



American Finance Association

Predictability and Transaction Costs: The Impact on Rebalancing Rules and Behavior

Author(s): Anthony W. Lynch and Pierluigi Balduzzi

Source: *The Journal of Finance*, Vol. 55, No. 5 (Oct., 2000), pp. 2285-2309

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/222490>

Accessed: 24/11/2009 12:14

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

Predictability and Transaction Costs: The Impact on Rebalancing Rules and Behavior

ANTHONY W. LYNCH and PIERLUIGI BALDUZZI*

ABSTRACT

Recent papers show that predictability calibrated to U.S. data has a large effect on the rebalancing behavior of a multiperiod investor. We find that this continues to be true in the presence of realistic transaction costs. In particular, predictability causes the no-trade region for the risky-asset holding to become state dependent and, on average, wider and higher. Predictability also motivates the investor to spend considerably more on rebalancing and to rebalance more often. In other results, we find that introducing costly liquidation of the risky asset for consumption lowers the average allocation to the risky asset, though only marginally early in life. Our experiments also vary the nature of the return predictability and introduce return heteroskedasticity.

WE EXAMINE THE PORTFOLIO DECISIONS of a long-lived investor with constant relative risk aversion (CRRA) in the presence of transaction costs. Two types of costs are evaluated: *proportional* to the change in the holding of the risky asset and a *fixed* fraction of portfolio value. With a constant opportunity set, it is well known that these costs induce optimal rebalancing rules that involve a no-trade region for the CRRA investor's risky-asset holding. When the investor's holding goes outside the no-trade region, the investor rebalances to the closest boundary (proportional cost only) or to a point inside the region (nonzero fixed cost).¹ In this paper, we focus on how return predictability affects portfolio choice and rebalancing behavior when transaction costs are nonzero.

Recent papers show that predictability calibrated to U.S. data has a large effect on the rebalancing behavior of a multiperiod investor. We find that this continues to be true in the presence of realistic transaction costs. When

* Lynch is from New York University, and Balduzzi is from Boston College. The authors thank John Campbell, Jennifer Carpenter, George Constantinides, Ned Elton, Silverio Foresi, John Heaton, John Liew, Robert Whitelaw, and especially René Stulz (the editor) and an anonymous referee for helpful conversations and comments. Financial support from New York University Summer Research Grants is gratefully acknowledged.

¹ See Constantinides (1986) and Davis and Norman (1990) for proportional costs and Schroder (1995) and Morton and Pliska (1993) for fixed costs. A number of other papers consider the effects of transaction costs on portfolio choice and market equilibrium: see, for example, Duffie and Sun (1990), Heaton and Lucas (1996), Vayanos (1996), Gennotte and Jung (1994), Luttmer (1996a, 1996b), and He and Modest (1995).

returns are predictable, the no-trade region is state dependent and, on average, wider and higher than when returns are independently and identically distributed (i.i.d.) with the same unconditional distribution. In fact, predictability motivates the investor to spend considerably more on rebalancing and to rebalance more often. The average allocation to the risky asset early in life is also higher with predictability.

In addition, the presence of realistic transaction costs is found to materially affect rebalancing behavior. For example, the presence of a realistic fixed cost causes rebalancing frequency to decline considerably. At the same time, rebalancing costs alone have almost no effect on the average allocation to the risky asset, whereas costly liquidation of the risky asset for consumption lowers average allocations, though only marginally early in life. Finally, the no-trade region widens late in life but only when death is certain and imminent. In other results, we find that varying the nature of the return predictability or introducing return heteroskedasticity can drastically affect the investor's portfolio choices.

A long line of literature has established that returns are predictable.² Recent work has attempted to assess the impact of return predictability on investor utility and portfolio choice. Kandel and Stambaugh (1996) explore the effects of ignoring predictability in a myopic setting, whereas Brennan, Schwartz, and Lagnado (1997) and Barberis (1999) analyze numerically the impact of myopic versus dynamic decision making when returns are predictable. 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; Kim and Omberg (1996) solve an analogous continuous-time problem but without intermediate consumption. Brandt (1999) estimates an investor's stock allocation using the investor's Euler equations and U.S. stock return data. None of these papers simultaneously considers the impact of return predictability and rebalancing costs on portfolio choice. The only paper to do so is Balduzzi and Lynch (1999), the focus of which is on the utility costs associated with ignoring predictability and behaving myopically. In contrast, the current paper focuses on how predictability and heteroskedasticity alter rebalancing rules and behavior in the presence of rebalancing and liquidation costs.

We solve the investor's problem numerically given a risky asset and a riskless asset. Using Tauchen and Hussey's (1991) quadrature approximation, the risky-asset return is calibrated to the value-weighted index of NYSE stocks in such a way as to capture the predictive ability of aggregate dividend yield. Other forms of predictability are also calibrated. Rebalancing frequencies, average costs incurred, and average holdings are obtained by simulating asset return time series and using the investor's portfolio rebalancing rules. The main results can be summarized as follows.

² Campbell (1987) and Fama and French (1989), among others, find that the dividend yield, the term premium, and the one-month T-bill rate all forecast future U.S. stock returns.

With the introduction of transaction costs, return predictability calibrated to U.S. returns continues to have large effects on optimal rebalancing behavior relative to that for i.i.d. returns with the same unconditional distribution. In particular, by increasing the benefits from rebalancing, return predictability causes rebalancing frequency to increase, and cost incurred to increase by an order of magnitude, at all points in the investor's life. These results complement those of Balduzzi and Lynch (1999), who find that the utility cost of ignoring predictability remains substantial (typically around 20 percent of wealth) in the presence of realistic rebalancing costs.

Consistent with results for the no-cost case (see, e.g., Barberis (1999)), the average allocation to the risky asset early in life is also higher with predictability. Further, no-trade regions early in life are wider when returns are predictable than when they are i.i.d., which can be explained as follows. Ignoring transaction costs, a less volatile risky-asset return means a smaller utility loss from keeping a given inherited risky-asset allocation rather than trading. For this reason, lower risky-asset volatility leads to wider no-trade regions (see Constantinides (1986) and Gennotte and Jung (1994), who discuss this point in the context of the volatility of i.i.d. returns). Because unconditional volatility is unchanged, predictability reduces conditional volatility. It is this reduction in conditional volatility that causes no-trade regions to widen relative to the i.i.d. case. Finally, predictability also causes the no-trade region to move around, which explains why the rebalancing frequency increases despite the wider no-trade region.

The presence of realistic rebalancing costs can considerably affect rebalancing rules and behavior over the life cycle. First, realistic fixed rebalancing costs cause rebalancing frequency to decline considerably.³ Second, no-trade regions widen dramatically close to the terminal date.⁴ On the other hand, realistic proportional and fixed costs have little effect on the average risky-asset holding and, in particular, on the no-trade midpoint unless liquidation costs differ across the two assets. When the investor faces a liquidation cost on only the risky asset, no-trade regions are lowered, but the lowering is only material late in life when consumption is a large fraction of wealth. In fact, as the terminal date approaches, the upper boundary of the no-trade region converges to the risky-asset holding in the absence of transaction costs.

The nature of the return predictability can have large effects on rebalancing behavior. Holding unconditional return moments fixed, we focus on two return parameters, each with a clear impact on portfolio choice: the magnitude of the single-period predictability and the persistence of the predictive variable. Increasing the magnitude of the single-period predictability causes no-trade regions to widen. Decreasing the persistence of the predictive vari-

³ These results are consistent with those of Schroder (1995), who finds for i.i.d. returns that a small fixed cost (but *no* proportional cost) can cause rebalancing frequency to drop dramatically from the no-cost case.

⁴ Gennotte and Jung (1994) report a similar result but for the more restrictive case where the risky asset return is an i.i.d binomial process.

able causes no-trade width, as a function of the predictive variable, to become more U-shaped; that is, the no-trade region is wider for extreme values of the predictive variable. Finally, return heteroskedasticity can also have big effects on rebalancing behavior. Allowing conditional volatility to be a steeper positive function of expected return causes no-trade width and midpoint to be less positive (or more negative) functions of expected return.

Thus, realistic transaction costs and return predictability can have a significant impact on rebalancing behavior. The implication is that their simultaneous presence is likely to affect the joint distribution of consumption and equity returns. Thus, calibration studies using transaction costs to explain the equity premium (e.g., Heaton and Lucas (1996)) may be sensitive to the introduction of predictability in returns.

Our paper also has implications for existing studies (He and Modest (1995), Luttmer (1996a)) documenting how proportional transaction costs weaken Hansen–Jagannathan (1991) bounds on intertemporal marginal rates of substitution. Because these bounds become weaker relative to the frictionless bound as the data frequency becomes higher, they are likely to be uninformative for data frequencies that are too much higher than the rebalancing frequency of individuals. Thus, our results on rebalancing frequencies can shed some light on the informativeness of the bounds reported in these papers. For example, with predictability and a proportional cost of 0.5 percent, our 20-year investor rebalances every 3.7 months, which is only slightly less frequently than every quarter. This rebalancing frequency suggests that the bound reported by Luttmer for quarterly data and a proportional cost of 0.5 percent is likely to be informative.

Samuelson (1969) was the first to show that a CRRA investor, facing a constant investment-opportunity set and zero transaction costs, makes the same portfolio choice irrespective of age. Interestingly, our results for the case of unpredictable returns indicate that Samuelson's irrelevance result is robust to the introduction of transaction costs except when death is certain and imminent. In this sense, Samuelson's (1969) original intuition extends to the case of transaction costs.

A number of recent papers emphasize the importance of nonfinancial wealth particularly early in life (see, e.g., Bodie, Merton, and Samuelson (1992), Koo (1995), Jagannathan and Kocherlakota (1996), and Heaton and Lucas (1997)). Because of the potentially large impact of nonfinancial wealth, care must be taken when applying our results to a young investor. For example, if most of an investor's wealth is nonfinancial, then a fixed cost proportional to the investor's financial wealth understates the cost of the investor's time. However, our results are directly applicable to an investor who has reached retirement age: our standard investor's death probabilities are calibrated so that the investor is 65 years and one month old in the first period.

The paper is organized as follows. Section I describes the investor's optimization problem, our solution technique, and the simulation details. Section II describes how the discrete return processes are calibrated. Section III presents the results, and Section IV concludes.

I. Optimal Portfolio Allocation with Transaction Costs

A. Constraints and Preferences

We consider situations where *two* assets are available for investment: a risky asset and a riskless asset. The risky-asset return from time t to $t + 1$, R_{t+1} , is either i.i.d. for all t or predictable using an instrument, D_t , available at t . The risk-free rate R^f is assumed to be constant. Also, the investor faces transaction costs that are proportional to wealth.

We consider the optimal portfolio problem of a CRRA investor with a finite life of T periods, a relative-risk-aversion parameter of γ , and a rate of time preference equal to β . The investor's probability of dying at time t conditional on survival until $t - 1$ is denoted by p_t , which may be nonzero. Expected lifetime utility is given by

$$E \left[\sum_{t=1}^T \left(\prod_{\tau=1}^t \beta_{\tau} \right) \frac{c_t^{1-\gamma}}{1-\gamma} | \mathbf{Z}_1 \right], \quad (1)$$

where c_t is consumption at time t , $\mathbf{Z}_t$ is the vector of state variables for the investor at time t , and $\beta_{\tau} = \beta(1 - p_{\tau})$. When transaction costs are nonzero, the inherited portfolio allocation $\hat{\alpha}_t$ is an element of $\mathbf{Z}_t$, because its value determines the transaction costs to be paid at time t . When returns are predictable, the predictive variable at time t , D_t , is also an element of $\mathbf{Z}_t$.

Letting κ_t be the fraction of wealth consumed at t and f_t be the transactions cost paid at t per dollar of portfolio value (which might be zero if no trading occurs at t), the law of motion of the investor's wealth, W , is given by

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

where $R_{W,t+1} = (1 - f_t)[\alpha_t(R_{t+1} - R^f) + R^f]$ and, α_t is the share of portfolio value allocated to the risky asset at t . So $R_{W,t+1}$ is the rate of return on the portfolio from t to $t + 1$, net of transaction costs. Dollar transaction costs at t are $W_t(1 - \kappa_t)f_t$ and are paid by costlessly liquidating the risky and the riskless assets in the proportions α_t and $(1 - \alpha_t)$.

In general, we model the cost of transacting, f , as a function of the difference between the chosen and inherited risky-asset allocations:

$$f_t = \phi_P |\alpha_t - \hat{\alpha}_t| + \phi_F I_{\alpha_t - \hat{\alpha}_t \neq 0}, \quad \phi_P, \phi_F \geq 0, \quad (3)$$

where $I_{\alpha_t - \hat{\alpha}_t \neq 0}$ is an indicator function that equals one if $\alpha_t - \hat{\alpha}_t \neq 0$ and equals zero otherwise. The first term is proportional to the change in the risky-asset holding, whereas the second term reflects the fixed cost of rebalancing one's portfolio, regardless of the size of the rebalancing. This fixed cost increases with the investor's wealth, because it is likely to depend on the opportunity cost of the investor's time. The inherited allocation satisfies

$$\hat{\alpha}_t \equiv \frac{\alpha_{t-1}(1 - \kappa_{t-1})W_{t-1}(1 - f_{t-1})R_t}{W_t} = \frac{\alpha_{t-1}R_t}{\alpha_{t-1}(R_t - R^f) + R^f}. \quad (4)$$

The law of motion for wealth in equation (2) implicitly assumes that consumption at time t is obtained by costlessly liquidating the risky and the riskless asset in the proportions $\hat{\alpha}_t$ and $(1 - \hat{\alpha}_t)$. However, rebalancing rules may be sensitive to any cost differential associated with using the risky rather than the riskless asset for consumption (see, e.g., Heaton and Lucas (1997)). Consequently, with the fixed cost parameter set to zero ($\phi_F = 0$), we introduce a proportional cost of ϕ_L to liquidate the risky asset for consumption. Letting $f_{L,t}$ be the liquidation cost incurred at t as a fraction of W_t , the law of motion for wealth becomes

$$W_{t+1} = W_t(1 - \kappa_t - f_{L,t})R_{W,t+1}, \quad (5)$$

where $f_{L,t} = \kappa_{L,t}\phi_L$ and $\kappa_{L,t} \geq 0$ is the component of consumption at t obtained by liquidating the risky asset. The liquidation cost is assumed to be paid from the risky asset, so $\kappa_{L,t} \leq \hat{\alpha}_t/(1 + \phi_L)$. The pre-rebalancing portfolio risky-asset weight at t now depends on the consumption decision at t . After consumption, the risky-asset holding is $\hat{\alpha}_t - \kappa_{L,t} - f_{L,t}$ and the portfolio value is $1 - \kappa_t - f_{L,t}$, each expressed as a fraction of W_t . Because the pre-rebalancing risky-asset weight is the ratio of these two, the rebalancing cost f_t becomes

$$f_t = \phi_P \left| \alpha_t - \frac{\hat{\alpha}_t - \kappa_{L,t} - f_{L,t}}{1 - \kappa_t - f_{L,t}} \right|. \quad (6)$$

We set ϕ_L equal to $\phi_P/(1 - \phi_P)$ so that the cost of liquidating the entire risky-asset holding equals the cost of rebalancing the risky-asset holding to zero. In this sense, our ϕ_L choice equates the cost of liquidating the risky asset with the cost of rebalancing.

B. Optimization Problem and Solution Technique

We now consider the investor's optimization problem. Given our parametric assumptions, the Bellman equation faced by the investor is given by

$$\begin{aligned} \frac{\alpha(\mathbf{Z}_t, t) W_t^{1-\gamma}}{1 - \gamma} &= \max_{\kappa_t, \alpha_t} \frac{\kappa_t^{1-\gamma} W_t^{1-\gamma}}{1 - \gamma} \\ &\quad + \beta_t \frac{(1 - \kappa_t - f_{L,t})^{1-\gamma} W_t^{1-\gamma}}{1 - \gamma} E[a(\mathbf{Z}_{t+1}, t+1) R_{W,t+1}^{1-\gamma} | \mathbf{Z}_t], \\ &\text{for } t = 1, \dots, T-1 \end{aligned} \quad (7)$$

where $f_{L,t} \equiv 0$ in the absence of a risky-asset liquidation cost. Because equation (7) only defines $a(\mathbf{Z}_T, t)$ recursively, we must specify its value at the terminal date: $a(\mathbf{Z}_T, T)$ is 1 in the absence of a liquidation cost, and $(1 - \phi_P \hat{\alpha}_T)^{1-\gamma}$ when ϕ_L is equal to $\phi_P/(1 - \phi_P)$. In all the optimizations, the holdings of both the risky and the riskless asset are constrained to be positive (the short-sale restriction), which implies that the inherited portfolio allocation also lies between zero and one, $0 \leq \hat{\alpha} \leq 1$. Constraining α to lie between zero and one is realistic, because individual investors typically face high costs in taking short positions.

The optimization problem is homogeneous of degree $(1 - \gamma)$ in wealth, which implies that the solution is invariant to wealth. Thus the Bellman equation can be rewritten as

$$\frac{a(\mathbf{Z}_t, t)}{1 - \gamma} = \max_{\kappa_t, \alpha_t} \frac{\kappa_t^{1-\gamma}}{1 - \gamma} + \beta_t \frac{(1 - \kappa_t - f_{L,t})^{1-\gamma}}{1 - \gamma} E[a(\mathbf{Z}_{t+1}, t + 1) R_{W,t+1}^{1-\gamma} | \mathbf{Z}_t],$$

for $t = 1, \dots, T - 1$. (8)

The Bellman equation (8) is solved by backward iteration, starting with $t = T - 1$. Thus, $a(\mathbf{Z}_t, t)$ is obtained by solving the optimization problem in equation (8) using $a(\mathbf{Z}_{t+1}, t + 1)$ from the previous iteration.

At each time t , the state variable $\hat{\alpha}_t$ is discretized, and the Bellman equation is solved at each $\hat{\alpha}_t$ grid point. The following grid of points on the interval $[0, 1]$ is used to discretize $\hat{\alpha}_t$ for all t : $\hat{\alpha} = 0, 0.02, 0.04, \dots, 0.96, 0.98, 1$. In solving the optimization problem in equation (8), the $(t + 1)$ value function for each D_{t+1} is linearly interpolated between $\hat{\alpha}_{t+1}$ points. Further details are available on request. The numerical technique yields an approximate solution that converges to the actual solution as the $\hat{\alpha}$ grid becomes finer.⁵

C. Unconditional versus Conditional Portfolio Choices

Facing a given return-generating process, the investor can make unconditional (U) or conditional (C) portfolio choices. When making unconditional choices, the investor uses the steady-state distribution and ignores any predictability of returns. In other words, the investor assumes returns are i.i.d. when making unconditional choices. In contrast, the investor exploits return predictability when making conditional choices. Thus, we can evaluate the impact of return predictability on portfolio choice when an investor faces proportional or fixed transaction costs.

D. Simulation Details

To calculate average portfolio holdings, average transaction costs incurred, and rebalancing frequencies, we simulate return histories with the discretized distribution and apply the individual's portfolio rebalancing rule

⁵ Increasing the number of grid points from 51 to 101 for $\hat{\alpha}$ leaves the results for a 20-year investor with $\phi_P = 0.5$ percent and $\phi_F = 0$ or 0.01 percent virtually unchanged.

at each time $t, t = 1, \dots, T$. This procedure is repeated 100,000 times, and average portfolio holdings, rebalancing costs incurred, and rebalancing frequencies at each time t are obtained by averaging over the 100,000 replications.

For unconditional portfolio choices, simulated returns are i.i.d. with the steady-state distribution. By making simulated returns i.i.d. in this way, we are assuming that the investor making unconditional choices is using the correct distribution. This approach allows us to obtain an understanding of how optimal rebalancing behavior is altered by return predictability. A different question, not answered in this paper, is the effect of using a policy rule optimized for i.i.d. returns when returns are in fact predictable. This question is addressed in Balduzzi and Lynch (1999) using utility cost as the metric to evaluate the effect of ignoring predictability.

II. Return Calibration

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 12-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. We calibrate real returns because they are more likely to be stationary; investors generally care more about real returns, and there is no money in our model. The stock return, interest rate, and dividend yield series are from CRSP; the CPI series is from CITIBASE. The data period used is from 1927:1 to 1996:11. 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^f of 0.04454 percent.

A. Approximating the Data Assuming Constant Variance

Letting $r \equiv \ln(1 + R)$ and $d \equiv \ln(1 + D)$, we assume that $[r, d]'$ follows the vector autoregressive model (VAR):

$$r_{t+1} = a_r + b_r d_t + e_{t+1}, \quad (9)$$

$$d_{t+1} = a_d + b_d d_t + v_{t+1}, \quad (10)$$

where b_r , b_d , a_r , and a_d are coefficients and $[e, v]'$ is a vector of mean-zero, serially uncorrelated, multivariate normal disturbances, with *constant* covariance matrix whose diagonal elements are σ_e^2 and σ_v^2 and whose off-diagonal element is $\rho_{ev}\sigma_e\sigma_v$. Similarly, the unconditional variances for r and d are σ_r^2 and σ_d^2 , respectively. Without loss of generality, we normalize the mean of d , μ_d , to be zero and its variance, σ_d^2 , to be 1. The specifications given in equations (9) and (10) assume that d_t is the only state variable needed to forecast r_{t+1} , which is in line with other papers on optimal portfolio selection (e.g., Barberis (1999), Campbell and Viceira (1996)).

The data VAR is estimated using ordinary least squares (OLS) and discretized using a variation of Tauchen and Hussey's (1991) Gaussian quadrature method; the variation is designed to ensure that d is the only state variable (for details, see Balduzzi and Lynch (1999)). We choose 19 quadrature points for the dividend yield and three points for the stock return innovations because Balduzzi and Lynch (1999) find that the resulting approximation is able to capture important dimensions of the predictability in the data.⁶

An interesting question is how portfolio choice varies with the nature of the return predictability. We assess this by holding the return parameters σ_r , μ_r , and ρ_{ev} constant and varying $b_r\sigma_d$ and b_d . The first parameter, $b_r\sigma_d$, is the single-period variation in r that is predictable. The second parameter, b_d , measures the persistence of the expected return process, because expected return is a linear function of d . We implement the Tauchen–Hussey approximation in such a way that the 19 values taken by d remain the same across different parameterizations of the return predictability. Doing so ensures that portfolio choice for a given state can be meaningfully compared across return-generating processes.

Table I presents VAR parameter values for both the data and the various quadrature approximations used. The unconditional mean return is successfully held constant across the return-generating processes and so is not reported. Six discretizations are considered and are labeled S0 through S5. Discretization S0 (second row) successfully replicates the predictability in the data (first row), except that the predictable single-period variation in $r(b_r\sigma_d)$ is understated. Thus, our results for S0 are likely to understate the impact of observed predictability on portfolio choice. The first four rows after the “Data” row show two pairs of comparisons: S0 versus S1 and S2 versus S3. Each comparison successfully fixes σ_r , b_d , and ρ_{ev} , while allowing $b_r\sigma_d$ to vary. Starting with a value of 0.26 each time, $b_r\sigma_d$ is halved going from S0 to S1 and roughly doubled going from S3 to S2. Similarly, moving from S4 to S5 to S3 keeps σ_r , $b_r\sigma_d$, and ρ_{ev} approximately constant, while allowing b_d to vary from 0.962 to 0.

B. Approximating the Data Assuming Heteroskedasticity

Equilibrium models such as Merton's (1980) dynamic CAPM imply a positive relation between expected return and conditional volatility.⁷ Such a relation is likely to have a large effect on portfolio rebalancing rules. We assess this effect by modifying the return-generating process to allow con-

⁶ Increasing the number of quadrature points to 20 for dividend yield or to six for stock return innovations leaves the results for a 20-year investor with $\phi_p = 0.5$ percent and $\phi_F = 0$ or 0.01 percent virtually unchanged.

⁷ Although recent papers have found a negative relation between conditional expected return and volatility (see, e.g., Glosten, Jagannathan, and Runkle (1993) and Whitelaw (1994)), Scruggs (1998) finds a positive partial relation once he includes long-term government bond return as a second factor.

Table I
Parameters for the Return-generating Processes and Estimates
for the Value-weighted Index of NYSE Stocks

A variety of return-generating processes are obtained using the quadrature approximation. Table I reports parameters for the various return processes and parameter estimates for the value weighted index of NYSE stocks. The following return parameters are reported: the magnitude of the single-period predictability ($b_r\sigma_d$), the conditional correlation between r and d (ρ_{ev}), the persistence of the predictability (b_d), and the heteroscedasticity parameter (ζ_1). The S0 process is calibrated to U.S. equity return assuming constant variance ($\zeta_1 = 0$).

Parameter Being Varied	Process	$\sigma_r \times 100$	$b_r\sigma_d \times 100$	b_d	ρ_{ev}	$\zeta_1 \times 100$
	Data	5.510	0.310	0.972	-0.925	1.358
$b_r\sigma_d$	S0	5.443	0.260	0.962	-0.923	0
	S1	5.445	0.130	0.962	-0.923	0
	S2	5.507	0.527	0	0	0
	S3	5.507	0.264	0	0	0
b_d	S4	5.507	0.260	0.962	0	0
	S5	5.507	0.260	0.386	0	0
	S3	5.507	0.264	0	0	0
ζ_1	H0	5.386	0.260	0.962	-0.904	1.279
	H1	5.443	0.260	0.962	-0.923	0.569
	H2	5.432	0.260	0.962	-0.919	0.110

ditional return volatility to be a linear function of d . In particular, $[r, d]'$ still satisfies equations (9) and (10), but now $e_{t+1} = \sigma_{e|d_t} z_t^e$ and $v_{t+1} = \sigma_{v|d_t} z_t^v$, where z^e and z^v have zero means, unit variances, and a correlation of $\rho_{ev|d}$; $[z^e, z^v]$ is multivariate normal white noise; and $\sigma_{e|d_t}$ and $\sigma_{v|d_t}$ are the conditional volatilities of e_{t+1} and v_{t+1} . The constant $\rho_{ev|d}$ is the conditional correlation between e and v . We take $\sigma_{v|d_t}$ to be a constant but allow $\sigma_{e|d_t}$ to be a linear function of d_t : $\sigma_{e|d_t} = \zeta_0 + \zeta_1 d_t$. The quadrature approximation technique is adapted to discretize this return process (details are available upon request).

Return distributions are discretized for three values of ζ_1 , holding the unconditional volatility of e , σ_e , constant. The first value is obtained from the data using the residuals from the OLS regression of r_{t+1} on d_t . The absolute value of these residuals is regressed on d_t , and the slope coefficient is taken as ζ_1 . The first row of Table I reports the data's ζ_1 as 1.358 percent. Because this ζ_1 value of 1.358 percent may be too large to be plausible in an equilibrium setting, the single-regime equilibrium model of Whitelaw (1998) is used to calibrate ζ_1 with positive values. In particular, Whitelaw's model implies ζ_1 values of 0.111 percent and 0.604 percent when the representative agent has a relative risk aversion of two and 40, respectively. The last three rows of Table I show that the three discretizations, H0, H1, and H2, closely match ζ_1 values of 1.358 percent, 0.604 percent, and 0.111 percent, respectively, while keeping the unconditional volatility for r , σ_r , close to that in the data.

III. Results

Optimal portfolio choice with fixed and proportional transaction costs involves a no-trade region for the risky-asset weight α and lower and upper return points depending on whether the upper or lower boundary is hit (Balduzzi and Lynch (1999)). Both return points lie inside the no-trade region. When the investor uses the conditional return distribution (C) rather than the unconditional (U), the position of the no-trade region varies across states (see Balduzzi and Lynch (1999)).

Given this structure for the rebalancing decision, portfolio choice in a given state can be characterized by the midpoint and width of the no-trade region together with the average distance between the return points and boundaries (return distance). These three parameters can be averaged across states at a given time t using the steady-state distribution to give the time- t no-trade midpoint, the time- t no-trade width, and the time- t return distance, respectively.⁸

For each t , three additional parameters are obtained from the simulations described in Section I: rebalancing frequency, average cost incurred, and average holding of the risky asset.⁹ Thus, there are six parameters all indexed by t that characterize portfolio choice by the investor. We also want to measure when the no-trade region first widens by a significant amount relative to its $t = 1$ width. The seventh parameter is such a measure and is the first time t that the width of the no-trade region is at least 0.05 larger than at time one. Each table and figure reports or plots some subset of these seven parameters.

A. Standard Investor Facing Returns Calibrated to U.S. Data

Table II and Figure 1 report results for the “standard” investor facing the S0 return-generating process which is calibrated to the value-weighted index of NYSE stocks. The standard investor has $\gamma = 4$, intermediate consumption, and a death probability at each t calibrated to mortality rates for U.S. males with $t = 1$ corresponding to 65 years and one month old. The mortality rates that we use are taken from the 1994 Group Annuity Mortality Table developed by the Society of Actuaries Group Annuity Valuation Task Force (1995) (GAM-94). Because these rates only assign a death probability of one when the male reaches 119 years of age, we take T to be 648. For comparison, Table II also reports results for a “20-year” investor who

⁸ To ensure that the average no-trade width and return distance are not understated, the averaging is performed across states for which the short-sale restriction is not binding. As expected, averaging over all states reduces no-trade width and does so by more late in life. For the standard investor, the reduction is only four percent late in life and less than two percent early in life. On the other hand, because bounds of zero or one still define ranges over which α moves, the averaging for the no-trade midpoint is performed across all states irrespective of whether a short-sale restriction is binding.

⁹ Risky-asset holding is determined after any rebalancing has occurred.

Table II
Portfolio Choice Parameters for the Standard Investor

The value-weighted index of NYSE stocks is discretized using the quadrature approximation to give the return process S_0 in Table I. The riskfree rate is assumed constant. The proportional cost parameter (ϕ_P) set equal to 0.25 percent and the fixed cost parameter (ϕ_F) is set equal to 0.01 percent. For the standard investor (CRRA utility, $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption), $t = 1$ is taken to be age 65 and 1 month when using CAM-94 to obtain death probabilities. The 20-year investor lives for $T = 240$ months and dies at T but is otherwise identical to the standard investor. The width and midpoint of the no-trade region, and the average return distance at a point in time t are obtained by averaging across states using the steady state distribution. Rebalancing frequency, average cost incurred per month and average holding of the risky asset are obtained by simulation of the investor's portfolio choices. Parameters for both unconditional (U) and conditional (C) portfolio choices are reported. The portfolio choice parameters are averaged over three two-year periods in the investor's life: $t = 96$ to $t = 119$; $t = 216$ to $t = 239$; and, $t = 624$ to $t = 647$.

Portfolio Choice Parameters		$t = 96$ to $t = 119$		$t = 216$ to $t = 239$		$t = 624$ to $t = 647$
		Standard	20-year	Standard	20-year	Standard
Rebalancing Frequency (%)	U	2.80	2.83	2.84	1.12	1.13
	C	10.33	10.35	10.27	8.71	8.70
Average Cost per Month (bp)	U	0.075	0.076	0.076	0.032	0.033
	C	0.571	0.576	0.573	0.517	0.517
Holding of Risky Asset-No Transaction Costs	U	0.562	0.562	0.562	0.562	0.562
	C	0.694	0.680	0.686	0.581	0.580
Average Holding of Risky Asset	U	0.574	0.574	0.574	0.583	0.585
	C	0.711	0.696	0.702	0.600	0.599
Midpoint of No-trade Region	U	0.560	0.560	0.560	0.563	0.562
	C	0.699	0.684	0.690	0.579	0.578
Width of No-trade Region	U	0.160	0.160	0.160	0.262	0.266
	C	0.258	0.258	0.261	0.315	0.316
Distance between Boundaries and Return Points	U	0.060	0.060	0.060	0.079	0.080
	C	0.083	0.083	0.085	0.092	0.092

lives for 240 months and only dies at T but who otherwise resembles the standard investor. Our standard investor is already 65 years old at $t = 1$ because the investor does not receive labor income and this is only realistic for an individual who has reached retirement age.

The investors are assumed to face both a fixed and a proportional cost. The fixed cost parameter, ϕ_F , is taken to be 0.01 percent, which translates into paying a fee of \$10 whenever a \$100,000 portfolio is reshuffled. Viewed as the opportunity cost for an individual to process information and instruct a broker to change portfolio composition, this value for ϕ_F seems small. The proportional cost parameter, ϕ_P , is set to 0.25 percent, which implies a round-trip transaction cost of 0.5 percent. Both Fidelity's Spartan Total Market Index Fund and the Schwab Total Market Index Fund attempt to track value-




Figure 1. Portfolio choice parameters: The standard investor. Portfolio choice parameters are plotted for the standard investor (CRRA utility, $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption, $T = 648$) for whom $t = 1$ is taken to be age 65 and one month when using CAM-94 to obtain death probabilities. The investor faces the return process S0 (quadrature approximation to U.S. data) and makes unconditional (U) and conditional (C) portfolio choices. There is a fixed cost (ϕ_F) of 0.01 percent, and the proportional cost (ϕ_P) is set equal to 0.25 percent.

weighted portfolios of U.S. equities and, as of May 5, 1999, charge a redemption fee of 0.5 percent on fund shares sold within three and six months of purchase, respectively. Thus, whereas our proportional cost structure does not capture the specific features of these redemption fees, our value for ϕ_P of 0.25 percent can be viewed as indicative of the proportional rebalancing cost faced by an individual using an index fund to hold U.S. equities.

A.1. Portfolio Choice over the Life Cycle

The first question that we address is portfolio choice over the life cycle. Table II shows that both average rebalancing frequency and average cost incurred are lower when certain death is imminent than early in the investor's life. Figures 1A and 1C show that these decreases are particularly pronounced just prior to the terminal date.

Consistent with this result, we find that the no-trade width is always much larger just prior to the terminal date ($T = 648$) than earlier in life. Figure 1B and Table II show that the average no-trade width for the standard investor using C goes from 0.258 early in life to more than 0.60 by time $t = 647$. The intuition for the widening of the no-trade region near the end of an investor's life is as follows. Early in the investor's life, a decision to rebalance may eliminate the need to rebalance for several periods in the future. In this sense, the transaction cost of rebalancing is being spread over the next several periods. However, near the end of the investor's life, this potential benefit from rebalancing is limited by the small number of remaining periods in the investor's life. Consistent with this intuition, Figure 1D shows that the widening does not occur until certain death is imminent. This intuition also suggests that the no-trade region widens earlier when the investor's rebalancing frequency is lower. The reason is that a lower rebalancing frequency indicates that the benefits from rebalancing now must be enjoyed over a longer period of time to make rebalancing now attractive.

Another question is whether rebalancing rules and behavior change much as the standard investor goes from being 65 to 85 years old. Table II and Figure 1 show that the portfolio choice parameters are largely unchanged over these early retirement years, despite the use of realistic death probabilities. Thus, the rebalancing behavior of older investors in the United States does not involve a rapid widening of the no-trade region unless certain death is imminent.

A.2. Impact of Return Predictability on Rebalancing

Previous work has shown that predictability calibrated to U.S. data causes a multiperiod investor to hold more of the risky asset early in life (see, e.g., Brennan et al. (1997), Barberis (1999), and Campbell and Viceira (1996)). Consistent with this work, we find in Table II and Figure 1D that early in life with zero transaction costs, the standard investor chooses a risky-asset holding of 0.562 using U and an average risky-asset holding of 0.694 using C. Table II and Figure 1C show that this result is robust to the introduction of realistic transaction costs. In fact, for both C and U, the midpoint of the no-trade region is roughly equal to the optimal risky-asset holding without transaction costs for all t , whereas the average holding is marginally higher but by the same amount for both.

Table II shows that rebalancing frequency over all three two-year periods is higher when the standard investor faces C rather than U. For example, rebalancing frequency early in life increases from 2.80 percent (or once every 35 months) with i.i.d. returns to 10.33 percent (or once every nine months)

when returns are predictable. Figure 1 shows that the rebalancing frequency is higher throughout the standard investor's life when facing C rather than U.

An even stronger result is the effect of return predictability on the actual cost incurred by the investor. Table II and Figure 1A show that the cost incurred per month is an order of magnitude higher when the investor faces C rather than U, irrespective of whether the investor is young or old. For example, early in life the standard investor spends 0.075 basis points of portfolio value per month when facing U but spends 0.571 basis points when facing C. The increased rebalancing and spending are not surprising because return predictability increases the benefits from trading.

At the same time, the no-trade width early in life is much wider when the investor uses the conditional rather than the unconditional distribution. Table II shows that the no-trade width early in life for the standard investor increases from 0.16 to 0.258 if returns are predictable. Near the end of the investor's life, the difference narrows, as can be seen in Figure 1B. Earlier work on transaction costs without predictability finds that no-trade regions increase when return volatility decreases (see Constantinides (1986) and Gennotte and Jung (1994)). The reason is that the cost of a suboptimal α choice is increasing in the return volatility. Going from the U to C distribution, the conditional volatility of return decreases because the unconditional volatility is being held fixed. Thus, this same reasoning explains why no-trade regions widen when returns are predictable.

When the investor is using C, the no-trade region varies considerably across states (unreported). This occurs because the conditional Sharpe ratio varies across states. This variability in the no-trade region explains why rebalancing frequency is higher for C than U, despite the wider no-trade region.

A.3. Comparing the Standard and 20-Year Investors

Comparing the standard and 20-year investors early in life ($t = 96$ to $t = 119$), we see that all six portfolio-choice parameters indexed by t are virtually identical. These parameters are also similar across the two investors when both are near the time of certain death: $t = 624$ to $t = 647$ for the standard investor and $t = 216$ to $t = 239$ for the 20-year investor. So whereas positive death probabilities are analogous to lower rates of time preference, the death probabilities faced by U.S. males from age 65 onward are not large enough to materially alter rebalancing rules. As a result, because computation time per investor problem is much lower when $T = 240$ than when $T = 648$, we use the 20-year investor as our canonical investor for the remainder of the paper.

B. Varying the Transaction Costs

B.1. Varying the Rebalancing Costs

Table III reports all seven portfolio choice parameters for the 20-year investor facing the S0 return-generating process. The proportional cost parameter (ϕ_P) is allowed to take four values, 0, 0.25 percent, 0.5 percent,

Table III
Portfolio Choice Parameters as Transaction Costs Vary

The value-weighted index of NYSE stocks is discretized using the quadrature approximation to give the return process S_0 in Table I. The riskfree rate is assumed constant. The 20-year investor (CRRA utility, $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption, $T = 240$, no early death) faces a variety of fixed (ϕ_F) and proportional (ϕ_P) cost combinations. The width of the no-trade region, the midpoint of the no-trade region, and the average return distance at a point in time t are obtained by averaging across states using the steady state distribution. Rebalancing frequency, average cost incurred per month and average holding of the risky asset are obtained by simulation of the investor's portfolio choices. Parameters for both unconditional (U) and conditional (C) portfolio choices are reported. The portfolio choice parameters are averaged over two-year periods in the investor's life: $t = 96$ to $t = 119$; and $t = 216$ to $t = 239$. The table also reports the first time that the width of the no-trade region changes by more than 0.05 relative to time $t = 1$.

Portfolio Choice Parameters	Proportional Cost ϕ_P	Fixed Cost ϕ_F							
		0%	0.01%	0.1%	0%				
Panel A:									
Rebalancing Frequency (%)	U	$t = 96 \text{ to } t = 119$				$t = 216 \text{ to } t = 239$			
		100.00	3.06	1.32	100.00	2.14	0.25	0.1%	0.1%
		25.25	2.83	1.07	14.83	1.12	0.17		
		21.17	2.32	1.11	8.30	0.69	0.09		
	14.70	1.71	0.72	0.13	0.00	0.00			
	C	100.00	15.03	5.29	100.00	13.72	4.62		
		35.52	10.35	4.11	30.46	8.71	3.35		
		27.07	8.85	3.96	21.12	6.36	2.86		
16.03		4.65	2.31	5.26	1.71	0.93			
Average Cost Incurred per Month (basis points of portfolio value)	U	0.000	0.031	0.132	0.000	0.021	0.025		
		0.058	0.076	0.143	0.033	0.032	0.024		
		0.096	0.106	0.182	0.037	0.030	0.016		
		0.309	0.300	0.280	0.002	0.000	0.000		
	C	0.000	0.150	0.529	0.000	0.137	0.462		
		0.510	0.576	0.778	0.466	0.517	0.656		
		0.854	0.898	1.098	0.702	0.700	0.830		
		1.612	1.576	1.524	1.063	1.036	0.995		
Average Holding	U	0.562	0.574	0.580	0.562	0.571	0.596		
		0.571	0.574	0.590	0.574	0.583	0.607		
		0.576	0.586	0.593	0.584	0.593	0.617		
		0.604	0.611	0.616	0.646	0.654	0.672		

Midpoint of No-trade Region (No-trade Midpoint)	C	0%	0.673	0.685	0.701	0.693	0.587	0.609
		0.25%	0.692	0.696	0.707	0.594	0.600	0.618
		0.5%	0.700	0.703	0.712	0.602	0.608	0.627
	U	2.5%	0.742	0.745	0.755	0.660	0.666	0.680
		0%	0.562	0.570	0.560	0.562	0.560	0.562
		0.25%	0.563	0.560	0.560	0.562	0.563	0.558
	C	0.5%	0.563	0.570	0.560	0.562	0.559	0.555
		2.5%	0.562	0.562	0.559	0.543	0.540	0.532
		0%	0.673	0.681	0.676	0.693	0.581	0.576
	U	0.25%	0.687	0.684	0.678	0.581	0.579	0.572
		0.5%	0.690	0.684	0.677	0.580	0.575	0.568
		2.5%	0.699	0.691	0.678	0.558	0.552	0.540
Width of No-trade Region (No-trade Width)	U	0%	0.000	0.140	0.240	0.000	0.172	0.407
		0.25%	0.068	0.160	0.280	0.121	0.262	0.495
		0.5%	0.092	0.180	0.280	0.217	0.348	0.562
	C	2.5%	0.188	0.276	0.352	0.690	0.762	0.875
		0%	0.000	0.197	0.376	0.000	0.204	0.455
		0.25%	0.127	0.258	0.389	0.159	0.315	0.535
	U	0.5%	0.176	0.295	0.409	0.264	0.403	0.603
		2.5%	0.349	0.452	0.542	0.726	0.802	0.902
		0%	0.000	0.070	0.120	0.000	0.204	0.455
	C	0.25%	0.000	0.060	0.130	0.159	0.315	0.535
		0.5%	0.000	0.059	0.124	0.264	0.403	0.603
		2.5%	0.000	0.058	0.114	0.726	0.802	0.902
Distance between Boundaries and Return Points (Return Distance)	U	0%	0.000	0.098	0.188	0.000	0.204	0.455
		0.25%	0.000	0.083	0.176	0.159	0.315	0.535
		0.5%	0.000	0.079	0.172	0.264	0.403	0.603
	C	2.5%	0.000	0.070	0.165	0.726	0.802	0.902
		Panel B: Time of Widening of No-trade Region						
		U	0%	239	235	217		
	0.25%		233	229	219			
	0.5%		227	221	212			
	C	2.5%	190	182	173			
		0%	239	238	229			
		0.25%	235	233	224			
	U	0.5%	231	230	220			
2.5%		192	190	182				

and 2.5 percent, whereas the fixed cost parameter (ϕ_F) is either 0, 0.01 percent, or 0.1 percent. All possible pairwise combinations of ϕ_P and ϕ_F are considered.

Table III highlights how most of the reported parameters vary systematically with the proportional cost (ϕ_P) and the fixed cost (ϕ_F). Rebalancing frequency is decreasing in both. This result is robust to the age of the investor and the presence of predictability and is consistent with the intuition that a larger cost per trade (ϕ_P or ϕ_F bigger) leads to fewer trades. The presence of a small fixed cost has a particularly large impact. In contrast, the cost per trade has an ambiguous impact on the average trading cost incurred per period.

Rebalancing frequency is reduced by a wider no-trade region. Consequently, no-trade width is increasing in both ϕ_P and ϕ_F . This result is robust to the age of the investor and the predictability of returns. On the other hand, return distance is increasing in the fixed cost parameter (ϕ_F) but is insensitive to changes in ϕ_P . Because the fixed cost must be nonzero for the return distance to be positive, this finding is not surprising.¹⁰

Table III shows that, for a given distribution (U or C) and time of life, the no-trade midpoint is unaffected by variation in ϕ_P or ϕ_F . Note that the cost to consume out of either asset is zero across all these specifications. On the other hand, the extent to which the average risky-asset holding exceeds the midpoint is monotonically increasing in both ϕ_P and in ϕ_F . The likely reason is that rebalancing occurs less frequently when transaction costs are high. Because the expected risky-asset return exceeds the riskless rate, less rebalancing is likely to result in larger holdings of the risky asset, holding the no-trade midpoint constant.

Finally, the time of widening of the no-trade region is monotonically decreasing in both ϕ_P and in ϕ_F . This finding is consistent with earlier intuition (in Section III.A.1) because a larger fixed or proportional cost implies less frequent rebalancing.

B.2. Adding a Cost to Liquidate the Risky Asset

We reported above that the no-trade midpoint is unaffected by the magnitude of either the fixed or the proportional cost parameter. However, this result is likely to be sensitive to the assumption of costless liquidation of both assets. To explore the impact of relaxing this assumption, Table IV and Figure 2 report choice parameters for the 20-year investor facing either zero liquidation costs ($f_{L,t} \equiv 0$) or a risky-asset liquidation cost (ϕ_L) of $\phi_P/(1 - \phi_P)$. The proportional cost parameter (ϕ_P) is taken to be 0.5 percent, and there is no fixed cost.

¹⁰ Note that when the proportional cost is zero, the average return distance always equals half the no-trade width, which is consistent with the result of Schroder (1995), who found a single return point.

Table IV

Portfolio Choice Parameters: Introducing a Liquidation Cost

The value-weighted index of NYSE stocks is discretized using the quadrature approximation to give the return process S_0 in Table I. The riskfree rate is assumed constant. The proportional cost parameter (ϕ_P) is set equal to 0.5 percent and the fixed cost parameter (ϕ_F) set equal to 0. The liquidation cost for the risky asset (ϕ_L) is either 0 or $\phi_P/(1 - \phi_P)$. The 20-year investor (CRRA utility, $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption, $T = 240$, no early death) either consumes each month or only at the terminal date. The width and midpoint of the no-trade region at a point in time t are obtained by averaging across states using the steady state distribution. Rebalancing frequency, average cost incurred per month and average holding of the risky asset are obtained by simulation of the investor's portfolio choices. Parameters for both unconditional (U) and conditional (C) portfolio choices are reported. The portfolio choice parameters are averaged over two 2-year periods in the investor's life: $t = 96$ to $t = 119$; and $t = 216$ to $t = 239$.

Portfolio Choice Parameters		$t = 96 \text{ to } t = 119$		$t = 216 \text{ to } t = 239$	
		Liquidation Cost ϕ_L			
		0	$\phi_P/(1 - \phi_P)$	0	$\phi_P/(1 - \phi_P)$
Rebalancing Frequency (%)	U	21.17	48.63	8.30	99.19
	C	27.07	39.98	21.12	80.92
Average Cost per Month (bp)	U	0.096	0.312	0.037	6.942
	C	0.854	0.988	0.702	7.157
Average Holding	U	0.576	0.575	0.584	0.564
	C	0.700	0.698	0.602	0.578
Midpoint of No-trade Region	U	0.563	0.531	0.562	0.428
	C	0.690	0.669	0.580	0.462
Width of No-trade Region	U	0.092	0.105	0.217	0.289
	C	0.176	0.178	0.264	0.302

Table IV and Figure 2 show that a risky-asset liquidation cost induces very different rebalancing behavior just prior to the terminal date when consumption as a fraction of portfolio value is high. The clearest illustration is contained in Figures 2E and 2F, which show the boundaries of the no-trade region facing either the U (Figure 2E) or C (Figure 2F) distribution of S_0 . Also plotted in each graph is the risky-asset holding chosen by the investor facing zero transaction and liquidation costs. With zero liquidation costs, the upper no-trade boundary moves up and the bottom boundary moves down as the region widens late in life. On the other hand, the upper boundary when $\phi_L > 0$ drops late in life, whereas the bottom boundary drops lower than that when $f_{L,t} \equiv 0$. In fact, when $\phi_L > 0$, the upper boundary converges to the no-cost rebalancing point as t approaches 239. It makes sense that the two boundaries drop as the terminal date approaches because consumption is becoming a larger and larger fraction of portfolio value. As one moves away from the terminal date, both bound-




Figure 2. Portfolio choice parameters: Impact of liquidation costs. Portfolio choice parameters are plotted for the 20-year investor (CRRA utility, $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption, $T = 240$, no early death). The proportional cost parameter (ϕ_P) is set equal to 0.5 percent, and the fixed cost parameter (ϕ_F) is set equal to zero. Either liquidation costs are zero ($f_{L,t} = 0$), or the liquidation cost for the risky asset (ϕ_L) is $\phi_P/(1 - \phi_P)$. The investors face the return process S_0 (quadrature approximation to U.S. data) and make unconditional (U) and conditional (C) portfolio choices.

aries increase as consumption expressed as a fraction of portfolio value gets smaller. When $t = 1$, consumption is such a small fraction of wealth that the presence of a risky-asset liquidation cost only lowers the no-trade region marginally: Table IV reports that the average midpoint is lower by




Figure 3. Width of no-trade region by state at $t = 1$: Varying return parameters. The width of the no-trade region at $t = 1$ is plotted by state for the 20-year investor (CRRA utility, $\gamma = 4$, $\beta = 1/R_f$, $T = 240$, intermediate consumption, no early death) making conditional (C) and unconditional (U) portfolio choices. There is no fixed cost, and the proportional cost (ϕ_P) is set equal to 0.5 percent. Parameters of the return processes are summarized in Table 1. The graphs are organized to facilitate comparative-static analysis for two return parameters: the magnitude of the single-period predictability ($b_r \sigma_d$), and the persistence of the predictability (b_d). The value of the parameter being varied is in parentheses for each process.

only 0.025 using U and by only 0.015 using C. Interestingly, Table IV and Figure 2C show that the average risky-asset holding at $t = 1$ (after the investor liquidates and rebalances) is unaffected by the liquidation cost.

C. Varying the Return-Generating Process

To better understand how predictability alters no-trade regions, we take the 20-year investor facing a proportional cost (ϕ_P) of 0.5 percent, and in Figures 3 and 4 plot $t = 1$ no-trade widths as a function of the dividend state for a variety of return processes. The investor's no-trade width for each associated unconditional distribution (U) is also graphed as a flat line.

C.1. Varying the Magnitude of the Single-Period Return Predictability

The parameter $b_r \sigma_d$ controls the magnitude of the single-period predictability. Figures 3A and 3B plot the no-trade width for the S0–S1 and S2–S3 comparisons, respectively, which vary $b_r \sigma_d$ while holding the other return




Figure 4. Width and midpoint of no-trade region at $t = 1$ by state: Introducing heteroskedasticity. The width and midpoint of the no-trade region for $t = 1$ are plotted by state for the 20-year investor (CRRA utility, $\gamma = 4$, $\beta = 1/R_f$, $T = 240$, intermediate consumption, no early death) making conditional portfolio choices. There is no fixed cost, and the proportional cost (ϕ_P) is set equal to 0.5 percent. Parameters of the return-generating processes are summarized in Table 1. The heteroskedasticity parameter, ζ_1 , which measures the extent to which conditional volatility varies with condition mean, ranges from zero to high moving across return-generating processes: S0 to H2 to H to H0. The value of ζ_1 is in parentheses for each process.

parameters fixed. The main result is the increase in no-trade width that accompanies an increase in single-period predictability, irrespective of whether b_d and ρ_{ev} have data values (as in S0 versus S1) or have values of zero (as in S2 versus S3). In particular, the no-trade width for S1 with $b_r\sigma_d = 0.13$ percent drops below the unconditional width in some states. As discussed in Section III.A.2, lower return volatility leads to wider no-trade regions because the cost of a suboptimal α is lower. This explains why no-trade regions are wider when $b_r\sigma_d$ increases because this increase in $b_r\sigma_d$ causes conditional single-period volatility to decrease, holding unconditional volatility fixed.

C.2. Varying the Persistence of the Predictive Variable

The parameter b_d controls the persistence of the predictive variable, dividend yield. Figure 3C reports results for the S4–S5–S3 comparison that fixes $b_r\sigma_d$ and ρ_{ev} and varies b_d . Looking across the return processes (S4–S5–S3), the average no-trade width (unreported) does not exhibit any clear pattern. However, the graph reveals that no-trade width is a very different function of the dividend state depending on the magnitude of b_d . In particular, as b_d decreases, going from S4 to S5 to S3, we see that the no-trade width goes from being flat to U-shaped as a function of the dividend state. In other words, as the predictive variable becomes less persistent, the investor is less inclined to rebalance when confronted with an extreme value for the dividend yield, and so the no-trade region widens for extreme d states. This disinclination to rebalance is understandable because a less persistent predictive variable implies that expected return is also less persistent. Consequently, the benefits of rebalancing in an extreme state are less likely to outweigh the cost because expected return reverts more quickly to its unconditional value.

C.3. Allowing the Risky-Asset Return to Exhibit Heteroskedasticity

Figure 4 reports results for return-generating processes that allow returns to exhibit heteroskedasticity: the slope for conditional volatility as a linear function of conditional mean, ζ_1 , is positive and increasing going from H2 to H1 to H0.¹¹ As discussed in Section III.A.2, the cost of a suboptimal α is increasing in the conditional return volatility. This reasoning implies that no-trade width is decreasing in conditional return volatility across states. Thus, we expect the slope of the no-trade width as a function of the state to become less positive going from S0 to H2 to H1 to H0 because ζ_1 is increasing going across those states. Figure 4A confirms this result.

Because ours is the first paper to consider the effect of return heteroskedasticity on portfolio choice in a multiperiod setting, we also report the no-trade midpoint by state in Figure 4B. As ζ_1 increases, the slope of the conditional Sharpe ratio as a function of the state becomes less positive or more negative. This implies that the slope of the no-trade midpoint as a function of the state decreases going from S0 to H2 to H1 to H0. Again, Figure 4B confirms this result except perhaps for extreme dividend states. In fact, conditional volatility in H0 increases so rapidly as a function of the conditional mean that the no-trade midpoint is a negative function at all but the extreme states. Thus, return heteroskedasticity can potentially have a big impact on rebalancing behavior. At the same time, it worth noting that the average cost incurred (unreported) is largely unaffected as ζ_1 increases going from S0 to H2 to H1 to H0.

IV. Concluding Remarks

Transaction costs and return predictability are realistic features of the environment facing a U.S. investor or mutual fund. How they jointly affect portfolio choices by entities with long horizons is not well understood. This paper considers the impact of transaction costs on the portfolio decisions of a long-lived agent with isoelastic preferences. A particular focus is how portfolio choice and rebalancing behavior are affected by return predictability. A number of extensions are feasible and of interest. The current paper has only two assets and two types of transaction costs. Extending the framework to consider portfolio choice with multiple assets and a variety of cost structures would be interesting. The techniques in this paper could also be used to consider portfolio choices by fund managers given the compensation functions that they face. Finally, the framework could be extended to assess the value of a piece of information, and how to rebalance a portfolio given that information, in a setting with transaction costs.

¹¹ Recall that conditional mean return in each state is held constant across return-generating processes by construction.

REFERENCES

- Balduzzi, Pierluigi, and Anthony Lynch, 1999, Transaction costs and predictability: Some utility cost calculations, *Journal of Financial Economics* 52, 47–78.
- Barberis, Nicholas, 1999, Investing for the long run when returns are predictable, *Journal of Finance*, forthcoming.
- Bodie, Zvi, Robert C. Merton, and William F. Samuelson, 1992, Labor supply flexibility and portfolio choice in a life cycle model, *Journal of Economic Dynamics and Control* 16, 427–449.
- Brandt, Michael, 1999, Estimating portfolio and consumption choice: A conditional Euler equations approach, *Journal of Finance*, forthcoming.
- Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, Strategic asset allocation, *Journal of Economic Dynamics and Control* 21, 1377–1403.
- Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373–399.
- Campbell, John Y., and Luis M. Viceira, 1996, Consumption and portfolio decisions when expected returns are time varying, Working paper, Harvard University.
- Constantinides, George, 1986, Capital market equilibrium with transaction costs, *Journal of Political Economy* 94, 842–862.
- Davis, M. H. A., and A. R. Norman, 1990, Portfolio selection with transaction costs, *Mathematics of Operations Research* 15, 676–713.
- Duffie, Darrell, and Tong-sheng Sun, 1990, Transaction costs and portfolio choice in a discrete-continuous-time setting, *Journal of Economic Dynamics and Control* 14, 35–51.
- Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.
- Gennotte, Gerard, and Alan Jung, 1994, Investment strategies under transaction costs: The finite horizon case, *Management Science* 40, 385–404.
- Glosten, Lawrence R., Ravi Jagannathan, and David Runkle, 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance* 48, 1779–1801.
- Hansen, Lars Peter, and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.
- He, Hua, and David M. Modest, 1995, Market frictions and consumption-based asset pricing, *Journal of Political Economy* 103, 94–117.
- Heaton, John, and Deborah Lucas, 1996, Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* 104, 443–487.
- Heaton, John, and Deborah Lucas, 1997, Market frictions, savings behavior and portfolio choice, *Macroeconomic Dynamics* 1, 76–101.
- Kandel, Shmuel, and Robert Stambaugh, 1996, On the predictability of stock returns: An asset allocation perspective, *Journal of Finance* 51, 385–424.
- Kim, Tong Suk, and Edward Omberg, 1996, Dynamic nonmyopic portfolio behavior, *Review of Financial Studies* 9, 141–161.
- Koo, Hyeng Keun, 1995, Consumption and portfolio section with labor income I: Evaluation of labor income, Working paper, Washington University.
- Jagannathan, Ravi, and Narayana R. Kocherlakota, 1996, Why should older people invest less in stocks than younger people? *Federal Reserve Bank of Minneapolis Quarterly Review* 20, 11–23.
- Luttmer Erzo, 1996a, Asset pricing in economies with frictions, *Econometrica* 64, 1439–1467.
- Luttmer, Erzo, 1996b, Consumption and portfolio choice in the presence of fixed costs, Working paper, Northwestern University.
- Merton, Robert C., 1980, On estimating the expected return on the market, *Journal of Financial Economics* 8, 323–361.
- Morton, Andrew, and Stanley Pliska, 1993, Optimal portfolio management with fixed transaction costs, Working paper, University of Illinois at Chicago.
- Samuelson, Paul A., 1969, Lifetime portfolio selection by dynamic stochastic programming, *Review of Economics and Statistics* 51, 239–246.

- Schroder, Mark, 1995, Optimal portfolio selection with fixed transaction costs, Working paper, Northwestern University.
- Scruggs, John, 1998, Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: A two-factor approach, *Journal of Finance* 53, 575–604.
- Society of Actuaries Group Annuity Valuation Table Task Force, 1995, 1994 GAM Table and 1994 GAR Table, *Society of Actuaries Transactions* 47, 865–913.
- Tauchen, George, and Robert Hussey, 1991, Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models, *Econometrica* 59, 317–396.
- Vayanos, Dimitri, 1996, Transaction costs and asset prices: A dynamic equilibrium model, Working paper, Stanford University.
- Whitelaw, Robert F., 1994, Time variations and covariations in the expectations and volatility of stock market returns, *Journal of Finance* 49, 515–541.
- Whitelaw, Robert F., 1998, Stock market risk and return: An equilibrium approach, Working paper, New York University.