 6.




Figure 2. LIQUIDITY PREMIA AND VOLATILITY OF WEALTH SHOCKS WITH DATA CALIBRATION. The graph reports annual liquidity premia in percent on the low liquidity portfolio when this portfolio is the only risky asset available to the investor against the unconditional volatility of the monthly log wealth shock. Calibrated values for the transactions costs and wealth shock dynamics in table 3 are used by setting $\rho_{l,d} = -0.21$, $\rho_{\phi,d} = 0.11$, $\rho_{\phi,l} = -0.18$. The liquidity premium is the increase in the low liquidity portfolio's expected return going from the no-transactions-cost-case to the current case that makes the investor indifferent between the two cases as defined in Constantinides (1986). Returns are predictable and are calibrated to data by a quadrature approximation. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. Data period is from February 1964 to December 1996. Let Φ denote the proportional transactions cost rate, L denote the monthly percentage wealth shock, d denote the log dividend yield, ρ denote unconditional correlation, μ denote unconditional mean and σ denote unconditional standard deviation. Further, define $\phi = \log(1+\Phi)$ and $l = \log(1+L)$. All graphs set $\mu_{\phi} = 2\%$, $\sigma_{\phi} = 0.76\%$. The Risk aversion parameter, γ , is either 6 or 8.




Figure 3. Portfolio allocations with labor income, data calibration, no transactions costs, averaged across dividend yield states. Graph plots average allocations in the presence of labor income versus time (time 1 (240) being the start (end) of the investment horizon), where the joint dynamics of labor income, transactions costs and dividend yield are calibrated to the data as described in table 3. We set $\rho_{g,d} = -0.21$, $\rho_{\phi,d} = 0.11$, $\rho_{\phi,g} = -0.18$. Returns are predictable and are calibrated to the data by a quadrature method. Average allocations for the wealth to lagged permanent labor income ratios of 0, 1, 10, 100 and 1000, across the nineteen dividend yield states using the unconditional probability distribution for the dividend yield are plotted. Decision interval is monthly. The only risky asset available is the low liquidity portfolio. g denotes monthly change in log permanent labor income, d denotes log dividend yield, ρ denotes unconditional correlation. The marginal distribution of g is calibrated to match PSID data as described in section 3.2. The Acharya & Pedersen (2002) data set provides 25 value-weighted portfolios sorted on ILLIQ, a liquidity measure suggested by Amihud (2002). The low liquidity portfolio is the value-weighted portfolio of the least liquid 13 portfolios. The data period is from February 1964 to December 1996. The risk aversion parameter, γ , is set to 6.