Journal of Financial Economics 52 (1999) 47–78

Transaction costs and predictability: some utility cost calculations☆

Pierluigi Balduzzi[a], Anthony W. Lynch[b,*]

[a]*Boston College, Chestnut Hill, MA, 02167, USA*
[b]*Stern School of Business, New York University, New York, NY 10012, USA*

Received 17 October 1996; received in revised form 2 January 1998

Abstract

We examine the loss in utility for a consumer who ignores any or all of the following: (1) the multi-period nature of the consumer's portfolio-choice problem, (2) the empirically documented predictability of asset returns, or (3) transaction costs. Both the costs of behaving myopically and ignoring predictability can be substantial, although allowing for intermediate consumption reduces these costs. Ignoring realistic transaction costs (fixed and proportional) imposes significant utility costs that range from 0.8% up to 16.9% of wealth. For the scenarios that we consider, the presence of transaction costs always increases the utility cost of behaving myopically, but decreases the utility cost of ignoring predictability. © 1999 Elsevier Science S.A. All rights reserved.

JEL classification: G11; G12

Keywords: Portfolio choice; Transaction costs; Return predictability; Utility cost

* Corresponding author contact information. Tel.: 212-998-0350.

☆This research was initiated while the first author was affiliated with New York University. The authors thank Jennifer Carpenter, Silverio Foresi, and Robert Whitelaw for many helpful conversations. They also thank Ravi Bansal, seminar participants at the 1996 Western Finance Association Meetings, a joint NYU-Columbia conference, and Boston University, and especially Nick Barberis (the referee) for useful comments. Financial support from New York University Summer Research Grants is gratefully acknowledged.

PII: S0304-405X(99)00004-5

1. Introduction

Recent financial research investigates the impact of transaction costs and return predicatability on asset allocation and consumption decisions. Studies that incorporate transaction costs typically assume that the opportunity set is constant through time. This assumption is inconsistent with the documented predictability of U.S. asset returns.[1] On the other hand, studies that examine the impact of return predictability on expected utility and portfolio rebalancing typically do not consider transaction costs.

Our paper contributes to this literature by considering a multiperiod individual who possibly faces both transaction costs and predictable returns. We numerically solve the individual's multiperiod problem in the presence of transaction costs and predictability. In particular, we characterize the consumer's optimal portfolio choice with both fixed and proportional transaction costs, and with return predictability similar to that observed for the U.S. stock market.

Armed with this solution technique, we consider the utility costs associated with three different dimensions of ignorance on the part of the consumer. First, the consumer can behave in either a *myopic* or a *dynamic* manner, depending on whether the chosen portfolio weights maximize a one-period or a multiple-period problem. A myopic investor does not hedge against future changes in the investment opportunity set. Second, the consumer can use either the *conditional* or the *unconditional* distribution of asset returns for decision-making. A consumer that uses the unconditional distribution is ignoring return predictability when making decisions. If the risk-free rate is constant and transaction costs are zero, myopic and dynamic consumers who use the unconditional distribution rebalance identically because their opportunity set is constant across states and over time (see Samuelson, 1969). Finally, the consumer can either account for *transaction costs* or ignore them (*no-transaction-costs* problem). Thus, we examine the impact of transaction costs on the utility costs associated with ignoring predictability.

Transaction costs can be modeled either as *proportional* to the change in the holding of the risky asset, or as a *fixed* fraction of portfolio value. Constantinides (1986) finds that proportional transaction costs affect portfolio choice since the optimal policy is a no-trade region with return to the closer boundary when rebalancing. Davis and Norman (1990) consider the same problem, and are able to solve it exactly, without imposing restrictions on the consumption process. Morton and Pliska (1993) and Schroder (1995) characterize portfolio choice with fixed transaction costs and find that the optimal policy is a no-trade region

[1] Campbell (1987) and Fama and French (1989), among others, find that stock return variation can be explained by the one-month Treasury bill rate, by a contemporaneous and a lagged measure of the term premium, and by the dividend yield.

with a single return point inside the region. All of these papers allow continuous decision-making and impose a constant opportunity set.[2]

Our paper extends this literature in several ways. First, our discrete-time setting allows the utility cost of ignoring transaction costs to be quantified. Earlier papers did not consider this utility cost because a consumer who ignores transaction costs, and trades continuously, would end up bankrupt. Second, our law of motion for the opportunity set incorporates the return predictability observed in the data. Finally, the solution technique allows the simultaneous presence of both proportional and fixed transaction costs. This is realistic, since investors are subject to brokerage fees and bid-ask spreads, which generate proportional costs of trading, as well as costs of gathering and processing information, which generate fixed costs.

A number of recent studies examine the effects of ignoring predictability or behaving myopically. Brennan and Schwartz (1996) and Brennan et al. (1996) analyze numerically the impact of myopic versus dynamic decision-making when the consumer bases decisions on the conditional distribution of returns. Campbell and Viceira (1996) obtain a closed-form solution to the consumer's multiperiod problem in the presence of predictability by using log-linear approximations to the budget constraint. Kim and Omberg (1996) solve the continuous-time analog without any approximations. Barberis (1996) considers the effect of asset-return predictability on myopic portfolio choices at different investment horizons; he also considers the effect of predictability on dynamic portfolio choices, when the investor rebalances every year. Finally, Kandel and Stambaugh (1996) explore the effects of predictability in a myopic setting, where the investor rebalances monthly.[3] All of these papers assume away transaction costs on the grounds that existing financial markets, especially for futures contracts, allow for inexpensive hedging. We relax this assumption and quantify the impact of realistic transaction costs for consumers who behave myopically or ignore predictability.

We calibrate both returns and transaction costs to those faced by U.S. investors, and find that the utility costs of behaving myopically and of ignoring predictability can be substantial. Allowing for intermediate consumption

[2] Several recent papers account for transaction costs in general equilibrium models, in empirical testing, and with different decision-making structures. Koo (1991) and Vayanos (1996) investigate the equilibrium implications of proportional transaction costs. Heaton and Lucas (1996) consider a dynamic equilibrium model where infinitely lived workers trade a bond and a stock, and show that large transaction costs are needed to generate a sizable equity premium in equilibrium. Hansen et al. (1995), Luttmer (1996), and He and Modest (1995) develop pricing-operator tests of asset-pricing models which explicitly account for bid-asked spreads and short-sale constraints. Duffie and Sun (1990) model consumers who face fixed transaction costs and decide the interval of time until their next rebalancing at the time of their current rebalancing. The authors find that the optimal interval between trades is a constant.

[3] These last two papers account for the effects of estimation error on portfolio choice.

reduces these utility costs. This result implies that the effects of inefficient portfolio choices are likely to be more important for institutional investors, who manage assets for the long run only, than for individual investors, who consume as time goes by. Further, ignoring realistic transaction costs (fixed and proportional) imposes significant utility costs which range from 0.8% up to 16.0% of wealth.

In all scenarios, the presence of transaction costs increases the utility cost of behaving myopically. So by not accounting for transaction costs, the utility cost estimates of Barberis (1996) and Brennan et al. (1996) represent a lower bound. In contrast, the large utility cost associated with ignoring predictability is reduced by the presence of transaction costs. Thus, the utility cost estimates of Kandel and Stambaugh (1996) represent an upper bound.

Regarding portfolio-choice policies, proportional transaction costs induce a no-trade region such that the consumer rebalances to the nearer boundary when the risky-asset weight goes outside the region. Fixed *and* proportional transaction costs also induce a no-trade region, but now rebalancing places the risky-asset portfolio weight inside the boundary. Thus, the basic nature of the optimal policy identified by previous studies carries over to the case when the risky return is predictable, at least for the parameter values that we consider.[4] Also, we find that realistically-small transaction costs induce sizable no-trade regions.

The paper also makes contributions at the methodological level. We show how the investor's multiperiod problem can be numerically solved in the presence of time varying conditional expected returns, and both proportional and fixed transaction costs. Since our solution technique is quite flexible, it can be applied to a variety of scenarios with different constraints facing the consumer and different laws of motion for the state variables.

The paper is organized as follows. Section 2 describes the consumer's optimization problem. Section 3 calibrates the state variables and asset returns to the U.S. economy. Section 4 describes the various scenarios studied in the paper. Section 5 discusses the portfolio choices, while Section 6 looks at the utility costs for different types of consumer ignorance. Section 7 concludes.

2. Portfolio allocation with transaction costs and predictable returns

This section lays out the preferences of and constraints faced by the consumer. We characterize the optimization problem for a consumer who is either dynamic or myopic, and who does or does not take into account asset-return predictability. We also describe the nature of the utility comparisons to be performed later in the paper, and the solution technique for numerically solving the consumer's problem.

[4] See Constantinides (1986) and Davis and Norman (1990) for proportional transaction costs, and Morton and Pliska (1993) and Schroder (1995) for fixed transaction costs.

2.1. Constraints and preferences

We consider the portfolio allocation between two assets, a risky asset and a riskless asset. The consumer faces transaction costs that are proportional to wealth. The law of motion of the consumer's wealth, W, is given by

$$W_{t+1} = (W_t - c_t)(1 - f_t)[\alpha_t(R_{t+1} - R^{\mathrm{f}}_t) + R^{\mathrm{f}}_t], \quad (1)$$

where c is consumption, α is the share of wealth allocated to the risky asset, R is the rate of return on the risky asset, R^{f} is the risk-free rate, and f is the transactions cost per dollar of portfolio value. We further define κ as the fraction of wealth consumed and R_W as the rate of return on wealth, net of the transaction costs incurred. Hence, we have

$$W_{t+1} = (1 - \kappa_t)W_t R_{W,t+1}. \quad (2)$$

The law of motion for wealth in Eq. (1) implicitly assumes that consumption at time t and any transaction costs to be paid at time t are obtained by liquidating costlessly the risky and the riskless asset in the proportions $\hat{\alpha}$ and $(1 - \hat{\alpha})$, where $\hat{\alpha}$ is the allocation to the risky asset inherited from the previous period:

$$\hat{\alpha}_{t+1} \equiv \frac{\alpha_t(1 - \kappa_t)W_t(1 - f_t)R_{t+1}}{W_{t+1}} = \frac{\alpha_t R_{t+1}}{\alpha_t(R_{t+1} - R^{\mathrm{f}}_t) + R^{\mathrm{f}}_t}. \quad (3)$$

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 risky asset's dividend exceeds the consumption out of the risky asset, $\kappa\hat{\alpha}W$, a dividend reinvestment plan can be used to costlessly reinvest the excess dividend in the risky asset.

In general, we model the cost of transacting, f, as a function of the difference between the end-of-period wealth allocation to the risky asset, $\hat{\alpha}$, and α. Specifically, we assume transaction costs to have *two* components:[5]

$$f = \phi_1|\alpha - \hat{\alpha}| + \phi_2 I_{\alpha - \hat{\alpha} \neq 0}, \quad (4)$$

where ϕ_1, $\phi_2 \geqslant 0$ and $I_{\alpha - \hat{\alpha} \neq 0}$ is an indicator function which equals one if $\alpha - \hat{\alpha} \neq 0$ and zero otherwise. The first term is proportional to the change in the value of the risky asset holding, as in Constantinides (1986). The second term is a fixed fraction of the total value of the portfolio, as in Morton and Pliska (1993) and Schroder (1995). This second term reflects the fixed cost of rebalancing the portfolio, regardless of the size of the rebalancing. This fixed cost increases with the investor's wealth, since it is likely to depend on the opportunity cost of the investor's time.

[5] To keep notation simple, we drop the time subscript when all variables in a mathematical expression are contemporaneous.

We assume there is a 'predictive' variable, D (to be explicitly identified later in the calibration exercise), which affects the conditional mean of the risky asset's return. We assume D follows a first-order Markov process. For simplicity, the riskless rate is assumed to be constant, and so $R^{\mathrm{f}}_t = R^{\mathrm{f}}$ for every t.

We consider the optimal portfolio problem of a consumer with a finite life of T periods. Preferences are time separable and exhibit constant relative risk aversion (CRRA). Since earlier papers consider situations where expected lifetime utility depends only on terminal wealth (e.g., Barberis, 1996; Brennan and Schwartz, 1996; Brennan et al., 1996), we use the following preferences in most simulations:

$$\mathrm{E}\left[\frac{W_T^{1-\gamma}}{1-\gamma} \mid D_1, \hat{\alpha}_1\right], \tag{5}$$

where γ is the relative-risk-aversion coefficient. Note that the expected lifetime utility depends on the state of the economy at time 1. Further, the inherited portfolio allocation $\hat{\alpha}_1$ is a state variable when ϕ_1 or ϕ_2 is greater than zero, since its value determines the transaction costs to be paid at time 1. Our framework can also be modified to allow expected lifetime utility to depend on intermediate consumption:

$$\mathrm{E}\left[\sum_{t=1}^{T} \frac{c_t^{1-\gamma}}{1-\gamma} \mid D_1, \hat{\alpha}_1\right], \tag{6}$$

where the rate of time preference equals one. These preferences have been extensively used in empirical work by Grossman and Shiller (1981), Hansen and Singleton (1982), and many others. Values of the time-discount parameter close to one are consistent with the empirical findings for the U.S. economy. Thus, we examine the sensitivity of our results and those in earlier papers to the presence of intermediate consumption.

2.2. Consumer policies

We now consider various optimization problems and the associated Bellman equations for a consumer who is aware of some, but not necessarily all, features of the environment. Throughout, we present Bellman equations for a consumer whose utility does not depend on intermediate consumption. The Appendix discusses the case with intermediate consumption. The fraction of portfolio value allocated to the risky asset at time t is denoted by $\alpha_t = \alpha(D_t, \hat{\alpha}_t, t)$, which is *time dependent* since the time horizon T is finite.

The first dimension of the consumer's problem that we vary is the consumer's time horizon when choosing the portfolio weights. The Dynamic consumer (D policy) correctly considers *all* remaining periods before death when choosing portfolio weights. In this case, the consumer anticipates future changes in the investment opportunity set, which generates hedging demand. On the other hand, the Myopic consumer (M policy) chooses portfolio weights using a *one-period* horizon, although the consumer's remaining lifetime may be longer. Consequently, future changes in the investment opportunity set are not accounted for and, hence, are not hedged against. When utility depends on intermediate consumption, the Myopic consumer is assumed to dynamically choose consumption, taking the portfolio choice to be the myopic choice. This assumption is made to make comparisons between the Myopic and Dynamic consumers meaningful.

Another dimension that we vary is the consumer's beliefs about the law of motion for the state variable D. The individual can ignore any predictability, and behave as if the risky return R is independently and identically distributed (i.i.d.) with a distribution equal to its steady-state or unconditional distribution (the U policy). On the other hand, the individual can incorporate predictability by using the conditional distribution for R when making portfolio choices (the C policy).

Combining these two distinctions gives rise to four optimization problems and associated policies. However, in the absence of transaction costs and given a constant riskless rate, it is well-known that the asset allocation for the M-U policy is the same as that for the D-U policy (see Samuelson, 1969). Thus, the M-U policy is the same as the D-U policy in the no-transaction-cost case.

2.2.1. Dynamic-conditional (D-C) policy

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

$$\frac{a(D_t, \hat{\alpha}_t, t)}{1-\gamma} W_t^{1-\gamma} = \max_{\alpha_t} \left\{ \frac{1}{1-\gamma} W_t^{1-\gamma} \mathrm{E}[a(D_{t+1}, \hat{\alpha}_{t+1}, t+1) R_{W,t+1}^{1-\gamma} \,|\, D_t, \hat{\alpha}_t] \right\},$$
$$\text{for } t = 1, \ldots, T-1, \tag{7}$$

where $\mathrm{E}[.|D, \hat{\alpha}]$ denotes the expectation taken using the conditional distribution given D. This form of the value function derives from the CRRA utility specification in Eqs. (5) and (6), and from the linearity in W of the budget constraint Eq. (2).

2.2.2. Myopic-conditional (M-C) policy

In this case, the optimal portfolio policy is not time dependent, and equals that of a dynamic optimizer in the next-to-last period of life. We have $\alpha_t = \alpha(D_t, \hat{\alpha}_t)$ for $t = 1, \ldots, T-1$, which solves the following problem:

$$\max_{\alpha_t} \left\{ \frac{1}{1-\gamma} \mathrm{E}[R_{W,t+1}^{1-\gamma} | D_t, \hat{\alpha}_t] \right\}. \tag{8}$$

Hence, the consumer takes into account the current value of D, but not any covariance between next period's risky asset return and future realizations of D.

2.2.3. Dynamic-unconditional (D-U) policy

The Bellman equation faced by the consumer is given by

$$\frac{a(\hat{\alpha}_t, t)}{1-\gamma} W_t^{1-\gamma} = \max_{\alpha_t} \left\{ \frac{1}{1-\gamma} W_t^{1-\gamma} \mathrm{E}^U[a(\hat{\alpha}_{t+1}, t+1) R_{W,t+1}^{1-\gamma} | \hat{\alpha}_t] \right\}$$
$$\text{for } t = 1, \dots, T-1, \tag{9}$$

where $\mathrm{E}^U[.|\hat{\alpha}_t]$ denotes expectations taken using the *unconditional* distribution for future realizations of R. This expectation does not depend on D_t, and neither does α_t. Note that the Bellman Eqs. (7) and (9) are solved by *backward iteration*, starting with $t = T-1$ and either $a(D, \hat{\alpha}, T) = 1$ or $a(\hat{\alpha}, T) = 1$.

2.2.4. Myopic-unconditional (M-U) policy

The consumer's portfolio composition policy $\alpha_t = \alpha(\hat{\alpha}_t)$ for $t = 1, \dots, T-1$ solves the following problem:

$$\max_{\alpha_t} \left\{ \frac{1}{1-\gamma} \mathrm{E}^U[R_{W,t+1}^{1-\gamma} | \hat{\alpha}_t] \right\}. \tag{10}$$

Note that the only difference between the conditional polices D-C and M-C and the unconditional policies D-U and M-U is the distribution used to take expectations.

2.2.5. Transaction costs

The last dimension along which the consumer can exhibit ignorance is transaction costs. Of course, this is not possible if transaction costs are zero. Thus, there are only three possible problems in the no-transaction cost economy: D-U-N, M-C-N, D-C-N, where N denotes that there are no transaction costs.

In the presence of transaction costs, the consumer could use the αs from the analogous no-transaction cost case (denoted N-T) or could choose αs taking transaction costs into account (denoted T). Thus, when utility only depends on terminal wealth, the N-T policy is the same as the analogous N policy described above. On the other hand, when utility depends on intermediate consumption, the consumer is assumed to make consumption decisions taking transaction costs into account, and so each of the N-T consumption policies differs from the analogous N policy. Combining the transaction cost dimension with the two above gives seven policies in the transaction cost economy: D-U-N-T, M-C-N-T, D-C-N-T, M-U-T, D-U-T, M-C-T, and D-C-T.

With transaction costs taken into account, each conditional (C) problem has an associated policy function $\alpha(D, \hat{\alpha}, t)$ for $t = 1, \dots, T-1$, while each unconditional (U) problem has an associated policy function $\alpha(\hat{\alpha}, t)$ for $t = 1, \ldots, T-1$.[6] In the no-transaction costs economy, neither α (nor κ in the intermediate consumption case) varies across $\hat{\alpha}$. In other words, $\hat{\alpha}$ is no longer a state variable, and we have $\alpha_t = \alpha$ in the D-U-N case, $\alpha_t = \alpha(D_t, t)$ in the D-C-N case, and $\alpha_t = \alpha(D_t)$ in the M-C-N case, for $t = 1, \ldots, T-1$.[7]

2.3. *Economies and utility comparisons*

Each of the consumer problems described above imply a policy function that, in turn, yields a particular level of expected lifetime utility. Specifically, the policy function $\{\alpha(D_t, \hat{\alpha}_t, t)\}_{t=1}^{T-1}$ can be substituted into the actual law of motion for wealth (Eq. (1)) (with $\{c_t\}_{t=1}^{T-1}$ set to 0) to obtain terminal wealth, W_T. The distribution of terminal wealth is then substituted into Eq. (5) to obtain the consumer's expected lifetime utility. Similarly, when utility depends on intermediate consumption, the policy functions $\{\alpha(D_t, \hat{\alpha}_t, t)\}_{t=1}^{T-1}$ and $\{\kappa(D_t, \hat{\alpha}_t, t)\}_{t=1}^{T-1}$ can be substituted into the actual law of motion for consumer'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. (6) to obtain the consumer's expected lifetime utility.

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, in the utility comparisons we assume that, for a given D_1, $\hat{\alpha}_1$ equals the optimal α_1 for the D-U-N case. Thus, expected lifetime utility only varies with D_1.

In our utility calculations, we are interested in the fraction of wealth that the consumer, who is adopting a sub-optimal policy, would be prepared to give up to be allowed to use the optimal dynamic-conditional (D-C) policy. We consider both the no-transaction-costs and the transaction-costs economies. For the no-transaction costs economy, we calculate the percentage of wealth that a consumer using the dynamic-unconditional (D-U-N), or myopic-conditional (M-C-N) policy functions would sacrifice to use the D-C-N policy function. Similarly, for the transaction costs economy, we calculate the percentage of wealth that a consumer using any of the suboptimal policies, D-U-N-T, M-C-N-T, D-C-N-T, M-U-T, D-U-T, or M-C-T, would give up to use the D-C-T policy.

[6] When utility depends on intermediate consumption, each conditional (C) problem has associated policy functions $\alpha(D_t, \hat{\alpha}_t, t)$ and $\kappa(D_t, \hat{\alpha}_t)$ for $t = 1, \ldots, T-1$, while each unconditional (U) problem has associated policy functions $\alpha(\hat{\alpha}_t, t)$ and $\kappa(\hat{\alpha}_t, t)$ for $t = 1, \ldots, T-1$.

[7] When utility depends on intermediate consumption, we have $\alpha_t = \alpha$ and $\kappa_t = \kappa(t)$ in the D-U-N case, $\alpha_t = \alpha(D_t, t)$ and $\kappa_t = \kappa(D_t, t)$ in the D-C-N case, and $\alpha_t = \alpha(D_t)$ and $\kappa_t = \kappa(D_t, t)$ in the M-C-N case, for $t = 1, \ldots, T-1$.

2.4. Solution technique

The dynamic programming problems are solved by backward recursion. The state variable $\hat{\alpha}$ is discretized and the value function is linearly interpolated between $\hat{\alpha}$ points. This technique yields an approximate solution that converges to the actual solution as the $\hat{\alpha}$ grid becomes finer. In all the optimizations, the holdings of both the risky and the riskless asset are constrained to be non-negative, $0 \leqslant \alpha \leqslant 1$, and to lie on the grid, $\alpha = \{0.00, 0.02, 0.04, \dots, 0.96, 0.98, 1.00\}$. These restrictions imply that the inherited portfolio allocation, $\hat{\alpha}$, also lies between zero and one, $0 \leqslant \hat{\alpha} \leqslant 1$. Consequently, the grid for α is also used as the grid for $\hat{\alpha}$. Constraining α to lie between 0 and 1 is realistic, since individual investors typically face high costs in taking short positions, while institutional investors are often *precluded* by their clients from taking short positions.

3. Calibration

This section describes how we calibrate the evolution of the investment opportunity set to U.S. data. The law of motion is a discrete approximation based on the Gaussian quadrature method developed by Tauchen and Hussey (1991). Several statistics are calculated to assess the closeness of this approximation.

3.1. Law of motion of the investment opportunity set

To estimate a law of motion for the investment opportunity set, we need to identify real-life counterparts for the three variables R, R^{f}, and D. Specifically, we use the monthly rate of return on the value-weighted NYSE index as a proxy for the risky return R, the one-month Treasury-bill rate as a proxy for the risk-free rate R^{f}, and the twelve-month dividend yield on the value-weighted NYSE index as a proxy for the predictive variable D. Both the stock return and interest rate series are deflated using monthly CPI inflation. Also, the dividend yield and stock return series are converted to a continuously compounded basis; hence, R is replaced by $r \equiv \ln(R)$ and the dividend yield D is replaced by $d \equiv \ln(1 + D)$. The stock return, interest rate, and dividend yield series are taken from the Center for Research in Securities Prices (CRSP) and the CPI series is taken from CITIBASE. We use $\mathbf{x}$ to denote the vector $[r, d]'$.

We assume that $\mathbf{x}$ follows the vector autoregressive model (VAR):

$$\mathbf{x}_{t+1} = \mathbf{b} + \mathbf{A}\mathbf{x}_t + \mathbf{e}_{t+1}, \tag{11}$$

where $\mathbf{b}$ is a coefficient vector, $\mathbf{A} = \{a_{i,j}\}$ is a coefficient matrix, and $\mathbf{e} = [e_1, e_2]'$ is a vector of mean-zero, serially uncorrelated, multivariate normal disturbances, with *constant* covariance matrix $\boldsymbol{\Sigma}^{e} = \{\sigma^{e}_{i,j}\}$. $\boldsymbol{\Sigma}^{x} = \{\sigma^{x}_{i,j}\}$ is the unconditional covariance matrix for **x**. We assume that d is the only state variable; i.e., that $a_{i,1}$ equals zero for $i = 1, 2$. This characterization of the investment opportunity set is in line with other papers on optimal portfolio selection (e.g., Barberis, 1996; Campbell and Viceira, 1996). The VAR is estimated using ordinary least squares (OLS).[8]

We discretize the VAR using a variation of the gaussian quadrature method described by Tauchen and Hussey (1991). First, Tauchen and Hussey's method is used to discretize the dividend yield, treating it as a first-order autoregressive process. Second, we exploit the fact that the VAR implies the following expression for stock returns:

$$r_{t+1} = b_1 + a_{1,2}d_t + \rho e_{2,t+1} + u_{t+1}, \tag{12}$$

where ρ is the regression coefficient from regressing e_1 on e_2 and u is i.i.d. normally distributed with variance σ^2_u and is uncorrelated with e_2. The quadrature method is used to calibrate a discrete distribution for the innovation u. We can then calculate a discrete distribution for r_{t+1} for each $\{d_{t+1}, d_t\}$ pair from the discretization of d, since $e_{2,t+1} = d_{t+1} - b_2 - a_{2,2}d_t$.

As to the number of quadrature points, we are constrained by computational considerations. The solution algorithm that we employ is numerically very intensive, so we chose a specification with *nineteen* quadrature points for the dividend yield and *three* points for the innovations in stock returns.[9] Since the pair $\{d_{t+1}, d_t\}$ can take on 19×19 values, stock returns can therefore take $19 \times 19 \times 3 = 1083$ values.

This modification of the Tauchen–Hussey approach has two advantages for our application. First, in contrast to the Tauchen–Hussey approach applied to **x**, the modification ensures that d is sufficient to describe the state of the world at time t. Second, a large number of values can be taken by the stock return variable, which improves accuracy without seriously increasing computation time since the stock return is not a state variable.

Note that the values of the parameters for the stochastic process (Eq. (11)) are taken as known. Thus, we are ignoring the important issue of parameter uncertainty to concentrate on the impact of transaction costs. The effects of parameter uncertainty on optimal portfolio choice have been addressed by Kandel and Stambaugh (1996) and Barberis (1996), among others.

---

[8] Hodrick (1992) argues that in small samples a VAR supplies long-horizon returns statistics that appear to be unbiased. This is especially important for our analysis, since the consumer is long-lived and hence is concerned with the properties of long-horizon asset returns.

[9] We experimented with larger numbers of quadrature points for both d and u, with essentially the same results.

Table 1
Stock returns and dividend yield: data vs. quadrature approximation

We estimate the model

$$r_{t+1} = b_1 + a_{1,2}d_t + e_{1,t+1},$$

$$d_{t+1} = b_2 + a_{2,2}d_t + e_{2,t+1},$$

where d is the log twelve-month dividend yield and r is the continuously-compounded real return, both on the value-weighted NYSE index, $a_{i,2}$ and $b_{i,2}$ ($i = 1, 2$) are coefficients, and e_i ($i = 1, 2$) are mean-zero, serially uncorrelated disturbances. Both r and d are measured in percent. The model is estimated using OLS for the period January 1927 through November 1991. Panel A reports the estimated coefficients and R^2s for the data and contrasts them with those implied by the quadrature approximation. Panel B reports the unconditional volatility of $e_{1,t+1}$ and $e_{2,t+1}$, and the correlation between them, for both the data and the quadrature approximation. Panel C reports the unconditional volatility of r_{t+1} and d_{t+1}, and the correlation between them, for both the data and the quadrature approximation. A full description of the quadrature approximation is contained in Section 3

Panel A: coefficients and R^2's

Equation		b	$a_{\cdot,2}$	R^2
r_{t+1}	Data	− 0.795	0.304	0.361%
	Quad.	− 0.795	0.304	0.362%
d_{t+1}	Data	0.133[a]	0.969[a]	93.8%
	Quad.	0.173	0.959	92.0%

Panel B: unconditional volatility and correlation for e_1 and e_2

		e_1	e_2
e_1	Data	5.66	− 0.907
	Quad.	5.61	− 0.905
e_2	Data	− 0.907	0.317
	Quad.	− 0.905	0.314

Panel C: unconditional volatility and correlation for r and d

		r	d
r	Data	5.67	− 0.159
	Quad.	5.62	− 0.197
d	Data	− 0.159	1.27
	Quad.	− 0.197	1.11

[a]For rows labeled 'Data', denotes significance at the 1% level in a two-sided test.

3.2. A comparison: the data vs. the quadrature approximation

We estimate the model in Eq. (11) and ρ using monthly data from January 1927 through November 1991. Assuming that the degree of return predictability in the U.S. is constant through time, the long time-series should produce precise parameter estimates. The continuously compounded risk-free rate is estimated to be the mean of the continuously compounded one-month Treasury-bill rate over this period, which gives a value for $(R^{\mathrm{f}} - 1)$ of 0.03961%. Panel A of Table 1 reports the estimated coefficients for the data together with the implied coefficients from the quadrature approximation. Panel B reports the unconditional volatility of $e_{1,t+1}$ and $e_{2,t+1}$, and the correlation between them, for both the data and the quadrature approximation. Panel C reports the unconditional volatility of r_{t+1} and d_{t+1}, and the correlation between them, for both the data and the quadrature approximation.

The estimates in Panel A highlight two salient features of the data. First, the dividend yield exhibits a marked and significant persistence. Second, the stock return series is positively, though not significantly, predicted by the dividend yield. These essential features are captured by the quadrature approximation, which reproduces quite well the coefficients of the stock return equation and the persistence in d. Panels B and C show that the quadrature approximation also replicates well the covariance matrices $\boldsymbol{\Sigma}^{\mathrm{e}}$ and $\boldsymbol{\Sigma}^{\mathrm{x}}$. The estimates in Panel B are used to obtain an estimate of -16.18 for ρ.

We also calculate the conditional Sharpe ratio (not reported) for each of the 19 states of the quadrature approximation. The conditional Sharpe measure S in each state (calculated using the continuously compounded rates) is given by

$$S_{it} \equiv \frac{\mathrm{E}[r_{t+1}|d_{it}] - r^{\mathrm{f}}}{\sigma[r_{t+1}|d_{it}]}, \tag{13}$$

where i denotes the state, and σ denotes the volatility of the corresponding variable. This measure describes the investment opportunity set available at time t, given state i, using the conditional distribution of asset returns. We calculate the conditional Sharpe ratio for each of the 19 states of the quadrature approximation using the transition probability matrix to calculate conditional means and variances. We then compare these to the conditional Sharpe ratios calculated at each of the same 19 values of d for the data assuming the VAR specification holds and using the estimated coefficients and innovations covariance matrix. This comparison is important because predictability mainly affects portfolio choices through the conditional Sharpe ratio varying across states. The Sharpe ratios for both the data and the quadrature approximation vary monotonically with d. While the monotone relation for the data follows immediately from the VAR specification and a positive value for $a_{1,2}$, the monotonicity for the approximation is further evidence that the quadrature approximation is capturing the predictability in the data. Indeed, the correlation between the two Sharpe ratios across states is quite high, 0.97.[10]

[10] The correlation is somewhat lower when the first and last state are included. This discrepancy between the data and the quadrature approximation is not a cause of concern because in these two states the asset allocation is likely to be at an endpoint (either zero or one) for both assets.

Since the consumer is long-lived, the predictability of long-horizon equity returns implied by the quadrature approximation relative to that implied by the data is also of interest. One measure of this predictability is the R^2 from a predictive regression of long-horizon stock returns on the previous period's dividend yield, assuming that the VAR in Eq. (11) is correctly specified for the data and the quadrature approximation. Panel A of Fig. 1 graphs the percentage of total variation in long-horizon equity returns that is predictable for the data and the quadrature approximation using the VAR parameters reported in Table 1. The pattern of the R^2 implied by the quadrature approximation essentially replicates that of the data. Since the quadrature approximation tends to understate the R^2, it is likely that our estimates of the utility cost from behaving myopically are a lower bound on the cost of such behavior.

As a second measure of predictability, Panel B reports the residual standard deviation of the predictive regression for different horizons assuming that the VAR is correctly specified. Since the unconditional variance of equity returns may vary with the horizon differently for the data VAR than the quadrature VAR, the residual volatility could behave very differently from the R^2 plotted in Panel A. However, we see that the residual volatility implied by the quadrature VAR replicates very well the pattern of residual volatility for the data VAR. We also report the residual volatility for different horizons assuming that one-period returns are i.i.d. The residual volatility for both the data and the quadrature approximation increases with the horizon at a substantially lower rate than for the i.i.d. case.

Fig. 1 also shows the conditional *long-horizon* Sharpe measure, again assuming the VAR in Eq. (11) is correctly specified. The conditional long-horizon Sharpe measure is the ratio of the conditional mean to the conditional volatility of the τ-period return on a portfolio long \$1 in the risky and short \$1 in the riskless asset:

$$S_{it}^{\tau} \equiv \frac{\mathrm{E}[\sum_{j=1}^{\tau} r_{t+j}|d_{it}] - \tau r^{\mathrm{f}}}{\sigma[\sum_{j=1}^{\tau} r_{t+j}|d_{it}]}. \tag{14}$$

Panel C reports, as a function of horizon, the mean (across states) conditional Sharpe ratio implied by the VAR for the data, by the VAR for the quadrature approximation, and by i.i.d. returns. Again, the quadrature method replicates well the pattern of long-horizon Sharpe ratios estimated from the data. Both the data and the quadrature approximation imply a substantially better trade-off between return and risk than the i.i.d. case.

As a final comparison of the data with the quadrature approximation, Panel D reports the volatility across states of the conditional Sharpe ratios for different horizons. The quadrature approximation would be distorting the variation in the risk-return trade-off across states if it produced a substantially

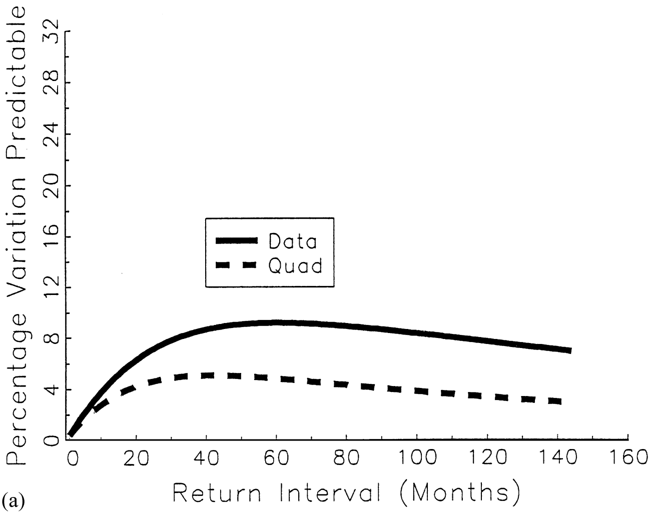

(a)

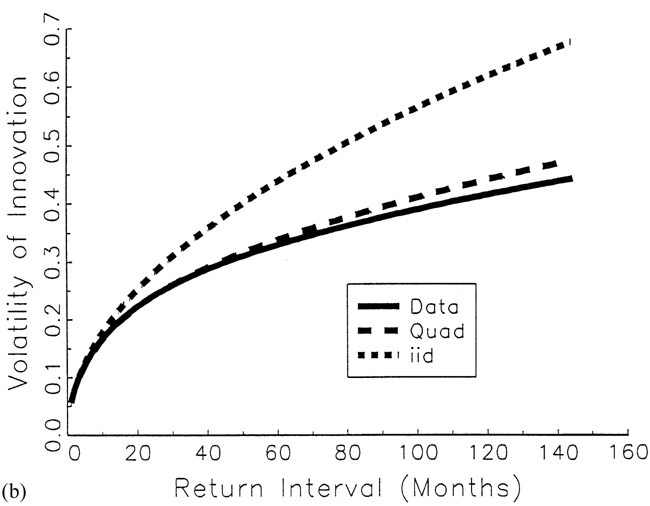

(b)

Fig. 1. Comparison of the properties of the long-horizon risky return and the long-horizon Sharpe ratio. A vector autoregression (VAR) with the lagged dividend yield as the only predictor is used to calculate conditional expectations of long horizon variables for the data, and, for the quadrature approximation of the data. The data, the quadrature approximation, the VAR specification, and the long-horizon Sharpe ratio are all described in Section 3.

different pattern of volatility than the data. However, the patterns exhibited by the data and the quadrature approximation are again similar. In sum, the quadrature approximation appears to capture important features of the data.

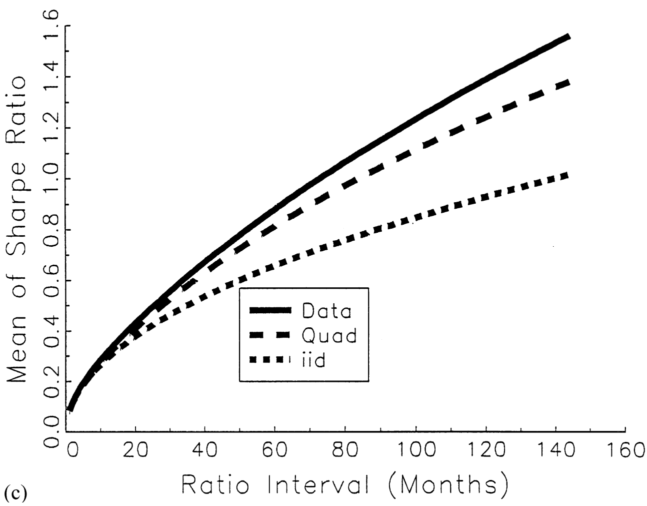

(c)

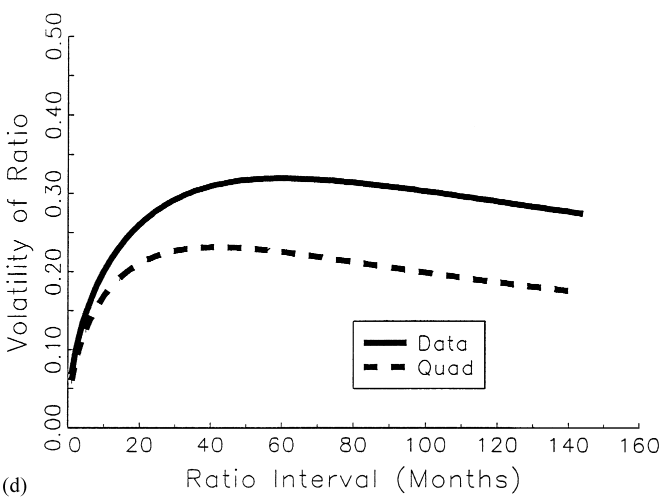

(d)

Fig. 1. Continued.

4. Scenarios and parameter choices

We examine the effects of transaction costs and predictability under several different assumptions about the investment horizon, risk aversion, and transaction costs.

In Scenario 1, which is the base case, the risk aversion parameter, γ equals six, the horizon T is 240, there is no intermediate consumption, the proportional transaction cost parameter, ϕ_1, equals 0.005, and fixed per-dollar transaction costs, ϕ_2, are zero. The other scenarios are variations on the base case. In Scenarios 2 and 3, we set γ equal to ten and two, respectively. In Scenarios 4 and 5, we set T to 360 and 120, respectively. Scenario 6 introduces intermediate consumption. Scenario 7 sets the fixed-cost parameter, ϕ_2, equal to 0.001. In the last three scenarios, the proportional-cost parameter, ϕ_1, is set equal to 0.0001, with the fixed cost parameter varying from zero to 0.0001 to 0.001 across Scenarios 8, 9 and 10, respectively.

In the choice of the relative risk aversion parameter, we follow Mehra and Prescott (1985), who argue that the existing evidence from macro and micro studies constitutes an a priori justification for restricting the value of γ to be a maximum of ten. The base-case investment horizon of 240 months (20 years) is a realistic investment horizon for an investor planning for retirement, for example. It is also the same horizon used in Brennan et al. (1996).

For the choice of the proportional-cost parameter, observe that the proportional component of transaction costs is given by $\phi_1|\alpha - \hat{\alpha}|(W - c)$. The cost of changing the holdings of the risky asset from $\hat{\alpha}W$ to αW, and back, amounts to $2\phi_1|\alpha - \hat{\alpha}|W$. Hence, $2\phi_1$ has the interpretation of the ratio of a bid–asked spread plus commission (plus any possible price impact) over the value of the asset. Bhardwaj and Brooks (1992), Lesmond et al. (1996), and Stoll and Whaley (1983) estimate the round-trip costs for individual stocks listed on the NYSE and the NASDAQ, and the value $2\phi_1 = 0.01$ (100 basis points) in Scenario 1 is realistic for an investor who trades individual stocks directly (rather than futures on the S&P 500). The value $2\phi_1 = 0.0002$ (2 basis points) in Scenarios 8 through 10, on the other hand, is ballpark for an investor who trades futures contracts on the S&P 500 index. For example, Fleming et al. (1995) calculate an average cost of \$3390 to trade 100 S&P 500 futures contracts during March 1991. In that month the index was at roughly 370, which translates into a relative cost of $3390/(100 \times 500 \times 370) = 0.00018$ (recall that each contract pays 500 times the index).

These two values for ϕ_1 should bracket the range of relevant proportional costs facing different investors. In particular, the denomination of the S&P 500 futures contract is such that using the contract to replicate a long position in the underlying securities is only viable for investors with high net wealth. For example, with a grid size of 0.02 for α and given the May 20, 1997 level of the S&P 500 index of about 840, an investor needs a portfolio value of \$21 million to be able to rebalance one grid point using a futures contract. While large institutional investors can use futures contracts to hold equity, we have not incorporated the costs of margin requirements and contract roll-over. Thus, small investors probably face a proportional transaction cost closer to the $\phi_1 = 0.005$ in Scenario 1, while large investors likely face costs greater than the $\phi_1 = 0.0001$ in Scenarios 8 to 10.

In regard to the fixed transaction costs parameter, the lower ϕ_2 value of 0.0001 translates into a fee of \$10 when a consumer reshuffles a \$100,000 portfolio. Viewed as the opportunity cost for an individual investor to process information and instruct a broker to change the portfolio composition, this value seems conservative. Thus, the actual fixed costs faced by an investor probably lie between the high and low values of 0.001 and 0.0001 that we use.

5. Portfolio choices

Fig. 2 illustrates the optimal portfolio choice for the base case, Scenario 1, for the D-C-N and D-C-T policies, for five of the nineteen dividend-yield states. Constantinides (1986) and Davis and Norman (1990) show that optimal portfolio choice in the case without predictability and *proportional* transaction costs involves a no-trade region with boundaries. When the portfolio weight in the risky asset goes outside this region, the CRRA individual trades to return the weight to the closest boundary. Fig. 2 shows that, for the particular set of parameter values considered here, the basic structure of the individual's portfolio choice is unaffected by return predictability.[11] Each state has a no-trade region whose size and position varies across states. These no-trade regions are represented here by the portions of the $\alpha(\hat{\alpha},\ .)$ policy function with a slope of one. While there is good intuition for the state-dependent no-trade regions that we report, our paper is the first to characterize the optimal portfolio choice in the presence of proportional transaction costs and return predictability.

Fig. 3 illustrates optimal portfolio choices in the presence of both *fixed* and *proportional* transaction costs (Scenario 7). There are two characteristics of the policy function in this Scenario. First, there is a state-dependent no-trade region, as in Scenario 1, and, second, when portfolio rebalancing does take place, α is brought to the *interior* of the no-trade region, not just to the boundary. In other words, since the fixed component of transaction costs does not increase with the size of the trade, a consumer who finds it optimal to rebalance does so by an amount that makes another portfolio change in the near future unlikely.

Table 2 reports statistics for portfolio choices in Scenario 1 for the seven sets of policies: D-U-N, M-C-N, D-C-N, M-U-T, D-U-T, M-C-T and D-C-T. Specifically, Panel A reports means and standard deviations (calculated using the unconditional distribution) for the *midpoint* of the no-trade region for the T (transaction-cost) policies, and the optimal portfolio choice for the N (no-transaction-cost) policies. Panel B reports means and standard deviations for the size of the no-trade region for the four T policies. Panel C reports the optimal portfolio choice for the two C-N policies, and the *midpoint* and size of the no-trade region for the two C-T policies, for the five states in Figs. 2 and 3.

Several features of the portfolio choice rules are worth noting. The portfolio choices for the N policies are qualitatively similar to those described in earlier

[11] Since we have only solved the consumer's problem numerically, the form of the consumer's portfolio choice that we find may be specific to our calibration of U.S. equity return.

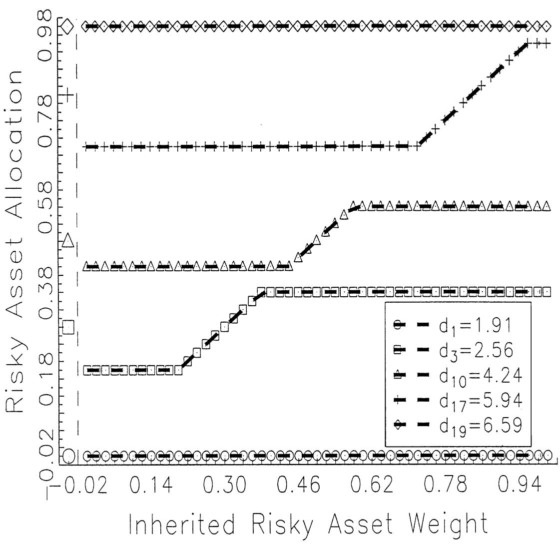

Fig. 2. Portfolio allocation to the risky asset in Scenario 1. The consumer lives for 240 months, possesses a risk aversion coefficient of six, consumes only at a final date, and faces a proportional transaction cost equal to 0.005, and no fixed cost. The time-1 policy for a dynamic-conditional consumer aware of the transaction costs (D-C-T) is shown to the right of the dashed line for five states where *d* denotes the log dividend yield and is measured in percent. The optimal portfolio allocation is a function of the inherited portfolio allocation. The left of the dashed line shows the optimal time-1 allocation for a D-U individual unaware of transaction costs. Details of the consumer's problem are contained in Section 2, while the return calibration is described in Section 3.

work. Specifically, a comparison of the M-C-N and D-C-N policies reveals that the allocation to the risky asset is on average higher with dynamic policies than with myopic policies (see Panel A). This result is consistent with the findings of Barberis (1996), Brennan et al. (1996), and Campbell and Viceira (1996) for risk aversion coefficients greater than one. The intuition for this result, described by Campbell and Viceira (1996), is that the dividend yield predicts future stock returns with a positive slope coefficient, while the innovations in the dividend yield and current stock return are negatively correlated. Therefore, expected future stock returns are low when current stock return is high, and vice versa. This implies that stocks provide a hedge against changes in the investment opportunity set. A dynamic investor with a risk aversion parameter greater than one appreciates the 'hedging virtues' of stocks, and allocates more to stocks than a myopic investor with the same risk aversion.

Our results are qualitatively similar to earlier papers in other respects. For example, Campbell and Viceira find that for power utility individuals, hedging demand (dynamic less myopic portfolio allocation) as a fraction of the dynamic portfolio allocation is increasing in risk aversion, while the slope of the dynamic

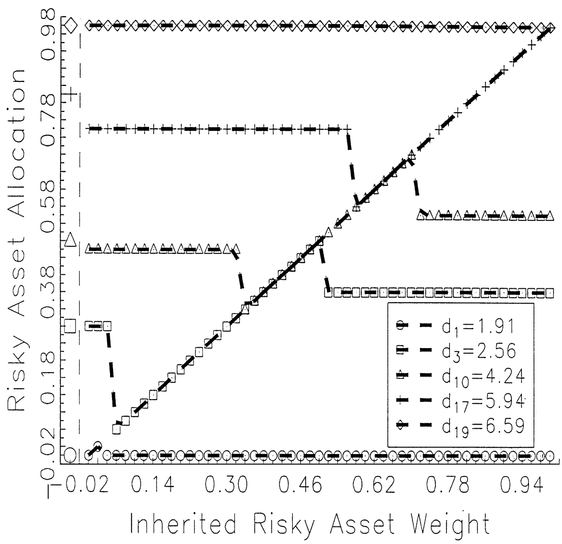

Fig. 3. Portfolio allocation to the risky asset in Scenario 7. The consumer lives for 240 months, possesses a risk aversion coefficient of six, consumes only at a final date, and faces a proportional transaction cost equal to 0.005, and a fixed cost of 0.001. The time-1 policy for a dynamic-conditional consumer aware of the transaction costs (D-C-T) is shown to the right of the dashed line for five states where *d* denotes the log dividend yield and is measured in percent. The optimal portfolio allocation is a function of the inherited portfolio allocation. The left of the dashed line shows the optimal time-1 allocation for a D-U individual unaware of transaction costs. Details of the consumer's problem are contained in Section 2, while the return calibration is described in Section 3.

portfolio allocation as a function of dividend yield is declining in risk aversion. Comparing the portfolio allocations (not reported) for Scenario 1 ($\gamma = 6$) to those for Scenario 2 ($\gamma = 10$), these results hold for our calibration of U.S. returns as well. Barberis (1996) finds, for $\gamma > 1$, that hedging demand is increasing in the dividend yield, and our results confirm this finding (see, for example, Table 2, Panel C).

Table 2 also shows that conditional myopic portfolio choices are on average similar to unconditional ones, but the conditional choices vary substantially across states (see Panel A). In fact, the standard deviation of the conditional portfolio choice across states can be up to 19.8%. Hence, as documented by Kandel and Stambaugh (1996), even little predictability in stock returns (the R^2 of the stock-return regression implied by the quadrature approximation is only 0.36%) significantly affects the consumer's portfolio decisions.

Panel B shows that the average size of the no-trade region for T policies is larger for myopic than for dynamic policies. This result is robust to whether the conditional (C) or unconditional (U) transition probability matrix is used by the

Table 2
Portfolio choices: Scenario 1

Transaction costs in Scenario 1 are proportional, giving rise to portfolio rebalancing rules with no-trade regions for the portfolio allocation to the risky asset. Panel A reports the mean and standard deviation (SD) across states of the midpoint of the no-trade region for the T policies, which account for transaction costs, and the portfolio choice for the N policies, which ignore transaction costs. Panel B reports the mean and standard deviation across states of the size of the no-trade region for the T policies. Means and standard deviations are calculated using the unconditional distribution. Panel C reports the midpoints of the no-trade region and the size of the no-trade region for five of the nineteen states, where d denotes the log dividend yield and is measured in percent. The table considers several consumer policies: the *myopic* (M) and *dynamic* (D) policies, the *unconditional* (U) and *conditional* (C) policies, and the *no-transaction-costs* (N) and *transaction-costs* (T) policies. We combine the three distinctions to obtain *seven* policies: D-U-N, M-C-N, D-C-N, M-U-T, D-U-T, M-C-T, and D-C-T. Panel C does not report information for the policies D-U-N, M-U-T and D-U-T, because they do not display any variation across states. Descriptions of the seven consumer problems are contained in Section 2

Panel A: *midpoint of the no-trade region*

	D-U-N	M-C-N	D-C-N	M-U-T	D-U-T	M-C-T	D-C-T
Mean	0.320	0.329	0.525	0.330	0.330	0.339	0.528
SD	0.000	0.148	0.198	0.000	0.000	0.140	0.190

Panel B: *size of no-trade region*

	M-U-T	D-U-T	M-C-T	D-C-T
Mean	0.540	0.060	0.489	0.159
SD	0.000	0.000	0.108	0.048

Panel C: *midpoint and size of no-trade region*

State	d	M-C-N	D-C-N	M-C-T	D-C-T	M-C-T	D-C-T
		Midpoints				No-trade region	
1	1.91	0.000	0.000	0.000	0.000	0.000	0.000
3	2.56	0.180	0.300	0.220	0.290	0.440	0.180
10	4.25	0.320	0.500	0.320	0.510	0.520	0.140
17	5.94	0.460	0.840	0.460	0.840	0.520	0.240
19	6.59	1.000	1.000	1.000	1.000	0.000	0.000

consumer. Balduzzi and Lynch (1997) find a similar result for the constant opportunity set case. The intuition for this result is described by Balduzzi and Lynch. Briefly, a myopic investor ignores the future gains from changing the current portfolio composition, and hence is reluctant to pay transaction fees. Consequently, the no-trade region is relatively wide. A dynamic investor, on the other hand, cares about the future gains from changing the current composition and so is more inclined to trade and incur fees. This reasoning does not depend on either predictability or hedging demands by the consumer, which explains why the no-trade region is larger for myopic than dynamic policies in the U as well as the C case.[12] Note that the average size of the no-trade region for T policies is substantial. The average no-trade region ranges from 0.540 for the M-U-T policy to 0.060 for the D-U-T policy.

Panel C shows that, even in the presence of hedging demands and transaction costs, there is a close positive correlation across states between the value of the $\alpha(., \hat{\alpha}, t)$ policy function for conditional investors and the Sharpe measure. In addition, the size of the D-C-T no-trade region varies substantially with the dividend yield.

To illustrate how portfolio choices differ across policies, Fig. 4 plots the portfolio choices that would have been made using the M-C-T, D-C-T and D-C-N policies from Scenario 1 over the sample period 1/27-11/91. The policies are those calculated as described in Sections 3 and 4 and summarized in Table 2. The real riskless rate is taken to be the sample average of 0.03961% and the real risky asset return is that of the value-weighted NYSE index. The inherited risky asset weight at the start of 1/27 is the midpoint of the no-trade region for the relevant policy given the d state at that time. The DCT and DCN rebalancing rules are implemented assuming a long-lived investor, which means that the $t = 1$ rule is used over the entire period. Each month, the calibration state whose dividend yield is closest to the dividend yield on the value-weighted NYSE is taken to be the prevailing state. The circles in Fig. 4 indicate when the investor is actually rebalancing.

We offer several observations about Fig. 4. First, the three conditional policies imply substantial variation in α over the sample period, with values ranging from 0.3 to one. Second, the range and pattern of variation is consistent with results in Campbell and Viceira (1996). For quarterly data and a relative risk aversion coefficient of four, they find that the risky asset allocation ranges from about 0.5 to five over a similar period (1947–1995), with little time being spent below one. We use a higher risk aversion coefficient and find that, while the pattern of the D-C-N allocation is similar to theirs, it is always substantially lower, and generally less than one. Third, the D-C-T allocation differs considerably from the D-C-N allocation over the period, even though the average D-C-N allocation and the average midpoint of the D-C-T allocation are similar (see Table 2, Panel A). Fourth, the D-C-T policy exhibits more variation in α and more frequent rebalancing than the M-C-T policy. These two results are consistent with the no-trade region being narrower for D-C-T policy than the M-C-T

[12] Since the result holds in the U case, it implies that transaction costs cause the portfolio rebalancing rule to change over the lifecycle even in the absence of predictability. Thus, Samuelson's irrelevance result (that portfolio choice by a CRRA investor is invariant over the lifecycle when returns are i.i.d.) is not robust to transaction costs. See Balduzzi and Lynch (1997) for further details.

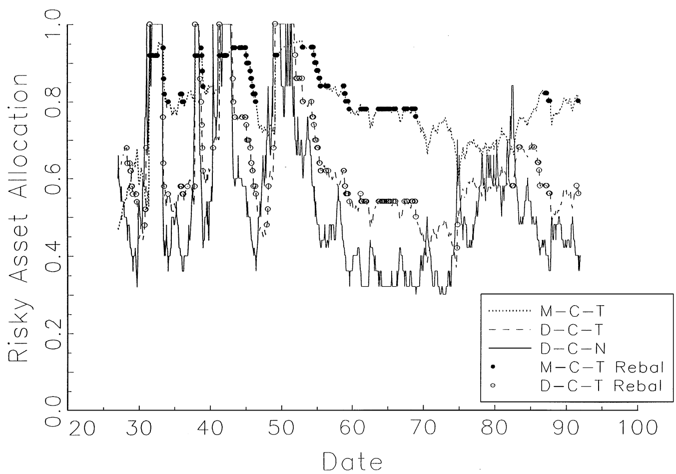

Fig. 4. Portfolio allocation over the period January 1927 through November 1991 by the consumer of Scenario 1. The consumer lives for 240 months, possesses a risk aversion coefficient of six, consumes only at a final date, and faces a proportional transaction cost equal to 0.005, and no fixed cost. The portfolio allocation for a myopic-conditional consumer aware of the transaction costs (M-C-T) is shown together with the allocations made by a dynamic-conditional consumer. The dynamic-conditional consumer is either aware of the transaction costs (D-C-T), or ignores them (D-C-N). Details of the consumer's problem are discussed in Section 2, while the calibration is described in Section 3. The portfolio rebalancing rules are applied given the sequence of realized real returns on stocks over the period and assuming a real riskless rate of 0.03961% as in the calibration. The inherited risky asset weight at the start of January 1927 is the midpoint of the trade region given the calibration state prevailing at that time. The last two rebalancing rules (D-C-T and D-C-N) are implemented for a long-lived investor, which means that the time-1 rule is used over the entire period. Circles indicate when the consumer is actually rebalancing.

policy, together with the state changing as the investor moves through the sample period. Finally, the D-C-N policy results in the most volatile α over the sample period.

6. Utility costs

Table 3 reports utility cost calculations for each of the ten scenarios. Calculations for the analogous no-transaction case are also reported for each scenario. Since the utility cost depends on the state at time 1, the table reports both the mean and the standard deviation of the utility cost using the unconditional distribution.

We have three possibilities for each scenario. First, transaction costs exist and are taken into account by the consumer, which corresponds to four policies,

Table 3
Utility comparisons

The table reports the mean and (underneath) the standard deviation of the percentage of wealth that a consumer, who is possibly only *partially* informed about the economy, would sacrifice to become *fully* informed. The individual consumes only at a final date (all Scenarios but 6) or as time goes by (Scenario 6). For each of the ten scenarios considered, ϕ_1 is the proportional-cost parameter, and ϕ_2 is the fixed-cost parameter. The individual lives for T months and has a relative-risk-aversion coefficient of γ. The table considers several consumer policies: the *myopic* (M) and *dynamic* (D) policies, the *unconditional* (U) and *conditional* (C) policies, and the *no-transaction-costs* (N) and *transaction-costs* (T) policies. We combine the three distinctions to obtain *seven* policies: D-U-N, M-C-N, D-C-N, M-U-T, D-U-T, M-C-T, and D-C-T. This table considers the *no-transaction-costs* (N) economy, and compares the performance of D-U-N and M-C-N, to that of D-C-N. This table also considers the *transaction-costs* (T) economy and compares the performance of D-U-N, M-C-N, D-C-N, M-U-T, D-U-T, and M-C-T, to D-C-T. For each of the 19 possible states, the initial value of the inherited portfolio allocation is set equal to the optimal portfolio allocation in the D-U-N case. The unconditional distribution for the states is used to calculate means and standard deviations. A full description of the seven consumer problems is contained in Section 2

	D-U-N	M-C-N	D-C-N	M-U-T	D-U-T	M-C-T	D-C-T
Scenario 1: $\gamma = 6$, $T = 240$, $\phi_1 = 0.005$, $\phi_2 = 0$							
N:	23.10	8.76	0.00	—	—	—	—
	2.19	0.62	0.00	—	—	—	—
T:	21.21	10.18	3.01	22.94	20.78	10.00	0.00
	2.15	0.69	0.02	1.52	2.12	1.33	0.00
Scenario 2: $\gamma = 10$, $T = 240$, $\phi_1 = 0.005$, $\phi_2 = 0$							
N:	22.40	11.10	0.00	—	—	—	—
	2.31	0.90	0.00	—	—	—	—
T:	19.77	10.67	2.43	21.93	19.82	11.27	0.00
	2.17	0.87	0.03	1.73	2.17	0.67	0.00
Scenario 3: $\gamma = 2$, $T = 240$, $\phi_1 = 0.005$, $\phi_2 = 0$							
N:	7.37	0.40	0.00	—	—	—	—
	1.62	0.04	0.00	—	—	—	—
T:	5.96	3.29	1.92	6.02	5.49	6.56	0.00
	1.31	0.31	0.27	1.36	1.35	2.37	0.00
Scenario 4: $\gamma = 6$, $T = 360$, $\phi_1 = 0.005$, $\phi_2 = 0$							
N:	33.22	13.55	0.00	—	—	—	—
	1.90	0.59	0.00	—	—	—	—
T:	30.79	15.61	4.46	32.87	30.20	15.27	0.00
	1.89	0.65	0.02	1.12	1.86	1.43	0.00
Scenario 5: $\gamma = 6$, $T = 120$, $\phi_1 = 0.005$, $\phi_2 = 0$							
N:	11.46	3.72	0.00	—	—	—	—
	2.49	0.64	0.00	—	—	—	—
T:	10.31	4.43	1.53	11.43	10.12	4.49	0.00
	2.40	0.71	0.02	2.04	2.36	0.74	0.00

Table 3. Continued.

	D-U-N	M-C-N	D-C-N	M-U-T	D-U-T	M-C-T	D-C-T
Scenario 6: $\gamma = 6$, $T = 240$, $\phi_1 = 0.005$, $\phi_2 = 0$, $c_t \neq 0$, $t < T$							
N:	10.97	3.29	0.00	—	—	—	—
	2.19	0.48	0.00	—	—	—	—
T:	9.97	4.02	1.39	10.94	9.79	3.89	0.00
	2.10	0.53	0.02	1.80	2.06	0.64	0.00
Scenario 7: $\gamma = 6$, $T = 240$, $\phi_1 = 0.005$, $\phi_2 = 0.001$							
N:	23.10	8.76	0.00	—	—	—	—
	2.19	0.62	0.00	—	—	—	—
T:	27.46	22.65	15.95	21.22	19.79	14.85	0.00
	1.93	0.56	0.20	1.53	2.09	4.12	0.00
Scenario 8: $\gamma = 6$, $T = 240$,$\phi_1 = 0.0001$, $\phi_2 = 0$							
N:	23.10	8.76	0.00	—	—	—	—
	2.19	0.62	0.00	—	—	—	—
T:	23.02	8.74	0.01	22.71	23.02	8.80	0.00
	2.19	0.62	0.00	2.15	2.19	0.65	0.000
Scenario 9: $\gamma = 6$, $T = 240$, $\phi_1 = 0.0001$, $\phi_2 = 0.0001$							
N:	23.10	8.76	0.00	—	—	—	—
	2.19	0.62	0.00	—	—	—	—
T:	23.16	9.50	0.78	23.80	22.92	8.83	0.00
	2.18	0.62	0.02	2.10	2.19	0.58	0.00
Scenario 10: $\gamma = 6$, $T = 240$, $\phi_1 = 0.0001$, $\phi_2 = 0.001$							
N:	23.10	8.76	0.00	—	—	—	—
	2.19	0.62	0.00	—	—	—	—
T:	28.27	20.46	12.30	22.84	21.47	10.19	0.00
	1.99	0.50	0.22	1.58	2.14	0.39	0.00

M-U-T, M-C-T, D-U-T and D-C-T. The utility costs for these policies are reported in the T row and the appropriate column. Second, transaction costs are zero, and so the M-U-N and D-U-N policies are the same. Thus, this case has three policies, D-U-N, M-C-N, and D-C-N, whose utility costs are reported in the N row.[13] Third, transaction costs exist but are ignored by the consumer

[13] Note that the only difference between Scenario 1 and Scenarios 7 through 10 is the magnitude of the transaction costs. Thus, the N rows for these four scenarios are the same.

when making portfolio choices, which corresponds to three policies, D-U-N-T, M-C-N-T, and D-C-N-T. Their utility costs are reported in the first three columns of the T row.[14]

6.1. The cost of being a myopic optimizer

In general, the utility cost of being a myopic, rather than a dynamic, optimizer can be substantial. In Scenario 1, behaving myopically using the conditional distribution has a utility cost of 10.00%. This utility cost can be as high as 15.27% for Scenario 4 (360 month horizon).

Also, a comparison of the M-C-N cost in the N row with the M-C-T cost in the T row shows that utility cost of being myopic is *always higher* in the presence of transaction costs. The intuition for this result is as follows. In the absence of transaction costs, the myopic optimizer is ignoring hedging demand. However, with transaction costs, this optimizer is ignoring both hedging demand and the impact of future transaction costs on portfolio choice today. As discussed above, future transaction costs narrow the no-trade region for the dynamic optimizer relative to the myopic optimizer. The failure by the myopic optimizer to narrow the no-trade region has an incremental effect over and above that of ignoring the hedging demand. This incremental effect causes a higher utility cost for myopic behavior when there are transaction costs.

6.2. The cost of ignoring predictability

The relative cost of ignoring the conditional distribution of asset returns (while still dynamically optimizing) can be quite high, both with and without transaction costs. In Scenario 1, the cost exceeds 23% for the D-U-N case and 20% for the D-U-T case. This high cost of ignoring the weak stock-return predictability in the data is consistent with the results of Kandel and Stambaugh for a one-period investor who uses available data and a diffuse prior. Kandel and Stambaugh find that the investor's ability to time the market can be quite valuable.

In addition, transaction costs affect the utility cost of ignoring predictability. For a dynamic optimizer, the utility cost of ignoring predictability in the absence of transaction costs (D-U-N policy in the N economy) can be compared to this same utility cost with transaction costs which are properly accounted for (D-U-T policy). This utility cost is *always* higher in the absence of transaction costs and the reason is as follows. Transaction costs cause the consumer to rebalance less frequently. Less frequent rebalancing means that the utility impact of using the unconditional rather than the conditional distribution is reduced.

[14] Recall that, in this third case, the consumer who cares about intermediate consumption (Scenario 6) takes the impact of transaction costs into account when making consumption choices.

6.3. The cost of ignoring transaction costs

The cost of ignoring realistically-small transaction costs when making portfolio choices can be quite substantial. When the proportional transaction cost is set at a value realistic for an investor trading in individual stocks ($\phi_1 = 0.005$, Scenarios 1–6), the utility cost of ignoring transaction costs ranges from a minimum of 1.39% (Scenario 6, intermediate consumption) to a maximum of 4.46% (Scenario 4, 360 month lifetime horizon). When a large fixed cost is added to this proportional cost, as in Scenario 7, a dynamic optimizer employing conditional policies is subject to a utility cost in excess of 15% by ignoring transaction costs. This result complements our previous finding that these realistically small transaction costs lead to wide no-trade regions.

When the proportional cost is set equal to the low value of 0.0001 and the fixed cost of trading is zero (Scenario 8), the utility cost associated with ignoring transaction costs is negligible. However, once realistic fixed costs are introduced, the utility cost of ignoring transaction costs becomes quite substantial. In particular, utility costs range from 0.78% when the fixed cost is 0.0001 (Scenario 9) up to 12.30% when the fixed cost is 0.001 (Scenario 10). Thus, even for those wealthy individuals rebalancing using futures contracts, the presence of fixed costs in addition to the proportional cost means that ignoring transaction costs can have a significant adverse effect on utility.

6.4. Changing the parameters

6.4.1. Changes in relative risk aversion (γ)

An increase in γ from six to ten reduces the utility cost of adopting unconditional policies both with and without transaction costs. The rationale for this result is that a high γ tends to shrink risky-asset portfolio allocations towards zero, which reduces the difference between conditional and unconditional policies. In worlds with and without transaction costs, the utility cost of behaving myopically rather than dynamically increases as γ increases from six to ten. This finding is consistent with hedging demand, as a percentage of dynamic demand, increasing as γ becomes larger (see the discussion in Section 5 above).

When γ decreases from six to two, the arguments are more complicated. With low risk aversion, the restriction on short-sales of the riskless asset is binding in a large number of states. This has two effects on the utility cost of behaving myopically and also on the utility cost of using the unconditional distribution. First, the binding constraint lowers the utility associated with the D-C policy. Consequently, the utility cost of behaving myopically and of using the unconditional distribution increases, since that cost depends on the percentage change in utility. In contrast, the second effect reduces the utility cost of behaving myopically and of using the unconditional distribution. If both the D-C and M-C optimal αs are one in a state, then the utility loss associated with behaving myopically in that state is zero. Similarly, if the D-C αs are hitting the riskless asset short-sale constraint in a number of states, the utility gain from using the conditional distribution is reduced. Since the utility cost of the M-C and D-U policies is lower for $\gamma = 2$ than $\gamma = 6$, in both the T and N-T economies, the second effect dominates for both myopic policies and for policies using the unconditional distribution.

6.4.2. Changes in the lifetime horizon T

An increase in T increases *all* utility costs. In fact, any inefficiency in the policies adopted by the consumer is amplified as the lifetime horizon increases: a higher number of inefficient decisions are taken.

6.4.3. Intermediate consumption

Introducing intermediate consumption decreases all utility costs. This finding suggests that inefficient portfolio policies have a greater adverse effect on investors only concerned with a terminal wealth (such as institutional investors managing assets for the long run) than those consuming as time goes by.

6.4.4. Changes in transaction costs (ϕ_1, ϕ_2)

As proportional transaction costs increase (going from Scenario 8 to 1), the cost of adopting the dynamic unconditional policy decreases. As mentioned before, unconditional policies induce less variability in portfolio choices. Hence, when proportional transaction costs increase and the individual trades less aggressively, the disadvantage of using unconditional, rather than conditional, policies is reduced. Similarly, the cost of adopting dynamic unconditional policies decreases as fixed transaction costs increase and the individual trades less frequently (going from Scenario 1 to 7, from Scenario 8 to 9, and from Scenario 9 to 10).

On the other hand, higher transaction costs always increase the cost of adopting inefficient conditional policies (M-C-N-T, D-C-N-T and M-C-T). In the case where the individual behaves myopically, larger transaction costs just make the no-trade region for M-C-T even bigger relative to that for D-C-T, causing the incremental utility cost associated with an excessively-wide no-trade region to increase. In the cases where the consumer ignores transaction costs (M-C-N-T and D-C-N-T), the variability in portfolio choice associated with conditional policies means that excessive rebalancing occurs. When the actual transaction cost parameters are larger, this excessive rebalancing leads to higher transaction costs being paid and greater utility losses.

7. Conclusions

This paper solves the optimal consumption–investment problem of a finitely-lived agent facing asset-return predictability and transaction costs, who possibly ignores some aspects of the economy. We find that the cost of behaving myopically, of ignoring predictability, or of ignoring realistic transaction costs can be substantial. These costs are reduced if the investor is allowed to consume as time goes by. The presence of transaction costs always increases the utility cost of behaving myopically, but reduces the large utility cost associated with ignoring predictability. Regarding portfolio-choice policies, proportional transaction costs induce a no-trade region such that the consumer rebalances to the nearer boundary when the risky-asset weight goes outside the region. Fixed and proportional transaction costs together also induce a no-trade region, but now rebalancing places the risky-asset portfolio weight inside the boundary. Also, realistically-small transaction costs induce sizable no-trade regions.

The paper also makes methodological contributions. We show how the investor's multiperiod problem can be numerically solved, in the presence of time-varying conditional expected returns, and proportional and fixed transaction costs. This methodology could be used to calculate the utility cost associated with using rules of thumb rather than the optimal rebalancing policy. For example, a rule that involves rebalancing at fixed intervals can be compared with the optimal rule in the presence of transaction costs.

This paper only considers time variation in the first moments of asset returns. A natural extension of our analysis would be to consider the impact of time variation in the conditional second moments as well. Such an extension is left for future research.

Appendix A

Let $\kappa_t = \kappa(D_t, \hat{\alpha}_t, t)$ denote the consumer's choice of consumption as a fraction of wealth at time t, where the policy is *time dependent* since the assumed time horizon T is finite.

A.1. Dynamic-Conditional Policy (D-C)

When the consumer consumes over time, her Bellman equation in the dynamic conditional case is given by

$$\frac{a(D_t, \hat{\alpha}_t, t)W_t^{1-\gamma}}{1-\gamma} = \max_{\kappa_t, \alpha_t} \frac{\kappa_t^{1-\gamma}W_t^{1-\gamma}}{1-\gamma} + \frac{(1-\kappa_t)^{1-\gamma}W_t^{1-\gamma}}{1-\gamma}\mathrm{E}[a(D_{t+1}, \hat{\alpha}_{t+1}, t+1)R_{W,t+1}^{1-\gamma}|D_t, \hat{\alpha}_t]$$

$$\text{for } t = 1, \ldots, T-1, \qquad \text{(A.1)}$$

where $\mathrm{E}[.|D, \hat{\alpha}]$ denotes the expectation taken using the conditional distribution given D. Note that the Bellman equation (A.1) is solved by *backward iteration*, starting with $t = T-1$ and $a(D, \hat{\alpha}, T) = 1$.

A.2. Myopic-Conditional Policy (M-C)

In the myopic case, the consumer's portfolio composition problem is separable from her consumption decision. Given the optimal portfolio decision, $\alpha(D,\hat{\alpha})$, the consumer takes into account the current state D and the number of periods she has left to live in determining her consumption as a fraction of wealth. Formally, the optimal consumption problem is

$$\begin{aligned}\frac{a(D_t, \hat{\alpha}_t, t)W_t^{1-\gamma}}{1-\gamma} = \max_{\kappa_t} & \frac{\kappa_t^{1-\gamma}W_t^{1-\gamma}}{1-\gamma} \\ & + \frac{(1-\kappa_t)^{1-\gamma}W_t^{1-\gamma}}{1-\gamma}\mathrm{E}[a(D_{t+1}, \hat{\alpha}_{t+1}, t+1)R_{W,t+1}^{1-\gamma}|D_t, \hat{\alpha}_t] \\ & \text{for } t = 1, \dots, T-1, \end{aligned} \tag{A.2}$$

which is also solved by backward iteration, starting with $t = T-1$ and $a(D, \hat{\alpha}, T) = 1$. Note that, while the consumer is myopic in her portfolio choice, she is a dynamic optimizer when it comes to consumption decisions. Hence, in the intermediate consumption scenario, the utility loss due to adopting a myopic policy can be entirely attributed to the less-than-fully-optimal portfolio choice.

A.3. Dynamic-Unconditional Policy (D-U)

The Bellman equation faced by the consumer is given by

$$\begin{aligned}\frac{a(\hat{\alpha}_t, t)W_t^{1-\gamma}}{1-\gamma} = \max_{\kappa_t, \alpha_t} & \frac{\kappa_t^{1-\gamma}W_t^{1-\gamma}}{1-\gamma} \\ & + \frac{(1-\kappa_t)^{1-\gamma}W_t^{1-\gamma}}{1-\gamma}\mathrm{E}^U[a(\hat{\alpha}_{t+1}, t+1)R_{W,t+1}^{1-\gamma}|\hat{\alpha}_t] \\ & \text{for } t = 1, \dots, T-1. \end{aligned} \tag{A.3}$$

A.4. Myopic-Unconditional (M-U)

The consumer's portfolio composition problem is again separable from her consumption decision. Given $\alpha(\hat{\alpha})$, consumption is determined based on

$$\begin{aligned}\frac{a(\hat{\alpha}_t, t)W_t^{1-\gamma}}{1-\gamma} = \max_{\kappa_t} & \frac{\kappa_t^{1-\gamma}W_t^{1-\gamma}}{1-\gamma} \\ & + \frac{(1-\kappa_t)^{1-\gamma}W_t^{1-\gamma}}{1-\gamma}\mathrm{E}^U[a(\hat{\alpha}_{t+1}, t+1)R_{W,t+1}^{1-\gamma}|\hat{\alpha}_t] \\ & \text{for } t = 1, \dots, T-1. \end{aligned} \tag{A.4}$$

References

Balduzzi, P., Lynch, A., 1997. Samuelson's irrelevance result: its breakdown in the presence of transaction costs. Unpublished working paper. New York University, New York, NY.

Barberis, N., 1996. Investing for the long run when returns are predictable. Unpublished working paper. Harvard University, Cambridge, MA.

Bhardwaj, R., Brooks, L., 1992. The january anomaly: effects of low share price, transaction costs and bid–ask bias. Journal of Finance 47, 553–574.

Brennan, M., Schwartz, E., 1996. The use of treasury bill futures in strategic asset allocation programs. Unpublished working paper. UCLA, Los Angeles, CA.

Brennan, M., Schwartz, E., Lagnado, R., 1996. Strategic asset allocation. Journal of Economic Dynamics and Control, forthcoming.

Campbell, J., 1987. Stock returns and the term structure. Journal of Financial Economics 18, 373–399.

Campbell, J., Viceira, L., 1996. Consumption and portfolio decisions when expected returns are time varying. Unpublished working paper. Harvard University, Cambridge, MA.

Constantinides, G., 1986. Capital market equilibrium with transaction costs. Journal of Political Economy 94, 842–862.

Davis, M., Norman, A., 1990. Portfolio selection with transaction costs. Mathematics of Operations Research 15, 676–713.

Duffie, D., Sun, T., 1990. Transaction costs and portfolio choice in a discrete–continuous-time setting. Journal of Economic Dynamics and Control 14, 35–51.

Fama, E., French, K., 1989. Business conditions and expected returns on stocks and bonds. Journal of Financial Economics 25, 23–49.

Fleming, J., Ostdiek, B., Whaley, R., 1995. Trading costs and the relative rates of price discovery in stock, futures, and option markets. Unpublished working paper. Duke University, Durham, NC.

Grossman, S., Shiller, R., 1981. The determinants of the variability of stock market prices. American Economic Review 71, 222–227.

Hansen, L., Heaton, J., Luttmer, E., 1995. Econometric evaluation of asset pricing models. Review of Financial Studies 8, 237–274.

Hansen, L., Singleton, K., 1982. Generalized instrumental variables estimation of nonlinear rational expectations models. Econometrica 50, 1269–1286.

He, H., Modest, D., 1995. Market frictions and consumption-based asset pricing. Journal of Political Economy 103, 94–117.

Heaton, J., Lucas, D., 1996. Evaluating the effects of incomplete markets on risk sharing and asset pricing. Journal of Political Economy 104, 443–487.

Hodrick, R., 1992. Dividend yields and expected stock returns: alternative procedures for inference and measurement. Review of Financial Studies 5, 357–386.

Kandel, S., Stambaugh, R., 1996. On the predictability of stock returns: an asset allocation perspective. Journal of Finance 51, 385–424.

Kim, T., Omberg, E., 1996. Dynamic nonmyopic portfolio behavior. Review of Financial Studies 9, 141–161.

Koo, H., 1991. Portfolio choice with linear transaction costs: theory, numerical solutions and applications. Unpublished working paper. Princeton University, Princeton, NJ.

Lesmond, D., Trzcinka, T., Ogden, J., 1996. A new measure of total transaction costs. Unpublished working paper. State University of New York at Buffalo, Buffalo, NY.

Luttmer, E., 1996. Asset pricing in economies with frictions. Econometrica 64, 1439–1467.

Mehra, R., Prescott, E., 1985. The equity premium: a puzzle. Journal of Monetary Economics 15, 145–161.

Morton, A., Pliska, S., 1993. Optimal portfolio management with fixed transaction costs. Unpublished working paper. University of Illinois at Chicago, Chicago, IL.

Samuelson, P., 1969. Lifetime portfolio selection by dynamic stochastic programming. The Review of Economics and Statistics 51, 239–246.
Schroder, M., 1995. Optimal portfolio selection with fixed transaction costs. Unpublished working paper. Northwestern University, Evanston, IL.
Stoll, H., Whaley, R., 1983. Transaction costs and the small firm effect. Journal of Financial Economics 12, 57–79.
Tauchen, G., Hussey, R., 1991. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. Econometrica 59, 317–396.
Vayanos, D., 1996. Transaction costs and asset prices: a dynamic equilibrium model. Unpublished working paper. Stanford University, Stanford, CA.