

DISCUSSION PAPER SERIES

No. 3795

PORTFOLIO CHOICE WITH ILLIQUID ASSETS

Miklós Koren and Adam Szeidl

TRANSITION ECONOMICS



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP3795.asp

PORTFOLIO CHOICE WITH ILLIQUID ASSETS

Miklós Koren, Harvard University
Adam Szeidl, Harvard University

Discussion Paper No. 3795
February 2003

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **TRANSITION ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Miklós Koren and Adam Szeidl

CEPR Discussion Paper No. 3795

February 2003

ABSTRACT

Portfolio Choice with Illiquid Assets*

The present Paper investigates the effects of incorporating illiquidity in a standard dynamic portfolio choice problem. Lack of liquidity means that an asset cannot be immediately traded at any point in time. We find the portfolio share of financial wealth invested in illiquid assets given the liquidity premium. Benchmark calibrations imply a portfolio share of 2-6% in cash. These numbers are in line with survey data and also with portfolio recommendations by practitioners. We also find that long horizon investors invest more in illiquid assets. Overall, our results suggest that differences between asset classes unrelated to standard price risk may influence portfolio shares.

JEL Classification: G11 and G12

Keywords: asset pricing, calibration, liquidity and portfolio choice

Miklós Koren
Department of Economics
Harvard University
324 Littauer Center
Cambridge MA 02138
USA
Tel: (1 617) 495 1045
Fax: (1 617) 495 7730
Email: koren@fas.harvard.edu

Adam Szeidl
Department of Economics
Harvard University
324 Littauer Center
Cambridge MA 02138
USA
Tel: (1 617) 495 1045
Fax: (1 617) 495 7730
Email: szeidl@fas.harvard.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=139516

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=153397

*We are grateful to Péter Benczúr, Laurent Calvet, John Campbell, David Laibson, Gérard Roland, Mark Schaffer, Jeremy Tobacman, János Vincze and participants of the CEPR Transition Economics Summer Workshop for Young Academics (Portorož) and the Central European University Workshop on Liquidity of Financial Markets (Budapest) for helpful comments and suggestions. Any remaining errors are ours. This research is supported by a grant from CERGE-EI Foundation under a program of the Global Development Network. The views expressed herein are ours and have not been endorsed by CERGE-EI or the GDN.

Submitted 05 November 2002

Introduction

The purpose of this paper is to examine portfolio investment decisions when available asset markets differ in their degrees of liquidity. It is a common observation that liquidity properties vary across markets. In our view liquidity can be an important determinant of portfolio investment decisions, complementary in nature to risk and return. To verify this intuition, we build a simple model to investigate the effect of asset liquidity on optimal portfolio shares of infinitely lived investors, and attempt to calibrate it.

Throughout we assume that liquidity refers to the ease with which an asset can be sold. Thus an asset or asset market is liquid if trade can take place at short notice in large quantities without a substantial price change. Out of these three aspects of liquidity, that is at short notice, in large quantities and without a substantial discount, we will focus on the first one. The source of illiquidity of an asset in our model is the difficulty of buying or selling the asset immediately when the necessity of trade is realized. Trade of an illiquid asset takes place with a lag; that is, some time elapses after the order to trade has been submitted to the market and before trade actually takes place. The waiting time till trade occurs once an order is submitted serves as a measure of the illiquidity of a particular asset market (see Lippman and McCall, 1986). In equilibrium, one expects that a less liquid asset will command a higher yield premium. Nowhere do we assume that the assets traded in our model are necessarily financial assets. This is partly because many examples of illiquid assets are non-financial, and partly to allow for greater flexibility and more generality.

This paper maintains the view that the utility possibilities set is not time invariant; rather during certain periods consumption is more rewarding than at other times. One reason might be that consumption is not infinitely divisible, but takes place in large pieces. In that case, the availability of a large piece of good at a reasonable price might correspond to more rewarding consumption. Alternatively, the utility stemming from a given level of consumption may vary over the life cycle (marriage, children, vacation). Thirdly, it may be the case that from time to time agents have a favorable private investment opportunity. To capture such effects, agents in this paper are sometimes faced with liquidity

shocks (Diamond and Dybvig, 1983). The liquidity shock is assumed to affect marginal utility, with the consumer having a higher marginal utility for the duration of the shock.

Given the presence of liquidity shocks and the difficulty of selling illiquid assets immediately, agents will face a non-trivial trade-off in their portfolio allocation decisions. Holding a less liquid asset will yield a higher premium; on the other hand, the agent will become more vulnerable to liquidity shocks. Indeed, when a shock comes, the agent with a higher share of illiquid assets in her portfolio will find it more difficult to liquidate sufficient wealth for current consumption. The analysis of this trade-off is the main purpose of the present paper.

We set up an infinite horizon model with power utility investors to examine portfolio investment decisions under illiquidity. As discussed above, the two main features of this model are the taste shocks and the lag in trade of the illiquid asset. We characterize the optimal consumption and portfolio decision in the model, and derive an almost completely analytic solution. The existence of such a semi-analytic solution is a big advantage of the model. It makes the analysis easy to perform and allows us to experiment with several calibration exercises. We rely on simple numerical techniques to get calibration results as well as comparative statics.

Calibration results imply that liquidity influences the portfolio decisions of an investor with borrowing constraints. Because many of the underlying parameters in our model are not directly observable, we experiment with a variety of specifications. Using a range of benchmark parameter values we find that the portfolio share of the liquid asset varies between 2% and 6%. This result is in line with empirical studies. Bertaut and Starr-McCluer (2001) find liquid wealth to be 3.5% of total assets in aggregate data. They also use survey data from the Survey of Consumer Finances, to find that average liquid wealth to total assets of U.S. households varies between 4.6% and 5.7% over time. Both estimates are in the range of our calibration results. Additionally, Elton and Gruber (2000) point out that the asset allocation recommendations of major investment banks contain 5% cash for liquidity reasons. This is again in line with our calibration

results. Angeletos, Laibson, Repetto, Tobacman and Weinberg (2001) report average liquid assets to total assets to vary between 8% and 16% depending on the exact definition of liquid assets, using data from the 1995 Survey of Consumer Finances. These estimates are somewhat larger than our results.

On the more qualitative side, we find a horizon effect on the optimal portfolio. Long horizon investors hold most of their wealth in illiquid assets, while short horizon investors prefer cash. This effect stems from the different optimal lifetime consumption patterns of short and long horizon investors. Short horizon investors consume more out of their current wealth, therefore need to hold a higher share of liquid assets in their portfolio. In contrast to the literature on optimal portfolios with risky assets (see Campbell and Viceira, 2002), the horizon effects in our model are present in a completely stationary environment, with no time variation in returns.

A somewhat surprising finding is that more risk averse investors hold less cash. To understand why, first note that in our model the coefficient of relative risk aversion is also the reciprocal of the elasticity of intertemporal substitution. A consumer with high risk aversion is less willing to accept a steady decline or growth in consumption over time, so her consumption policy will be such that her wealth does not change dramatically. In our parametrization this will imply that a more risk-averse investor consumes less out of her wealth, hence a lower buffer stock of liquid wealth is sufficient. There is an opposite effect more related to risk aversion itself. As the liquidity shock hits, the consumer will not be able to choose her first-best level of consumption because of the borrowing constraint. Such deviations from the unconstrained first best consumption are less tolerated by a more risk-averse investor, so she will tilt her portfolio towards the liquid asset to avoid them. Nevertheless, the first effect dominates in our benchmark parametrization.

Also interesting is how the portfolio share of liquid assets changes in the liquidity premium. The demand for liquid asset turns out to be inelastic in the premium. What is surprising, however, is that a higher liquidity premium implies a higher share of cash in the optimal portfolio. The intuitive reason is that a higher liquidity premium in effect increases the portfolio return of

the consumer. This influences the optimal consumption plan via a substitution effect and an income effect. The income effect dominates, therefore optimal consumption increases each period. But in order to finance that, the agent needs to hold more liquid assets in her portfolio.

We interpret our findings as evidence that non-risk related asset characteristics have some influence on investment decisions. It is particularly insightful to compare our calibration results to Merton (1971) and Merton (1973), who finds implausibly high portfolio weight on the risky asset in a calibration exercise. In contrast our study yields plausible results making use of asset characteristics not directly related to price risk.

The main driving force of our model is the recurring needs for liquidity. There are several ways to formalize these needs. In our setup liquidity shocks directly influence the consumer's marginal utility. This formulation is closest to Diamond and Dybvig (1983); in that paper, as here, the liquidity shock shifts the consumer's intertemporal marginal rate of substitution. Note that in their paper the shock is a lot more drastic. Also related is Baldwin and Meyer (1979), where investment opportunities with different rates of return arrive at stochastic intervals, and Huang (2001), where agents are forced to leave the economy after an exponentially distributed lifetime. At a formal level our framework is very similar to multiplicative models of habit formation (Abel, 1990), but in those models habit is predetermined, whereas here the liquidity shock is an innovation.

A less standard assumption of the present paper is that trade of the illiquid asset takes place with a lag. We tend to think about this assumption as a parable that represents one aspect of illiquidity. One expects that no asset is impossible to sell immediately if the price discount is large enough. However, if the discount is too high, it may be optimal to wait, as in Lippman and McCall (1986). The reason why we focus on this particular aspect of illiquidity is threefold. First, this formulation yields a tractable model and we are able to find a semi-analytic solution. Second, lags in trade are an empirically relevant aspect of several asset markets. Third, to our knowledge, this formulation of illiquidity has not been examined in the literature. A more standard way of modelling illiquidity is by means of transactions costs. This line of research focuses on

the price discount flavor of illiquidity. Ultimately, focusing on just one aspect of liquidity is not satisfactory. More specifically, our liquidity shocks would interact differently with transaction costs than they do with trade lags. Trade lags force the consumer to keep some of her wealth in liquid assets, thereby influencing her optimal portfolio. This is closer to a buffer-stock interpretation of liquid wealth. Transactions costs, on the other hand, would make it costly for her to liquidate her portfolio at any point in time, independently of the liquidity shock (see Huang, 2001 for similar liquidity shocks and transactions costs).

There are many examples of illiquid financial markets. These include small stocks on developed country exchanges, stock markets in developing countries, stocks of not publicly traded companies; and also publicly traded long bonds and emerging country sovereign bonds. Examples of illiquid non-financial assets are privately owned companies and real estate. The large variety of illiquid markets suggests that liquidity should matter in portfolio decisions and asset pricing.

Numerous authors investigated empirically whether the liquidity risk of a security is systematically related to its expected return. These studies include Amihud and Mendelson (1986), Brennan and Subrahmanyam (1996), and Fiori (2000). The studies tend to find, as expected, that less liquid stocks command a higher return. For assets less liquid than stocks, a study by Benczúr (2002) attempts to identify empirically part of the premium on emerging country government bonds that is caused by illiquidity.

On the theoretical side, many papers investigate the relationship between liquidity and asset pricing. A close relative to the present paper is Huang (2001), who also studies the effects of liquidity shocks in the presence of illiquid assets. Illiquidity in that paper is introduced by means of transaction costs. Another feature common with our model is that the volume of trade in the illiquid asset can be large. This feature is also present in Lo, Mamaysky and Wang (2001). Neither of these papers focus on portfolio choice. Other studies include Constantinides (1986), Grossman and Laroque (1990), Dumas and Luciano (1991), Heaton and Lucas (1996), Vayanos (1998). In all of these papers illiquidity is captured by means of transactions costs. Additionally, in contrast with the present study, the volume of trade in these latter papers is small.

There is also a considerable amount of literature on market microstructure, that studies the possible determinants of liquidity; Campbell, Lo and MacKinley (1997) is a good overview.

There has been a recent revival in the literature on portfolio choice, concerning long-horizon investment decisions (see Canner, Mankiw and Weil, 1997, Campbell and Viceira, 2002). That line of research investigates long horizon portfolio choice where the risk-return characteristics of assets are time-varying. The present study is complementary, in that it looks at long horizon portfolio choice, but in the presence of illiquid assets and time-varying utility possibilities set.

The rest of this paper is organized as follows. Section 1 presents our model. We set up the Bellman equation and characterize the optimal policy. We set up a problem that can be solved analytically, and characterize the solution of the original problem as a solution to this modified problem. This gives us the semi-analytic solution. In Section 2 we turn to calibration. Because the solution is not completely analytic, we need to rely on some simple numerical techniques. The calibration results are encouraging. Finally, Section 3 concludes.

1. The Model

In this section we analyze the portfolio decision problem of a consumer in the presence of liquidity shocks, when available assets differ in their degree of liquidity. We spell out the portfolio optimization problem and derive the Bellman equation. Due to the nature of illiquidity in our model, the value function admits a relatively simple functional form. This enables us to find a semi-analytic solution for the value function as well as the optimal consumption and portfolio policy. We turn to calibrating the model in the next section.

We study the portfolio decision problem of an infinitely lived consumer with exogenous initial wealth and no labor income. Time is discrete. As is customary, we assume that the per period utility function exhibits constant relative risk aversion. This ensures that portfolio shares will be independent of wealth. The consumer maximizes the lifetime discounted value of her utility over consump-

tion. Her (per period) marginal utility is subject to a multiplicative taste shock, the liquidity shock. The reasonableness of this assumption has been discussed in the introduction.

The consumer has access to two assets, a liquid and an illiquid one. Lack of liquidity is modelled as follows. The market for the illiquid asset operates with a lag. The consumer is free to place any order to buy or sell at the prevailing price of 1 at any point in time. However, an order placed at the beginning of period t will only be executed at the end of that period, after consumption has already taken place. As we have argued earlier, this formulation captures the time to sell aspect of liquidity. In what follows, we shall call the illiquid asset “bond,” and the liquid asset “cash.” This is not to say that we exclusively focus on financial assets, rather, for purposes of exposition only.

We assume that short sale and borrowing are prohibited, and that consumption can only be financed out of the liquid asset. This implies a liquidity constraint on consumption: at any point in time, consumption cannot be larger than total cash holdings plus the return earned on asset holdings.

The liquid asset pays a cash interest rate of r per unit on each date. The illiquid asset pays a dividend of $r + \pi$ in cash, π being the liquidity premium. Notice that there is no uncertainty about the price or the cash-flow of the assets; for the purposes of this study we have shut down issues related to price risk. The only reason why the illiquid asset is not a risk free bond is that there is a lag in the execution of orders: the asset is not perfectly liquid.

The consumer solves the following program

$$\max_{\{C_t, \alpha_{t+1}\}} E_0 \sum_{t=0}^{\infty} \beta^t \chi_t^\theta \frac{C_t^{1-\theta}}{1-\theta}, \quad (1)$$

where C_t is consumption at time t , α_{t+1} is the cash to wealth ratio once the order placed at time t realizes, $\beta < 1$ is the subjective discount factor, χ_t is the taste shock at time t , and θ is the coefficient of relative risk aversion (we assume $\theta > 1$). Maximization is subject to the budget constraint and the liquidity

constraints:

$$\begin{aligned} W_{t+1} &= R_{pt}W_t - C_t \\ C_t &\leq (\alpha_t + r_{pt})W_t \\ R_{pt} &= 1 + r_{pt} = 1 + r + (1 - \alpha_t)\pi. \end{aligned}$$

Here W_t and α_t stand for total wealth (liquid and illiquid) respectively the share of liquid wealth in the portfolio, both measured at the beginning of period t . We write R_{pt} for the gross, and r_{pt} for the net portfolio return from period t to period $t + 1$. Although both R_{pt} and r_{pt} depend on the portfolio share α_t , we suppress this dependence in notation, hoping that it will not cause confusion.

The taste shock χ_t is independently and identically distributed over time, and can take two values only:

$$\chi_t = \begin{cases} \gamma > 1 & \text{with probability } \mu, \\ 1 & \text{with probability } 1 - \mu. \end{cases}$$

We say that the consumer experiences a liquidity shock in period t if $\chi_t = \gamma$. The assumption that $\gamma > 1$ implies that the marginal utility of consumption is higher during a liquidity shock than at normal times. Therefore the consumer wishes to consume more during a liquidity shock.

Because the taste shock is independently distributed over time, there are only two state variables, wealth (W_t) and the share of liquid assets in the portfolio (α_t). Consumers are allowed to choose different consumption levels depending on their current taste shock. This gives three control variables to be chosen by the consumer: consumption levels in either of the two states and next period's portfolio share.

The timing of asset trade, consumption and interest earned is as follows. At the end of period $t - 1$, the order placed in that period is executed. Interest earned on the new portfolio between periods t and $t + 1$ is paid out in advance, at the beginning of period t . Then the taste shock is realized. Given the taste shock, the consumer chooses consumption subject to her current liquidity constraint, and also the portfolio share for next period. At the end of period t , her order is executed on the market. This timing is illustrated in Figure 1.




Figure 1: Timing of events

The value function and optimal program of the consumer can be characterized by the following Bellman equation:

$$\begin{aligned}
 V(W, \alpha) = \max_{\alpha'} \left\{ \mu \max_{C_1 \leq (\alpha + r_p)W} \left[\gamma^\theta \frac{C_1^{1-\theta}}{1-\theta} + \beta V(R_p W - C_1, \alpha') \right] + \right. \\
 \left. (1 - \mu) \max_{C_2 \leq (\alpha + r_p)W} \left[\frac{C_2^{1-\theta}}{1-\theta} + \beta V(R_p W - C_2, \alpha') \right] \right\}. \quad (2)
 \end{aligned}$$

Here C_1 stands for consumption if a liquidity shock occurs, whereas C_2 is consumption in the normal state. α' denotes the fraction of wealth in cash after this period's order is executed.

The intuition behind the Bellman equation is the following. A liquidity shock takes place with probability μ . When it occurs, current consumption will deliver higher marginal utility, hence the multiplier γ^θ on the per period utility. Consumption cannot be more than current liquid wealth, hence the constraint in the maximization. Next period wealth is given by wealth plus interest earned this period, minus consumption this period (recall that R_p is gross interest). The second maximization, describing the case when there is no liquidity shock, can be interpreted similarly. Since the taste shock is independent over time, in either of the two states the consumer chooses the same optimal portfolio share for next period, namely α' .

Because both the budget constraint and the liquidity constraint are linear in consumption and wealth and utility is power utility, the value function will be homogeneous of degree $1 - \theta$ in wealth.

Proposition 1. *The value function is homogeneous of degree $1 - \theta$ in wealth, i.e., there exists a function $\phi(\alpha)$ such that*

$$V(W, \alpha) = \phi(\alpha)^{-\theta} \frac{W^{1-\theta}}{1-\theta}. \quad (3)$$

Proof. Fix $k > 0$ any positive constant. Because all constraints are linear, the optimal contingent plan for initial wealth kW and portfolio α will be just k times the optimal contingent plan for initial wealth W and portfolio share α . Since per period utility is homogeneous of degree $1 - \theta$, it follows that total value with initial wealth kW will be equal to $k^{1-\theta}$ times total value with initial wealth W . Therefore the value function of the problem will also be homogeneous of degree $1 - \theta$. It is also a smooth function of W , therefore it has to admit the functional form indicated in the proposition. The proof is complete. \square

Homogeneity implies that when the investor chooses her portfolio share for next period, her choice will be independent of wealth. Hence there is an optimal portfolio share of wealth to be held in cash. We denote this optimal share by α^* . Now (3) implies that the value function is maximal for any wealth level exactly when $\phi(\alpha)^{-\theta}/(1 - \theta)$ is maximal. Assuming $\theta > 1$, this amounts to maximizing $\phi(\alpha)$, so $\alpha^* = \arg \max_{\alpha} \phi(\alpha)$.

So far we have proved that the optimal policy involves a constant share of liquid wealth, and also that the value function is separable in W and α . To find the optimal portfolio α^* , we adopt the following methodology. We consider a subset of all feasible policies, such that the optimal policy is an element of this subset. We characterize all policies in this subset, and then choose the one that gives most utility to the consumer. This is going to be the first best policy.

In order to make this subset as small as possible while still including the first best, we try to characterize the optimal policy in more detail. As we have seen above, optimality requires holding a constant portfolio share in the liquid asset. Here we also argue that under the optimal policy, once a liquidity shock takes place, the consumer will spend all her liquid wealth on consumption.

This is because the consumer can rebalance her portfolio before the next period anyway, so holding more cash than needed during a liquidity shock is unnecessary. Indeed, it is harmful if the liquidity premium π is positive. Therefore the optimal policy certainly satisfies two conditions: first that it maintains a constant share of the liquid asset, and second that it involves consuming all liquid wealth during a liquidity shock. The fact that consumption during a liquidity shock is determined by the constraint means that the agent is consuming

less than what her unconstrained first best would be; in other words, she would choose to consume more if more cash were available. But ex ante, she trades off this (infrequent) pain caused by the constraint during a liquidity shock against the yield premium of investing more in the illiquid asset.

Given our knowledge that the optimal policy is of the above mentioned form, we can restrict our attention to the set of policies that are also of that form. Specifically, consider contingent plans such that (a) the consumer always rebalances her portfolio to hold a share of wealth $\tilde{\alpha}$ in cash; and (b) she consumes all her liquid wealth during a liquidity shock, otherwise chooses consumption optimally. Here $\tilde{\alpha} \in [0, 1]$ is any fixed value. As we have argued above, the optimal policy is an element of this set, attained for $\tilde{\alpha} = \alpha^*$.

Definition 1. The *modified problem* with portfolio share $\tilde{\alpha}$ is the lifetime consumption problem of an agent when she is restricted to

- (a) always hold a share of wealth $\tilde{\alpha}$ in cash, and
- (b) consume all liquid wealth during a liquidity shock.

For any fixed value of $\tilde{\alpha}$, we can solve the modified problem, which is just solving the optimal problem with two extra restrictions. Because these two restrictions eliminate two of the three control variables to be chosen by the consumer in the original problem, the *modified problem* will turn out to be considerably easier to solve.

As we have shown above, once we manage to solve the modified problem for each value of $\tilde{\alpha}$, all we need to do is maximize total value over $\tilde{\alpha}$ to get the optimal value attainable. The argument of this maximum will give the optimal portfolio share α^* . The objective is, then, first to characterize the value function of the modified problem for any given $\tilde{\alpha}$, and then to maximize over $\tilde{\alpha}$ to get α^* .

Let $\tilde{V}(W, \tilde{\alpha})$ denote the value function of the modified problem. Then the previous arguments are summarized in the following proposition.

Proposition 2. *The value of the modified problem is less than or equal to the value of the original problem, with equality only at the optimal portfolio share α^* . Formally,*

$$\tilde{V}(W, \tilde{\alpha}) \leq V(W, \tilde{\alpha}) \tag{4}$$

with equality if and only if $\tilde{\alpha} = \alpha^*$.

Using homogeneity of the per period utility function, it is easy to show that the value function of this modified problem will also be homogeneous of degree $1 - \theta$. Therefore with some suitable function $f(\tilde{\alpha})$, we can write

$$\tilde{V}(W, \tilde{\alpha}) = f(\tilde{\alpha})^{-\theta} \frac{W^{1-\theta}}{1-\theta}. \quad (5)$$

The next step is to characterize the function $f(\cdot)$. This amounts to solving the modified problem. Once we have done this, all we need to do is maximize $f(\tilde{\alpha})$ in $\tilde{\alpha}$ to find α^* and the first best rule.

The value function of the modified problem satisfies the following Bellman equation

$$\begin{aligned} \tilde{V}(W, \tilde{\alpha}) = \mu \left[\gamma^\theta \frac{C_1^{1-\theta}}{1-\theta} + \beta \tilde{V}(R_p W - C_1, \tilde{\alpha}) \right] + \\ (1-\mu) \max_{C_2 \leq (\alpha + r_p)W} \left[\frac{C_2^{1-\theta}}{1-\theta} + \beta \tilde{V}(R_p W - C_2, \tilde{\alpha}) \right], \quad (6) \end{aligned}$$

where by assumption $C_1 = (\tilde{\alpha} + r_p)W$. Indeed, this is just rewriting the original Bellman equation (2), imposing the restrictions that the portfolio share has to be equal to $\tilde{\alpha}$ and that the agent consumes all her liquid wealth during a shock. In the modified problem only consumption in the normal state, C_2 , remains as a choice variable. This consumption level is determined as follows. If the consumer's liquidity constraint does not bind in the normal state (because she has enough cash), then C_2 will be chosen according to the first-order condition of the problem. From the Bellman equation (6), using the expression about the functional form (5), we find this first order condition for C_2 to be

$$C_{2,\text{foc}}^{-\theta} = \beta f(\tilde{\alpha})^{-\theta} [(R_p - C_2)W]^{-\theta}.$$

If the liquidity constraint happens to bind, that is, the consumer has too little cash to consume at the unconstrained first best, then she will choose to consume all of her liquid wealth. This gives

$$C_{2,\text{constr}} = (\tilde{\alpha} + r_p)W.$$

Clearly, $C_{2,\text{foc}}$ is only feasible if $C_{2,\text{foc}} \leq C_{2,\text{constr}}$. Hence in general the consumer will choose C_2 to be the smaller of these two consumption levels, that is

$$C_2 = \min \left\{ (\tilde{\alpha} + r_p)W, \frac{f(\tilde{\alpha})\beta^{-1/\theta}}{1 + f(\tilde{\alpha})\beta^{-1/\theta}} R_p W \right\}. \quad (7)$$

Let $\underline{\alpha}$ denote the threshold level of the portfolio where the liquidity constraint becomes binding in the normal state. In other words, $\underline{\alpha} \in (0, 1)$ is the value of $\tilde{\alpha}$ that exactly equates $C_{2,\text{foc}}$ and $C_{2,\text{constr}}$. We then have the following proposition.

Proposition 3. *Using $\underline{\alpha}$ determined above, the value function of the modified problem is characterized as follows. If $\tilde{\alpha} < \underline{\alpha}$ then $f(\tilde{\alpha})$ is given by*

$$f(\tilde{\alpha}) = \left[\frac{1 - \beta(1 - \tilde{\alpha})^{1-\theta}}{\mu\gamma^\theta + (1 - \mu)} \right]^{1/\theta} (\tilde{\alpha} + r_p)^{1-1/\theta}. \quad (8)$$

If $\tilde{\alpha} \geq \underline{\alpha}$ then $f(\tilde{\alpha})$ is implicitly determined as the root of the following equation

$$1 - \mu\beta(1 - \tilde{\alpha})^{1-\theta} - (1 - \mu)\beta R_p^{1-\theta} (1 + f(\tilde{\alpha})\beta^{-1/\theta})^\theta = f(\tilde{\alpha})^\theta \mu\gamma^\theta (\tilde{\alpha} + r_p)^{1-\theta}. \quad (9)$$

Proof. The structure of the proof is the following. We look at two cases depending on whether the liquidity constraint of the agent is binding or not. The formulas for $C_{2,\text{constr}}$ respectively $C_{2,\text{foc}}$ give the optimal consumption choice in these two cases. Once we have the optimal consumption choice, we just substitute into both sides of the Bellman equation (6) in both cases. Given the functional form (5), straightforward calculations imply equations (8) and (9) in the two cases.

We know that for $\tilde{\alpha} < \underline{\alpha}$ the constraint is binding, hence equation (8) holds. Similarly, for $\tilde{\alpha} > \underline{\alpha}$ the constraint is not binding, and so equation (9) holds. This is the statement of the Proposition. \square

The last Proposition gives a complete characterization of the consumer's optimal choice in the modified problem. In order to find the optimal policy in the original problem, all we need to do is maximize the function $f(\tilde{\alpha})$ implicitly determined in the proposition. Unfortunately, this last step of the derivation cannot be performed analytically; we will need to rely on numerical methods to get $f(\tilde{\alpha})$ for each value of $\tilde{\alpha}$, and then also to maximize in $\tilde{\alpha}$. Nevertheless, the fact that finding the solution only takes the numerical maximization of a

function in one variable can be considered a substantial simplification; especially compared to other numerical methods that rely on iterative procedures approximating the whole value function.

2. Calibration Results

In this section we outline our numerical strategy and attempt to calibrate the model. We start out using a set of benchmark parameter values to solve for the optimal portfolio share of cash. Then we investigate the effects of changing the underlying parameters on the optimal portfolio. We look at changes in each parameter separately, and in some cases also look at interaction effects.

Our numerical strategy is the following. For each candidate portfolio share $\tilde{\alpha}$ we solve for $f(\tilde{\alpha})$ using equation (9), that is, assuming that the liquidity constraint is not binding. Given $f(\tilde{\alpha})$, we verify whether the assumption that the liquidity constraint is not binding, $C_{2,\text{foc}} \leq C_{2,\text{constr}}$ indeed holds. This is equivalent to checking the condition $\tilde{\alpha} \geq \underline{\alpha}$. If the condition is satisfied, then the true value of $f(\tilde{\alpha})$ is indeed given by equation (9). If the condition is not satisfied, then it has to be that $\tilde{\alpha} < \underline{\alpha}$ and hence the true value of $f(\tilde{\alpha})$ is given by (8).

Once $f(\tilde{\alpha})$ has been determined for each value of $\tilde{\alpha}$, we maximize the function in the range $\tilde{\alpha} \in (0, 1)$. By Proposition 2, the argument α^* corresponding to the maximum gives the value function of the original problem. Also, the number α^* is just the optimal portfolio share in the original problem. Once we have the value function, we can easily derive the optimal choice of consumption in the two states.

Given this simple numerical procedure we can evaluate the predictions of the model for any set of underlying parameters. In the calibration exercise our most important target variable will be the optimal portfolio share α^* . We will also look at the second moment of consumption in order to better match available data.

2.1. BENCHMARK CALIBRATION

Our benchmark set of underlying parameters is reported in Table 1.

β	θ	r	π	Δt	μ	γ
0.9	2	0.02	0.04	1/12	0.25	11

Table 1: Benchmark parameter values (annualized)

Whenever possible, we annualized the parameters. These benchmark values can be interpreted as follows.

The assumption of $\beta = 0.9$ is somewhat unusual, a value of 0.96 is more common in the macro calibration literature. We have chosen a lower discount rate because we would like to look at consumers with short horizons, and then make a comparison with longer horizons. Additionally, relatively low discount factors (in the range of 0.85-0.90) have been estimated in Carroll and Samwick (1997).

We take the coefficient of relative risk aversion, θ to be equal to 2. This is at the lower end of the range of plausible values for risk aversion. This choice emphasizes the fact that our results are not driven by high risk aversion. Indeed, it is going to turn out that increasing θ has the somewhat counterintuitive implication of decreasing the share of cash in the optimal portfolio.

The annual risk free rate is $r = 0.02$. This is close to the US average real risk free rate of 2.02% during the time period 1891–1998, as reported in Campbell (2001).

The liquidity premium is chosen at a yearly four percent, $\pi = 0.04$. Unfortunately there are not many studies that actually come up with the liquidity premium of a typical household portfolio. For long horizon U.S. treasury bonds, Campbell (2001) finds an average excess return not exceeding 1.5%; this he attributes at least partly to liquidity reasons. However, we think that most of the assets in a typical household’s portfolio are less liquid than long bonds, and accordingly carry a larger yield premium. Indeed, we believe that the effects of illiquidity on returns may be comparable to those of risk. Accordingly we have

chosen 4% as an intermediate value between the term premium of 1.5% and the equity premium of about 6%.

The length of a time interval is 1 month, that is, $\Delta t = 1/12$. This means that an order is executed one month after it is placed. This seems like a considerable amount of illiquidity, partially justifying the high liquidity premium. We shall see how varying this waiting time changes our results. Again, due to lack of data on the illiquidity of a household portfolio, we have to rely on varying the interesting parameters, in this case the time interval, in a reasonable range.

Due to the setup of the model, the length of the time interval is also the length of the liquidity shock. Although this might seem like a serious weakness, it is not. The reason is that if a liquidity shock lasts longer than the waiting time before trade takes place, the consumer can optimally trade to counteract her continuing liquidity shock after the trade is executed. Thus a longer liquidity shock will not change the optimal portfolio substantially.

Liquidity shocks occur once in four years on average, that is $\mu = 0.25$. The size of these shocks ($\gamma = 11$) is such that the agent consumes roughly ten times as much during liquidity shocks as at normal times.

The last two benchmark values have been chosen by looking at their implied moments on consumption growth. At the current parameter values there is one month in every four years with a large unforeseen liquidity shock, that raises monthly consumption to 9.8 times its normal level. This corresponds to a 2.43% standard deviation of annual consumption growth, $\sigma(\Delta \ln C_t)$. In Campbell (2001) U.S. aggregate consumption growth is reported to have a standard deviation of just above 1%. However, the same paper reports international aggregate consumption growth to be considerably more volatile, with standard deviations ranging from 1.7% (Italy) to 2.9% (France). On top of that, individual consumption is probably a lot more volatile than aggregate consumption. Hence our yearly standard deviation of 2.43% can be considered as a reasonable match with the data.

Solving the model using the numerical procedure outlined earlier, we get that the optimal portfolio share of liquid assets is $\alpha^* = 5.0\%$. This number is roughly in line with empirical studies. Bertaut and Starr-McCluer (2001) find

liquid wealth to be 3.5% of total assets in aggregate data, a slightly lower value than our result. They also look at U.S. household survey data from the Survey of Consumer Finances (SCF), and find that average liquid wealth to total assets varies between 4.6% and 5.7% over time. Our calibration result is roughly in the middle of this range. Angeletos et al. (2001) report that average liquid assets to total assets varies between 8% and 16% depending on the exact definition of liquid assets. They use data from the 1995 Survey of Consumer Finances. These estimates are somewhat larger than our results.

Additionally, Elton and Gruber (2000) point out that the asset allocation recommendations of major investment banks contain 5% cash for liquidity reasons. This is exactly in line with our calibration results.

2.2. COMPARATIVE STATICS

Next we turn to comparative statics. Our strategy is to let one of the parameters vary while holding the others fixed at their benchmark values. We are interested in how each of the parameters of the model affect the share of the liquid asset in the optimal portfolio.

First let us examine how the two main building blocks of the model, that is, the liquidity shock and the market lag, affect the optimal portfolio rule. We shall examine separately how the share of cash varies as we change the magnitude and frequency of the taste shock (γ and μ) and also how it varies as we change the degree of illiquidity (Δt). Additionally, to get a better understanding of the model, we also look at the interaction of these two effects. In other words, we also examine whether illiquidity matters more when liquidity shocks are larger.

Figure 2 shows how the optimal portfolio rule changes as a function of the size of the liquidity shock (γ). Clearly there is a very strong positive relationship between γ and the optimal share of cash in the portfolio; moreover, this relationship is close to linear. Intuitively, the larger the size of the shock, the more the consumer wishes to consume when it actually hits; accordingly, a higher buffer-stock of liquid assets is required to finance that higher level of consumption. As the size of the liquidity shock vanishes, that is, γ goes to one, the share of liquid wealth in the portfolio becomes very small.

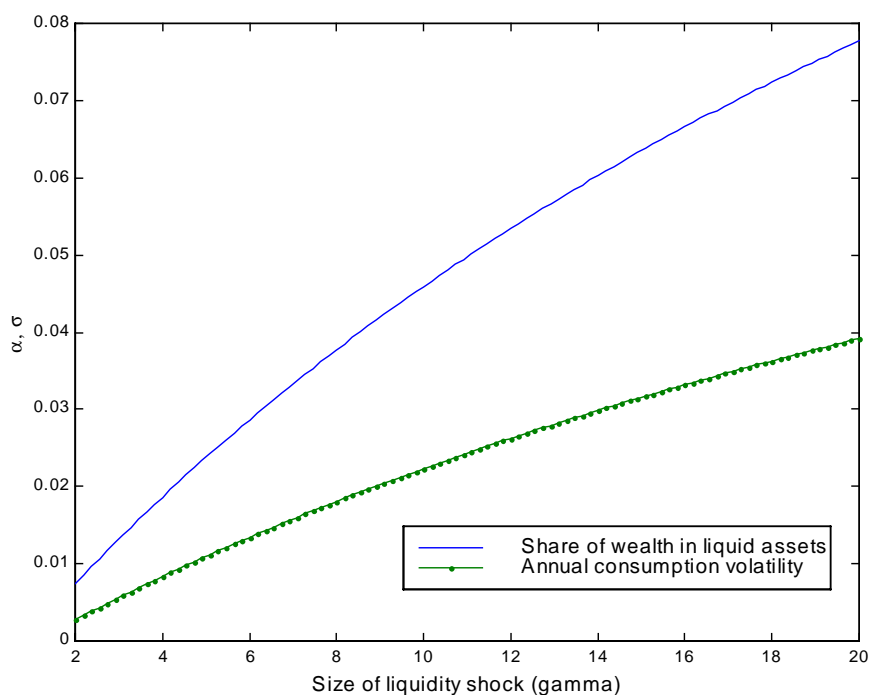


Figure 2: Optimal portfolio share and consumption volatility as a function of the size of the liquidity shock

Notes: The fraction of wealth in liquid assets (α) can vary between 0 and 1. Consumption volatility (σ) is the unconditional standard deviation of annual log consumption growth.

Because the parameters governing the liquidity shock (its frequency, μ , and its size, γ) are difficult to interpret directly, we also report the implied standard deviation of annual consumption growth. The figure shows that there is a clear one-to-one correspondence between the size of the shock and the standard deviation of consumption. A larger liquidity shock, unsurprisingly, gives rise to more volatile consumption. The magnitude of the standard deviation of consumption, as depicted in the figure (in the range of 1–4%), suggests that the size of shocks we consider here is not unrealistic.

The exact relationship between liquidity shocks, portfolio choice and consumption volatility depends on the distribution of taste shocks. This in turn is governed not only by γ , but also by μ , the frequency of liquidity events.

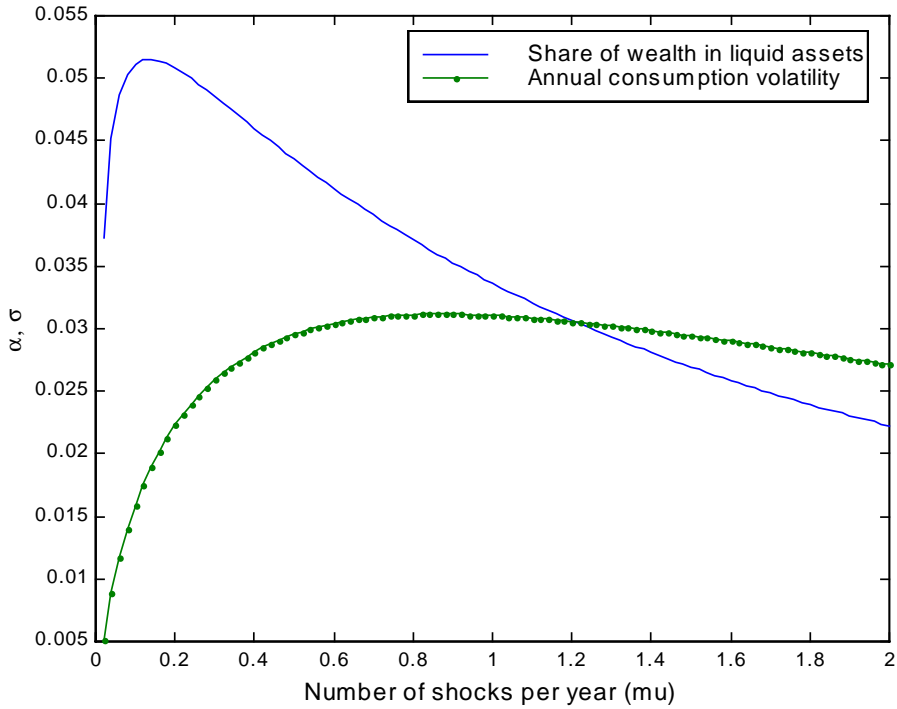


Figure 3: Optimal portfolio share and consumption volatility as a function of the frequency of the liquidity shock

Notes: The fraction of wealth in liquid assets (α) can vary between 0 and 1. Consumption volatility (σ) is the unconditional standard deviation of annual log consumption growth. Frequency of liquidity shocks (μ) is annualized so that $\mu = 2$ means 2 shocks a year on average.

Figure 3 depicts cash holdings to total assets and the standard deviation of consumption as a function of μ . A striking observation is that both functions are hump-shaped; they first increase with μ and then slowly decrease. The intuition for why the portfolio share of cash increases with μ for μ small is straightforward. As liquidity shocks become more frequent, holding too little buffer-stock hurts the investor more often, so she will tilt her portfolio towards cash. On the other hand, once the cash reserve is high enough to finance an almost optimal consumption plan, another effect starts to set in: more frequent liquidity shocks raise the effective discount factor of the consumer. This is because she expects to have on average higher marginal utility in the future. Thus it is as if the

consumer became more patient. She then modifies her optimal consumption pattern accordingly, consuming less in both possible states. But then a lower buffer-stock is sufficient to finance her lower consumption. A similar intuition applies as to why consumption volatility is increasing for μ small and decreasing afterwards.

With the benchmark parameters the portfolio share of cash is maximal for $\mu = 0.136$, that is, when roughly one liquidity shock occurs in every seven years. This suggests that even rare events have a sizable effect on portfolio choice when assets are illiquid. It seems, therefore, that consumers hold a significant share of their portfolio in liquid assets in order to prepare for unforeseen (or, at least very unlikely) contingencies.

Let us now turn to the question of how the degree of illiquidity affects portfolio choice. The relationship between the waiting time after an order is placed and the optimal share of wealth held in cash is depicted in Figure 4. When varying the length of time periods, Δt , we adjust all the other parameters such that their annualized value remains fixed at the benchmark level.

As expected, the more illiquid the asset is, that is, the more the investor has to wait before her order is fulfilled, the higher the share of liquid assets in the optimal portfolio. This is because when the investor suffers from a liquidity shock, she would wish to consume more during the entire waiting period, and in order to do that she needs to hold a higher buffer stock in cash.

One might think that this result strongly depends on the assumption that the order lag lasts exactly as long as the liquidity shock. However, the result would not change qualitatively if we allowed the length of the liquidity shock to be larger than the order lag. The reason for this is that once the taste shock hits, the investor can immediately place a sell order and consume its proceeds from the next period on. Thus the only reason for keeping a cash reserve is the first period of the taste shock.

Concerning the actual numbers, note that at the benchmark parameter values, even a small degree of illiquidity (say, 1 week) has a noticeable impact on the portfolio decision of the investor.

It is also instrumental to see how the effects of asset illiquidity interact with

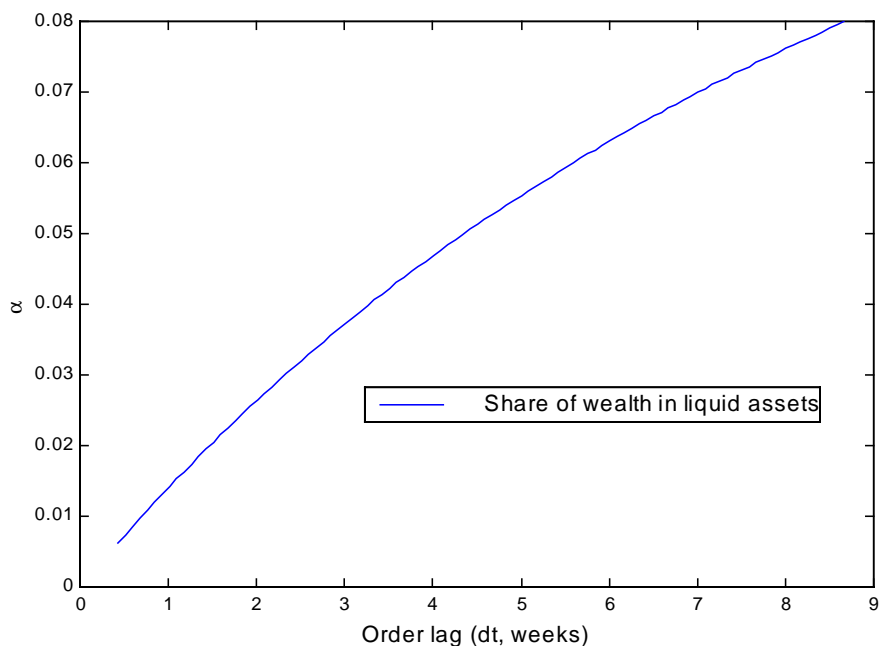


Figure 4: Optimal portfolio share as a function of the order lag

Notes: The fraction of wealth in liquid assets (α) can vary between 0 and 1. Order lag (dt) is reported in weeks.

the liquidity shock. Figure 5 shows the optimal portfolio share of cash as a function of both the order lag (Δt) and the size of the liquidity shock (γ). The interesting finding is that illiquidity matters more when liquidity shocks are large. That is, the size of the cash reserve increases faster in the degree of illiquidity when γ is high. For a small liquidity shock, or, at the extreme, without any liquidity shocks ($\gamma = 1$), not even a large degree of illiquidity has a significant impact on portfolio holdings. This finding underlines the importance of time-varying consumption needs in explaining liquidity-related phenomena.

Next we turn to examine how asset returns influence portfolio choice. Ultimately, one would wish to assess how the liquidity premium affects the demand for both liquid and illiquid assets.

We first look at the impact of a simultaneous increase in the return of the two assets. In Figure 6 we depict the portfolio share of cash as a function of r , the return on the liquid asset. Recall that the return on the illiquid asset is just

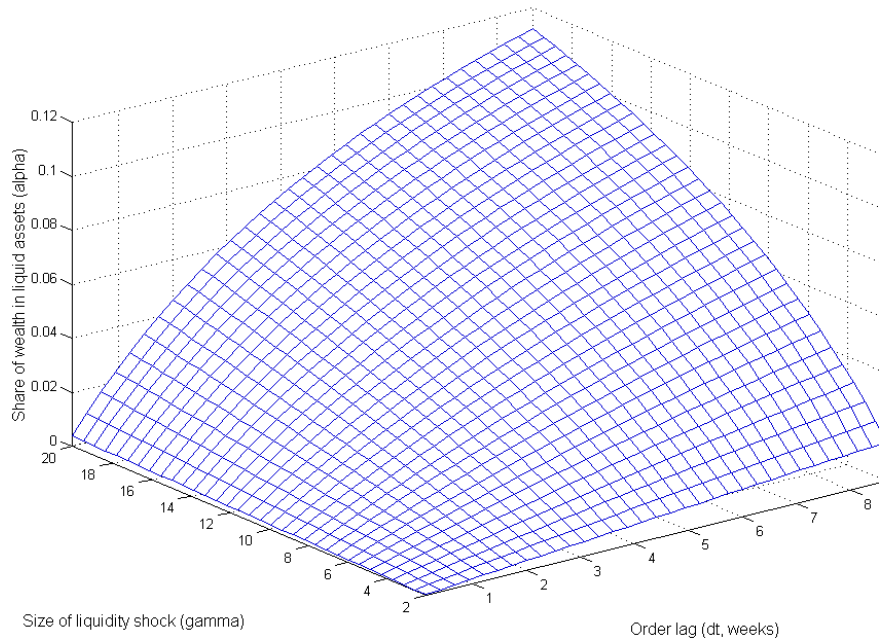


Figure 5: Interaction of the order lag and the liquidity shock

Notes: The fraction of wealth in liquid assets (α) can vary between 0 and 1. Order lag (dt) is reported in weeks.

the liquid return plus a fixed premium. We can see that cash holding increases in the interest rate. The reason for that is when the return on the portfolio increases, this influences the optimal consumption plan via a substitution effect and an income effect. For $\theta > 1$, the income effect dominates, therefore optimal consumption increases each period. But in order to finance that, the agent needs to hold more liquid assets in her portfolio. The relationship between portfolio return and the demand for cash is practically linear, with a 1%point increase in return causing 0.23%point increase in demand.

A more exciting question is how the liquidity premium affects portfolio choice. Figure 7 shows that a higher premium on the illiquid asset increases the demand for cash. This result runs counter to our first intuition but can be explained as follows. On the one hand, as the liquidity premium increases, so does the total return on the portfolio (recall that the investor holds over 90% of her wealth in the illiquid asset). The income effect described above still

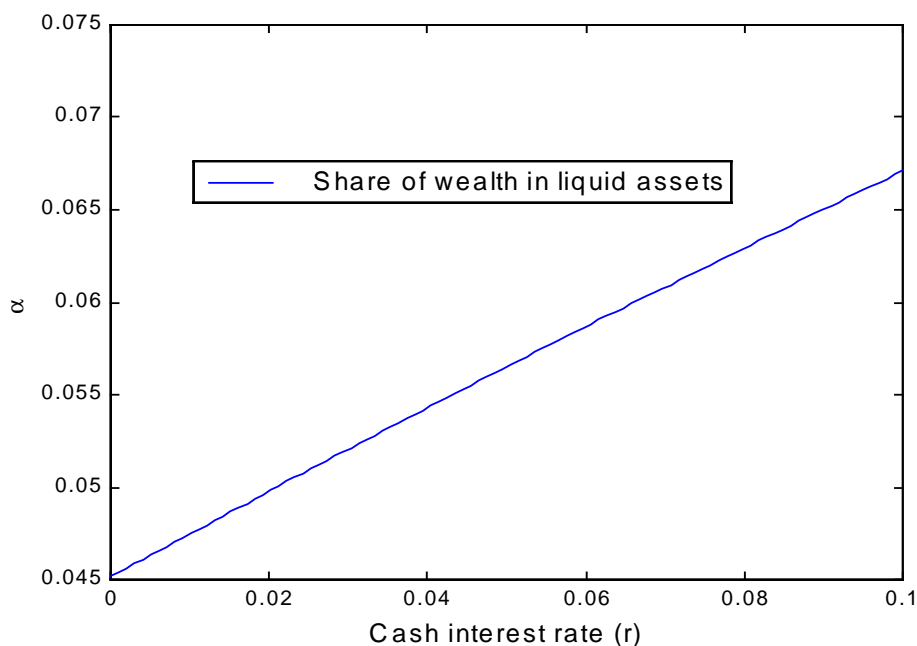


Figure 6: Optimal portfolio share as a function of the return on the liquid asset

Notes: The fraction of wealth in liquid assets (α) can vary between 0 and 1. Return (r) is measured in annual percentage terms, $r = 0.02$ means a 2% yearly return.

raises consumption and hence the required cash holding. On the other hand, because the illiquid asset is more attractive, the agent wishes to allocate less of her wealth to cash. It turns out that the income effect dominates this trade-off in the empirically relevant range, tilting the optimal portfolio towards the liquid asset. It follows that cash demand is increasing in the liquidity premium. Note, however, that the size of the increase is considerably smaller than in the case of a pure income effect. A 1%point rise in the excess return of the illiquid asset brings about a rise of only 0.1%point in the portfolio demand for cash, which is roughly half of what we found for an overall increase in returns. This suggests that, as expected, there is some substitution away from cash happening in the background.

Another question is how risk aversion affects the demand for liquid assets. One might expect that more risk-averse investors hold more cash since they tolerate less the risk of being constrained by lack of liquidity. However, because

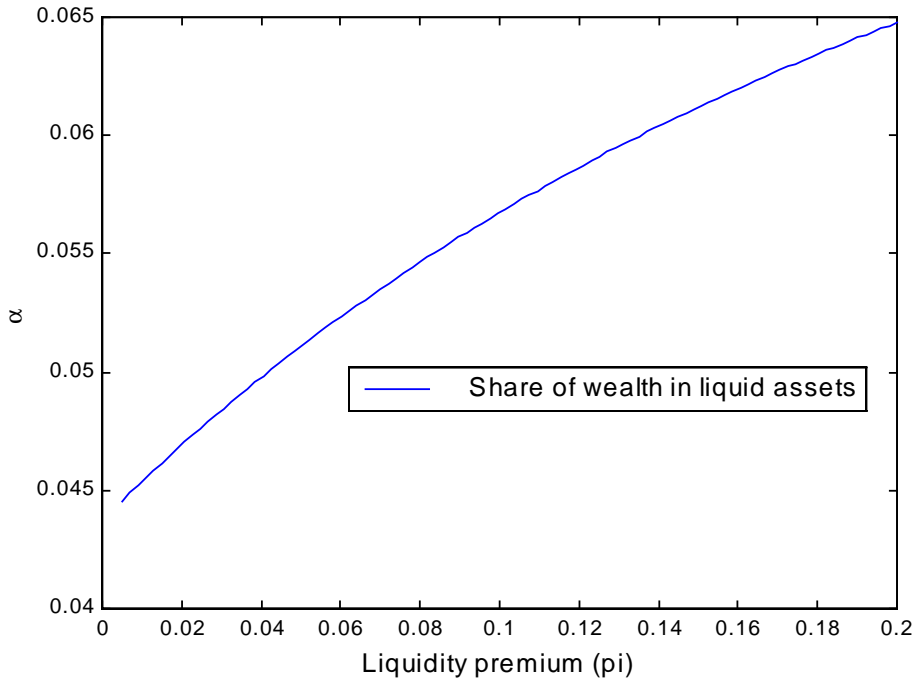


Figure 7: Optimal portfolio share as a function of the premium on the illiquid asset

Notes: The fraction of wealth in liquid assets (α) can vary between 0 and 1. Premium (π) is measured in annual percentage terms, $\pi = 0.04$ means a 4% annual excess return.

of the power utility specification, the coefficient of relative risk aversion, θ , is also the inverse of the elasticity of intertemporal substitution. Thus varying θ brings about intertemporal considerations, too. In particular, a consumer with high θ is less willing to accept a steady decline or growth in consumption over time. Therefore her optimal consumption policy will be such that her wealth does not change dramatically. In our parametrization this implies that a more risk-averse investor will consume less out of her current wealth, hence a lower buffer stock of liquid assets is sufficient.

Figure 8 depicts liquid asset holdings as a function of θ , showing that more risk averse investors hold less cash. This means that the effect of a lower intertemporal substitution dominates that of a higher risk aversion in our benchmark case.

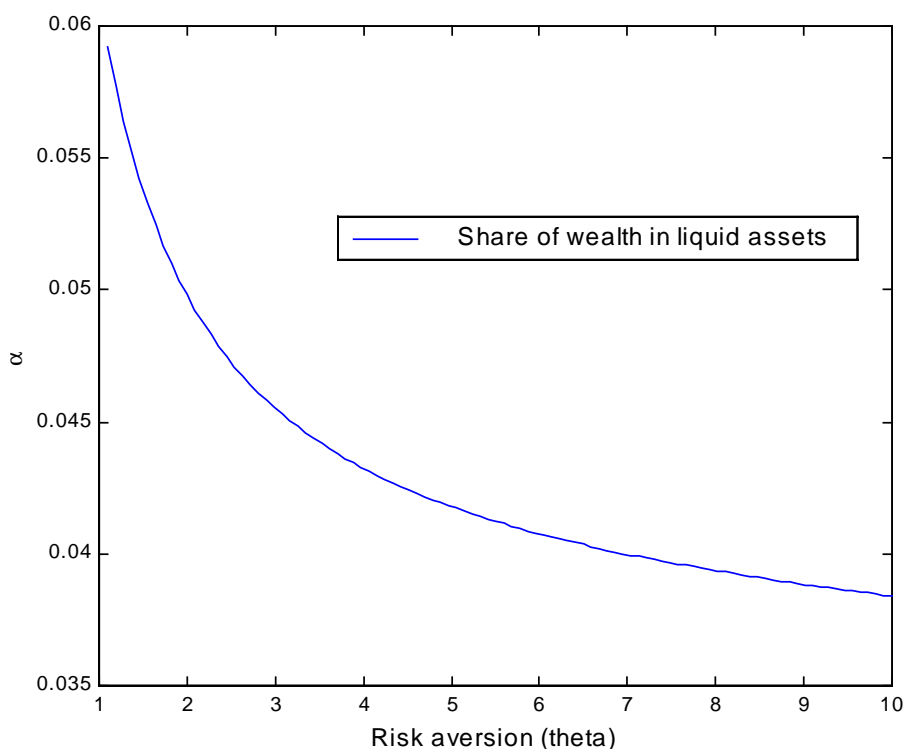


Figure 8: Optimal portfolio share as a function of risk aversion

Notes: The fraction of wealth in liquid assets (α) can vary between 0 and 1.

Finally, let us discuss the effects of investor patience. Figure 9 shows that more patient investors, that is, those with a higher annual discount factor (β) hold less of their wealth in liquid assets. Consumers with a shorter horizon (low β) value future consumption less and therefore allocate more of their consumption to the present. But since consumption is financed from liquid wealth, short-horizon investors should hold a larger fraction of their wealth in liquid form. Note that unlike in Campbell and Viceira (2002), this horizon effect arises in a completely stationary environment.

3. Conclusion

We developed a model of portfolio optimization in the presence of illiquid assets. Our results suggest that illiquidity is an important factor in portfolio investment




Figure 9: Optimal portfolio share as a function of the discount factor

Notes: The fraction of wealth in liquid assets (α) can vary between 0 and 1. The reported discount factor is annualized.

decisions. Our calibration results imply a portfolio demand for cash that is in the range of 2 – 6% of total wealth. This is in line with asset allocation recommendations by practitioners and also with survey data from the Survey of Consumer Finances. Qualitatively, we find that long horizon investors invest more in illiquid assets, whereas short horizon investors prefer cash. We also find, somewhat surprisingly, that the portfolio share of liquid wealth is increasing in the liquidity premium over the empirically relevant range. Additionally, even a very infrequent liquidity shock can give rise to a sizeable buffer stock cash holding, verifying the intuition that consumers hold liquid assets mostly to prepare for unforeseen (or at least unlikely) contingencies. Overall the present study indicates that differences across assets unrelated to risk and return may have a sizeable influence on portfolio demand.

Directions for future research may include extending the model to incorporate price risk and/or risky labor income as well as illiquidity to study the joint effect of these forces on portfolio decision problems. Another desirable extension would be to close the model and solve for equilibrium.

References

- Abel, A. B. (1990). Asset prices under habit formation and catching up with the Joneses, *American Economic Review Papers and Proceedings* **80**: 38–42.
- Amihud, Y. and Mendelson, H. (1986). Asset pricing and the bid-ask spread, *Journal of Financial Economics* **17**: 223–249.
- Angeletos, G.-M., Laibson, D., Repetto, A., Tobacman, J. and Weinberg, S. (2001). The hyperbolic consumption model: Calibration, simulation, and empirical evaluation, *Journal of Economic Perspectives* **15**(3): 47–68.
- Baldwin, C. Y. and Meyer, R. F. (1979). Liquidity preference under uncertainty, *Journal of Financial Economics* **7**: 347–374.
- Benczúr, P. (2002). Identifying sovereign bond risks, *National Bank of Hungary Working Paper* **2002**(2).
- Bertaut, C. C. and Starr-McCluer, M. (2001). Household portfolios in the United States. *in* Guiso, Haliassos and Jappelli (eds), *Household Portfolios*, MIT Press, Cambridge, MA.
- Brennan, M. J. and Subrahmanyam, A. (1996). Market microstructure and asset pricing: On the compensation for illiquidity in stock returns, *Journal of Financial Economics* **41**: 441–464.
- Campbell, J. Y. (2001). Consumption-based asset pricing, Manuscript. Harvard University.
- Campbell, J. Y. and Viceira, L. M. (2002). *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press.

- Campbell, J. Y., Lo, A. and MacKinley, A. C. (1997). *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.
- Canner, N., Mankiw, N. G. and Weil, D. N. (1997). An asset allocation puzzle, *American Economic Review* **87**: 181–191.
- Carroll, C. D. and Samwick, A. A. (1997). The nature of precautionary wealth, *Journal of Monetary Economics* **40**.
- Constantinides, G. M. (1986). Capital market equilibrium with transaction costs, *Journal of Political Economy* **94**(4): 842–862.
- Diamond, D. W. and Dybvig, P. H. (1983). Bank runs, deposit insurance and liquidity, *Journal of Political Economy* **91**(3): 401–419.
- Dumas, B. and Luciano, E. (1991). An exact solution to a dynamic portfolio choice problem under transaction costs, *Journal of Finance* **46**(2): 577–595.
- Elton, E. J. and Gruber, M. J. (2000). The rationality of asset allocation recommendations, *Journal of Financial and Quantitative Analysis* **35**: 27–41.
- Fiori, F. (2000). Liquidity premia in the equity markets: An investigation into the characteristics of liquidity and trading activity. Unpublished paper, University of Chicago.
- Grossman, S. J. and Laroque, G. (1990). Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods, *Econometrica* **58**(1): 25–51.
- Heaton, J. and Lucas, D. J. (1996). Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* **104**: 443–487.
- Huang, M. (2001). Liquidity shocks and equilibrium liquidity premia. Unpublished paper, Stanford University.
- Lippman, S. A. and McCall, J. J. (1986). An operational measure of liquidity, *American Economic Review* **76**(1): 43–55.

- Lo, A. W., Mamaysky, H. and Wang, J. (2001). Asset prices and trading volume under fixed transaction costs. Unpublished paper, MIT.
- Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* **3**: 373–413.
- Merton, R. C. (1973). An intertemporal capital asset pricing model, *Econometrica* **41**(5): 867–887.
- Vayanos, D. (1998). Transaction costs and asset prices: A dynamic equilibrium model, *Review of Financial Studies* **11**: 1–58.