NEW YORK UNIVERSITY
NYU STERN
LEONARD N. STERN SCHOOL OF BUSINESS

Department of Finance
Working Paper Series 1998

FIN-98-049

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

Anthony Lynch, Pierluigi Balduzzi

October 22, 1998

This Working Paper Series has been generously supported by a grant from

CDC Asset Management
AMERICAS

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

by

Anthony Lynch* and Pierluigi Balduzzi+

22 October, 1998.

* New York University.

+ Boston College

The authors thank John Campbell, Jennifer Carpenter, George Constantinides, Ned Elton, Silverio Foresi, John Heaton, and Robert Whitelaw, and especially the editor and referee for helpful conversations and comments. Financial support from New York University Summer Research Grants is gratefully acknowledged.

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

Abstract

We consider the impact of transaction costs on the portfolio decisions of a long-lived agent with isoelastic preferences. In particular, we focus on how portfolio choice, rebalancing frequency and average cost incurred change over the lifecycle and are affected by return predictability. Two types of costs are evaluated: *proportional* to the change in the holding of the risky asset and a *fixed* fraction of portfolio value. We find that realistic transaction costs can materially affect rebalancing behavior, creating no-trade regions that widen near the investor's terminal date. At the same time, realistic proportional and fixed costs have little effect on the midpoint of the no-trade region, unless liquidation costs differ across assets. Return predictability calibrated to U.S. stock returns is found to have large effects on rebalancing behaviour relative to independent and identically distributed (i.i.d.) returns with the same unconditional distribution. For example, 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. No-trade regions early in life are wider when returns are predictable than when they are not. Finally, we find that the nature of the return predictability, including the presence or not of return heteroscedasticity, can have large effects on rebalancing behavior.

JEL # G12

Introduction

We consider the impact of transaction costs on the portfolio decisions of a long-lived agent with isoelastic preferences. In particular, we focus on how portfolio choice, rebalancing frequency and average cost incurred change over the lifecycle and are affected by return predictability. Two types of costs are evaluated: *proportional* to the change in the holding of the risky asset and a *fixed* fraction of portfolio value. Transaction costs, investor characteristics and the return-generating process are all calibrated to the U.S. experience.

We have several major results. 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, realistic proportional and fixed costs have little effect on the midpoint of the no-trade region, unless liquidation costs differ across assets. We find that rebalancing behavior changes over the investor's life, particularly near the terminal date when no-trade regions widen, rebalancing frequency drops, and less is spent on rebalancing. The presence of predictability calibrated to U.S. equities is found to have a big effect on rebalancing behavior. In particular, at all points in the investor's life, return predictability causes rebalancing frequency to increase, and cost incurred to increase by an order of magnitude. In addition, no-trade regions early in life are wider when returns are predictable than when they are not.

The nature of the return predictability can also have a large effect on portfolio choice. With the unconditional distribution held fixed, increasing the magnitude of the single-period predictability causes no-trade regions to widen. On the other hand, increasing the persistence of the predictive variable causes no-trade width, as a function of the predictive variable, to become more u-shaped; i.e., the no-trade region is wider for extreme values of the predictive variable. Finally, we calibrate return heteroscedasticity to an equilibrium model and find that its presence has a big impact on the slopes of no-trade width and no-trade midpoint expressed as functions of expected return.

Portfolio choice in a multiperiod setting with realistic transaction costs and realistic return predictability is of interest for several reasons. First, a mutual fund manager typically has an horizon that is much shorter than that of investors (see, for example, Brown, Harlow and Starks (1996) who discuss the calendar year focus of fund managers). Assessing how this shorter horizon affects the manager's portfolio choices relative to the investor's optimal choices allows the cost of this shorter horizon for the investor to be qualitatively evaluated. This cost is likely to be sensitive to the presence of transaction costs and return predictability.

Second, transaction costs have been advanced as a potential explanation of the equity premium puzzle (see He and Modest (1995) and Luttmer (1996a and b)).[1] The ability of transaction costs to explain this puzzle depends on the frequency of portfolio rebalancing and the form of the rebalancing rule that they induce. The presence of return predictability is likely to affect these dimensions of portfolio choice.

Third, portfolio rebalancing rules in the presence of predictability and transaction costs are of interest in their own right. Return predictability using publicly available instruments has been well documented, and possession of private information implies an ability to predict returns. Utilizing predictability to increase utility, especially in the presence of transaction costs, is a difficult problem, but the gains from doing so are potentially large.[2] In particular, we would like to understand the impact of varying the form of the predictability. Doing so will help us understand how predictability and transaction costs jointly affect portfolio choices.

Portfolio choice with realistic transaction costs and no return predictability is also of interest. In his seminal 1969 paper, Samuelson disproves the notion that different portfolio choices are appropriate at different stages of an investor's life. Samuelson's argument is cast in a framework where the investor exhibits isoelastic preferences (constant relative risk aversion CRRA) and faces a constant investment-opportunity set with zero transaction costs.

[1]The equity premium puzzle says that aggregate consumption is too smooth and not sufficiently correlated with equity return to explain the magnitude of the premium on U.S. equities (relative to T-bills) in the context of the time-separable representative agent model with a reasonable risk aversion parameter (Mehra and Prescott (1985)).

[2]See Kandel and Stambaugh (1996) who discuss the gains in a one period setting, Brennan, Schwartz and Lagnado (1996), and Barberis (1997) who discuss the gains in a multiperiod setting, and Balduzzi and Lynch (1998) who discuss the gains in a multiperiod setting with and without transaction costs.

An interesting question is whether his result is robust to the presence of transaction costs.[3]

The impact of proportional and fixed transaction costs on the portfolio choice of an infinite-lived CRRA investor facing a constant opportunity set is well-understood. With proportional transaction costs, Constantinides (1986) and Davis and Norman (1990) find that the optimal rebalancing rule involves a no-trade region with an upper and lower bound for the holding of the risky asset in the portfolio. When the holding goes outside the no-trade region, the investor rebalances to bring the holding back to the closest bound. With fixed transaction costs alone, Schroder (1995) and Morton and Pliska (1993) also find a no-trade region, but rebalancing puts the risky-asset holding inside the region.[4]

At least two papers have addressed the issue of portfolio choice over the lifecycle in the presence of transaction costs. Schroder shows how a longer horizon reduces the no-trade region and increases rebalancing frequency in the presence of fixed costs. Gennotte and Jung (1994) show how proportional transaction costs imply no-trade regions that widen near the end of an investor's life. However, neither of these papers considers fixed and proportional costs simultaneously, intermediate consumption, or, most importantly, return predictability.[5]

A long line of literature has established that returns are predictable.[6] 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

[3]Recently, other authors have qualified Samuelson's result by adding realistic features to his set-up. Bodie, Merton, and Samuelson (1992) and Jagannathan and Kocheralokota (1996), for example, show how portfolio allocations to the risky asset may decline over an investor's lifetime because of human capital.

[4]Duffie and Sun (1990) analyze the impact of transaction costs that are a fixed fraction of total portfolio value for consumers who decide the interval of time until their next rebalancing at the time of their current rebalancing.

[5]There are a number of related papers that consider the effects of transaction costs on portfolio choice and market equilibrium, and on the individual's first order condition. Heaton and Lucas (1996), Vayanos (1996) and Koo (1991) investigate the equilibrium implications of proportional transaction costs and other frictions. Hansen, Heaton, and Luttmer (1995), Luttmer (1996a), and He and Modest (1995) develop pricing-operator tests of asset-pricing models which explicitly account for bid-asked spreads and short-sale constraints. Luttmer (1996b) explores the impact of fixed transaction costs on the joint distribution of consumption and asset returns. None of these papers studies how portfolio choice changes over an individual's life when she faces transaction costs and return predictability of the magnitude observed in the U.S. economy.

[6]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.

a myopic setting, where the investor rebalances monthly. Brennan and Schwartz (1996), Brennan, Schwartz, and Lagnado (1996) and Barberis (1997) analyze numerically the impact of myopic versus dynamic decision-making, when the consumer is allowed to use 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 an analogous continuous-time problem but without intermediate consumption.

None of these papers simultaneously consider the impact of return predictability and transaction costs on portfolio choice. The only paper to do so is Balduzzi and Lynch (1998) whose focus is on the utility costs associated with ignoring predictability and behaving myopically. The focus in this current paper is on the form of the rebalancing rule, average portfolio holdings, average spending on transaction costs and rebalancing frequency over the lifecycle. Also, the current paper assesses the impact on portfolio choice of varying the return predictability, while the earlier paper only considers one form of predictability calibrated to U.S. equities. Finally, the current paper considers the impact on portfolio choice of return heteroskedasticity, realistic death probabilities and differential liquidation costs across assets. None of these are considered in the other paper.

We solve the consumer's problem numerically and calibrate asset returns to the U.S economy using the quadrature approximation techniques of Tauchen and Hussey (1991) as applied by Balduzzi and Lynch. The opportunity set consists of two assets: a riskless and a risky asset. When returns are allowed to be predictable, the economy has one state variable designed to capture the predictive ability of dividend yield for U.S. equity returns. The return-generating process is calibrated to the predictibility of U.S. equity return, but other forms of predictabiliy 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 of the paper can be summarized as follows.

No-trade regions widen dramatically close to the terminal date. For an investor facing a proportional cost of 50 basis points and no fixed cost, the no-trade region does not start widening (more than 0.05 wider) until the last year of life. Adding a fixed cost of 0.01% of portfolio value only causes the widening to start 6 months earlier.[7] There is good intuition for why no-trade regions widen late in life. Rebalancing early in life provides benefits over several subsequent periods, but rebalancing late in life can only provide benefits until the investor dies. Thus, the investor is more inclined to rebalance early in life than late in life, which implies smaller no-trade regions early in life.

Return predictability calibrated to U.S. returns has large effects on rebalancing behaviour relative to independent and identically distributed (i.i.d.) returns with the same unconditional distribution. In particular, 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. Further, no-trade regions early in life are wider when returns are predictable than when they are i.i.d.. As an example, consider the rebalancing behavior early in life of an investor facing a proportional cost of 50 basis points and a fixed cost of 0.01% of portfolio value. When facing the unconditional, i.i.d. return, the investor rebalances once every 43 months, incurs an average monthly rebalancing cost of 0.106 basis points of portfolio value and uses an average no-trade width of 0.180. On the other hand, with predictability calibrated to U.S equities, the same investor, early in life, rebalances once every 11 months, incurs an average monthly rebalancing cost of 0.898 basis points of portfolio value and uses an average no-trade width of 0.295. Predictability increases the benefits from incurring transaction costs and rebalancing, which explains why the investor rebalances more frequenctly and spends more on transaction costs. The widening of the no-trade region can be explained by two effects. First, cross-state variation in the position of the no-trade region means that this period's trade may be reversed next period. Second, predictability reduces conditional volatility which reduces the cost of a suboptimal risky-asset holding (see Constantinides (1996) and Gennotte and Jung (1994) who discuss this point but in the context of the volatility of i.i.d. returns).

[7]Gennotte and Jung report a similar result but for the more restrictive case where the risky asset return is a binomial process.

We are also able to characterize hedging demand over the lifecycle. When predictability is calibrated to U.S equities, average risky-asset holding and no-trade midpoint both increase moving backward from the terminal date. This decline in risky-asset holding as the terminal date approaches indicates that hedging demand decreases as the investor's horizon becomes shorter.

Rebalancing behavior is sensitive to investor characteristics. Increasing the investor's risk aversion coefficient causes average risky-asset holding and no-trade midpoint to drop considerably, while no-trade width is also lowered. These results are consistent with the high risk-aversion investor caring more about the dispersion of wealth. The absence of intermediate consumption causes the no-trade region to widen later in life. Further, when returns are predictable, average risky asset holding and no-trade midpoint rise more rapidly and to a higher level moving backwards from the terminal date. By increasing future wealth relative to current wealth at any given time τ, the absence of consumption makes the future more important (which increases hedging demand), and makes the investor's time-τ problem less like the single-period problem.

The presence of a realistic fixed cost causes rebalancing frequency to decline considerably. For example, when a fixed cost of 0.01% of portfolio value is added to a 50 basis point proportional cost, rebalancing frequency drops from once every 4.7 months to once every 43 months for i.i.d returns and from once every 3.7 months to once every 11 months for return predictability calibrated to U.S. equities.[8] However, realistic proportional and fixed costs have little effect on average risky-asset holding and, in particular, no-trade midpoint unless liquidation costs differ across the two assets. Thus, the magnitude of an investor's hedging demand is robust to the presence of transaction costs.

On the other hand, no-trade regions late in life are lowered when the investor without intermediate consumption faces a liquidation cost on the risky but not the riskless asset. In fact, the upper boundary of the no-trade region converges to the risky-asset holding in the

[8]These results are consistent with those of Schroder who found that a small fixed cost (but *no* proportional cost) could cause rebalancing frequency to drop dramatically from the no-cost case.

absence of transaction costs. The lowering of the no-trade region late in life is understandable since a reduction in the risky-asset holding involves a cost that must otherwise be incurred at the terminal date. Thus, differential liquidation costs across the two assets can have a big impact on rebalancing behavior, consistent with recent work by Heaton and Lucas (1997).

The nature of the return predictability can have large effects on rebalancing behavior. We focus on three return parameters: the magnitude of the single-period predictability; persistence of the predictive variable; and, the contemporaneous correlation between return and the predictive variable (which we refer to as the hedgeability parameter). Holding the unconditional distribution fixed, increasing the magnitude of the single-period predictability causes no-trade regions to widen. The same two effects that cause wider no-trade regions when returns are predictable also cause no-trade regions to widen when the magnitude of the single-period predictability increases: increased cross-state variation in the position of the no-trade region which increases the likelihood that this period's trade will be reversed next period; and, reduced conditional volatility which reduces the cost of a suboptimal risky-asset holding.

Decreasing the persistence of the predictive variable causes no-trade width, as a function of the predictive variable, to become more u-shaped; i.e., the no-trade region is wider for extreme values of the predictive variable. This result makes sense because an investor facing an extreme expected return is less likely to rebalance if the extreme value reverts quickly to its unconditional mean: the benefits of rebalancing are less likely to outweigh the cost.

However, the correlation between return and the predictive variable does not appear to have a first order effect on the no-trade width. This result is a little surprising since the magnitude and sign of the hedgeability parameter determines the magnitude and sign of return autocorrelation. Intuition suggests that negative (positive) return autocorrelation narrows (widens) the no-trade region. The idea is that, when return autocorrelation is negative, a high risky-asset holding is likely to be followed by a negative return that makes next period's holding less extreme. But for a given value of the predictive variable, the expected return next period is the same irrespective of this period's return or risky-asset holding. Since portfolio decisions are made knowing the predictive variable, the intuition does not apply. Instead, the magnitude of the single-period predictability is the most important determinant of no-trade width. In fact, only one process we consider has no-trade widths that are below the unconditional no-trade width in some states. This return-generating process has low single-period predictability and a negative value for the hedgeability parameter.

Finally, return heteroscedasticity 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 no-trade midpoint to be less positive (or more negative) functions of expected return. These two effects are explained in turn. As discussed above, higher volatility increases the cost of a sub-optimal risky-asset holding. Thus, an increase in conditional volatility going from one state to another is associated with a reduction in no-trade width. With respect to no-trade midpoint, the investor's risk aversion causes the no-trade midpoint to be a decreasing function of conditional volatility.

Thus, realistic transaction costs and return predictability can have a significant impact on the way individuals rebalance their portfolios over their lives. The implication is that the simultaneous presence of transaction costs and return predictability is likely to affect the joint distribution of consumption and equity return. Thus, calibration studies using transaction costs to explain the equity premium may be sensitive to the introduction of predictability in returns.

Further, our results on rebalancing frequencies have implications for existing studies (He and Modest and Luttmer (1996a)) documenting how proportional transaction costs weaken Hansen-Jagannathan (1991) bounds on intertemporal marginal rates of substitution (IMRSs). These bounds follow from necessary conditions on IMRSs in equilibrium and become weaker relative to the frictionless bound as the data frequency becomes higher. Consequently, they are likely to be uninformative for data frequencies that are too much higher than the rebalancing frequency of individuals. For example, saying an IMRS lies inside the one-month transaction cost bound may not be very informative if investors are rebalancing every year.

Thus, our results shed some light on the informativeness of the bounds in these papers. For example, Luttmer finds, using quarterly data, that a proportional cost of 0.5% causes the bound to be within one standard error of zero. With predictability and a proportional cost of 0.5%, we find that our investor rebalances every 3.7 months which is only slightly less frequently than every quarter. While our reported rebalancing frequency is averaged across states and starting risky asset weights, it suggests that the bounds reported by Luttmer are informative.[9]

Further, our results indicate that the shorter horizons of fund managers (relative to those of investors) reduce their incentive to trade. An overconfident manager is likely to trade too much, leading to the payment of excessive transaction costs. However, by reducing the manager's incentive to trade, the shorter horizon limits the payment of unnecessary transaction costs by the fund, holding all else equal. So while the shorter horizon of a fund manager is likely to have a number of disadvantages, its ability to discourage trading may be an advantatge that has not been emphasized before.

Finally, our results for the case of unpredictable returns indicate that Samuelson's (1969) irrelevance result is not robust to the introduction of transaction costs. Even so, we find that the average allocation to the risky asset tends to remain constant over the investor's lifetime, unless liquidation costs differ across assets. In this limited sense, Samuelson's (1969) original intuition extends to the case of transaction costs.

A number of recent papers emphasize the importance of non-financial wealth particularly early in life (see, for example, Koo, 1995, Bodie, Merton and Samuelson, 1992, and Heaton and Lucas, 1997). Because of the potentially large impact of non-financial wealth, care must be taken when applying our results to a young investor. The impact of non-financial wealth is likely to depend on its correlation with risky-asset return and its impact on the transaction costs faced by the investor. For example, if most of an investor's wealth is non-financial, then a fixed cost proportional to the investor's financial wealth understates the cost of the

[9]Our results are only suggestive since the rebalancing interval our investor faces depends on the sequence of return realizations since the last rebalancing.

investor's time. However, our results are directly applicable to an investor who has reached retirement age: the investor could be 80 years and 1 month old in the first period and 101 years old at the terminal date.

The paper is organized as follows. Section I describes the optimization problem that we solve. Section II calibrates asset returns to the U.S. economy, transaction costs to those faced by U.S. investors, and the investor's profile to that of a U.S. investor. Section III describes the technique used to solve the investor's problem and discusses the details of the simulations used to obtain rebalancing frequencies, average holdings and average transaction costs incurred. Section IV presents the results while Section V concludes.

I. Optimal Portfolio Allocation with Transaction Costs

This section lays out the constraints and preferences faced by the consumer and characterizes the optimization problem.

A. *Constraints and Preferences*

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

We consider the optimal portfolio problem of a consumer with a finite life of T periods. Preferences are assumed to be of the constant-relative-risk-aversion type (CRRA) and time separable with a rate of time preference equal to β. Expected lifetime utility is given by

$$E\left[\sum_{t=1}^{T}\beta^t \frac{c_t^{1-\gamma}}{1-\gamma}|\mathbf{Z}_1\right] \tag{1}$$

where c_t is consumption at time t, γ is the relative-risk-aversion coefficient and $\mathbf{Z}_t$ is the vector of state variables for the consumer at time t. When transaction costs are non-zero, the inherited portfolio allocation $\hat{\alpha}_t$ is an element of $\mathbf{Z_t}$, since 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 $\mathbf{Z_t}$.

The formulation in (1) allows the investor to live until the terminal date with probability 1. A death probability can be incorporated by generalizing the investor's objective function:

$$E\left[\sum_{t=1}^{T}(\Pi_{\tau=1}^{t}\beta_\tau)\frac{c_t^{1-\gamma}}{1-\gamma}|\mathbf{Z_1}\right] \tag{2}$$

where $\beta_\tau = \beta(1-p_t)$ and p_t is the probability of dying at time t given that the investor survived until $t-1$.

Since other studies of optimal portfolio choice have considered situations where expected lifetime utility only depends on terminal wealth (for example, Brennan, Schwartz, and Lagnado, 1996), we also consider the following preferences:

$$E\left[\frac{W_T^{1-\gamma}}{1-\gamma}|\mathbf{Z_1}\right] \tag{3}$$

where W_t is the investor's wealth at time t. Thus, we can assess the sensitivity of our results to allowing utility to depend on intermediate consumption.

The law of motion of the consumer's wealth, W, is given by

$$W_{t+1} = (W_t - c_t)(1-f_t)\left[\alpha_t(R_{t+1} - R^f) + R^f\right] \tag{4}$$

where α is the share of wealth allocated to the risky asset, and f_t is the transactions cost paid at t per dollar of wealth (which might be zero if no trading occurs at t). 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

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

In general, we model the cost of transacting f as a function of the difference between the chosen risky asset allocation,α, and allocation to the risky asset inherited from the previous period, $\hat{\alpha}$. The inherited allocation satisfies:

$$\hat{\alpha}_t \equiv \frac{\alpha_{t-1}(1-\kappa_{t-1})W_{t-1}(1-f_{t-1})R_t}{W_t} = \frac{\alpha_{t-1}R_t}{\alpha_{t-1}(R_t - R^f) + R^f}. \tag{6}$$

Specifically, we assume that transaction costs have *two* components, the first *proportional* to the change in the holding of the risky asset, and the second a *fixed* fraction of the total value of the portfolio:[10]

$$f = \phi_1|\alpha - \hat{\alpha}| + \phi_2 I_{\alpha-\hat{\alpha}\neq 0},\ \phi_1, \phi_2 \geq 0 \tag{7}$$

where $I_{\alpha-\hat{\alpha}\neq 0}$ is an indicator function which equals one if $\alpha - \hat{\alpha} \neq 0$, and equals zero otherwise. 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, since it is likely to depend on the opportunity cost of the investor's time.

The law of motion for wealth in (4) implicitly assumes that consumption at time t is obtained by liquidating costlessly the risky and the riskless asset in the proportions $\hat{\alpha}_t$ and $(1-\hat{\alpha}_t)$. In the case of the terminal-wealth investor in (3), it is sufficient if the proportional liquidation cost is the same for the risky and the riskless assets. Further, transaction costs at time t are based on portfolio value after time-t consumption, and are paid by costlessly liquidating the risky and the riskless assets in the proportions α_t and $(1-\alpha_t)$.

The costless-consumption 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. However, portfolio rebalancing rules are likely to be highly sensitive to any cost differential associ-

[10]In order to keep notation simple, when all variables in a mathematical expression are contemporaneous we drop the time subscript.

ated with using the risky rather than the riskless asset for consumption (see, for example, Heaton and Lucas, 1997). The impact of relaxing this assumption is assessed by allowing the terminal-wealth investor in (3) to pay a proportional cost to liquidate the risky asset at the terminal date. In particular, we assume proportional costs of ϕ_T and 0 to liquidate the risky and riskless assets respectively. Thus, transaction costs paid at the terminal date are given by

$$f_T = \phi_T \hat{\alpha}_T. \tag{8}$$

B. *Optimization Problem*

We now consider the consumer's optimization problem. Given our parametric assumptions, the Bellman equation faced by the consumer is given by[11]

$$\frac{a(\mathbf{Z_t}, t)W_t^{1-\gamma}}{1-\gamma} = \max_{\kappa_t, \alpha_t} \frac{\kappa_t^{1-\gamma} W_t^{1-\gamma}}{1-\gamma} + \beta_t \frac{(1-\kappa_t)^{1-\gamma} W_t^{1-\gamma}}{1-\gamma} E\left[a(\mathbf{Z_{t+1}}, t+1) R_{W,t+1}^{1-\gamma} | \mathbf{Z_t}\right], \quad \text{for } t = 1, \ldots, T-1 \tag{9}$$

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:

$$\frac{a(\mathbf{Z_t}, t)}{1-\gamma} = \max_{\kappa_t, \alpha_t} \frac{\kappa_t^{1-\gamma}}{1-\gamma} + \beta_t \frac{(1-\kappa_t)^{1-\gamma}}{1-\gamma} E\left[a(\mathbf{Z_{t+1}}, t+1) R_{W,t+1}^{1-\gamma} | \mathbf{Z_t}\right], \quad \text{for } t = 1, \ldots, T-1 \tag{10}$$

The Bellman equation (10) is solved by *backward iteration*, starting with $t = T-1$. Thus, $a(\mathbf{Z_t}, t)$ is obtained by solving the optimization problem in (10) using $a(\mathbf{Z_{t+1}}, t+1)$ from the previous iteration. In the first iteration, $a(\mathbf{Z_T}, T)$ is taken to be 1 for the intermediate-

[11]This form of the value function derives from the CRRA utility specification in (1) and (3), and from the linearity in W of the budget constraint, equation (5).

consumption investor in (2). For the terminal-wealth investor in (3),

$$a(\mathbf{Z_T}, T) = (1 - \phi_T \hat{\alpha}_T)^{1-\gamma}. \tag{11}$$

This specification for $a(\mathbf{Z_T}, T)$ collapses to the usual $a(\mathbf{Z_T}, T) = 1$ when the cost of liquidating the risky asset is set equal to 0; i.e., $\phi_T = 0$.

II. Calibration Details

In this section, we describe how the model is calibrated to U.S. data along three dimensions: transaction costs; investor profile; and, return-generating process.

A. *Transaction Costs*

Given our transaction cost function in (7), one can observe that the proportional component of transaction costs is given by $\phi_1|\alpha - \hat{\alpha}|W_p$ where W_p is the portfolio value after consumption. The cost of changing the holdings of the risky asset to αW_p, and then of changing them back to $\hat{\alpha} W_p$, amounts to $2\phi_1|\alpha - \hat{\alpha}|W_p$. 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 traded. We allow ϕ_1 to take values of 0.5% and 2.5%. Lesmond, Trzcinka, and Ogden (1996) calculate average round-trip proportional transaction costs to be 1.2% and 10.3% for large and small firms respectively. These estimates bracket the ϕ_1 values that we use. In particular, the low ϕ_1 value of 0.5% is consistent with their average cost of trading a large firm while a ϕ_1 of 2.5% implies a round-trip cost of 5%. Thus, the high ϕ_1 value is chosen to lie midway between the average cost of trading large and small firms. Their measure of transaction costs is typically 15% to 20% smaller than the spread-plus commissions estimator employed by Stoll and Whaley (1983) and Bhardwaj and Brooks (1992), suggesting that their estimates are unlikely to overstate the actual costs.

In regard to the fixed transaction costs parameter ϕ_2, the lower value of 0.01% that we use falls within a range of fixed-cost lower bounds that make observed U.S. per-capita consumption consistent with asset returns (see Luttmer (1996b)).[12] Further, this value of ϕ_2 translates into paying a fee of $10 whenever a $100,000 portfolio is reshuffled. Viewed as the opportunity cost for an individual investor to process information and instruct a broker to change portfolio composition, this value for ϕ_2 seems small. Thus, the actual fixed costs faced by an investor probably lie between the high and low values of 0.1% and 0.01% that we use.

The three possible ϕ_1 values of 0, 0.5%, and 2.5%, are combined with three possible ϕ_2 values of 0, 0.01%, and 0.1%, to give nine transaction cost combinations, which are listed in Panel A of Table 1. The first two characters of the label identify the proportional cost with P0, P1 and P2 corresponding to ϕ_1 values of 0, 0.5%, and 2.5%, respectively. The label's last two characters identify the fixed cost parameter with F0, F1 and F2 corresponding to ϕ_2 values of 0, 0.01%, and 0.1%, respectively. Thus, the nine labels P0F0, P0F1, P0F2, P1F0, P1F1, P1F2, P2F0, P2F1, P2F2, correspond to the nine possible $[\phi_1, \phi_2]$ combinations (expressed in %) [0,0], [0,0.01], [0,0.1],[0.5,0], [0.5,0.01], [0.5,0.1],[2.5,0], [2.5,0.01], [2.5,0.1]. The P1F0 setting is usually used when other aspects of the investor's problem are varied.

B. *Investor Profiles*

The various investor profiles considered are listed in Panel B of Table 1. The standard investor, who is denoted "si", has a relative risk aversion coefficient γ equal to 4, receives utility from intermediate consumption, has a zero probability of death prior to the terminal date, and pays a zero cost to liquidate either asset for consumption. The "si" investor's first and last decisions are made at $t = 1$ and $t = 239$ respectively: so the investor lives for 240 months. The rate of time preference for the "si" investor is taken to be $1/R^f$.

The standard investor is perturbed along a number of dimensions. The high risk-aversion investor denoted "ha" is identical to the standard investor except γ is set equal to 7.

We also consider the impact of a non-zero probability of death prior to the terminal date.

[12]Further details are available from the authors on request.

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). The reported mortality rate for a male of age t is the probability that the male dies during year t, given that he is alive at the start of year t. Similarly, the monthly death rate at time t is the conditional probability of death between t and $t+1$ given that the investor is alive at time t. The reported annual rates are converted to monthly death rates assuming a constant monthly death rate over the year.

Since we do not model labor income, we take male mortality rates for the retirement ages of 81 onward. The "pd" investor is 81 years and 1 month at time $t = 1$. GAM-94 mortality rates are used until $t = 240$ at which point the "pd" investor dies with probability 1. In other words, upon reaching 101 years of age, the "pd" investor dies with probability 1. The GAM-94 annual mortality rate starts off at 7.378% for the 81 year-old man and rises to 34.112% for the 100 year old man. Thus, the monthly death rate starts off at 0.637% for the first 11 months and finishes at 3.417% for the last 12.[13] On the other hand, the "dl" investor experiences the GAM-94 mortality rates through until $t = 468$ or age 120, which is the year that the mortality rate becomes 1 in GAM-94. The GAM-94 annual mortality rate has climbed to 50% by year 119.

The "nc" investor is identical to the standard investor except the investor's utility only depends on terminal wealth. The final consumer profile is identical to the "nc" profile except that the entire proportional transaction cost of ϕ_1 must be paid to liquidate the risky asset. The cost of liquidating the riskless asset is zero and the profile is denoted "cl". Portfolio rebalancing rules for the "nc" and "cl" profiles can be compared to assess the impact of imposing a liquidation cost to consume from the risky asset.

[13]Regarding time 1 as the end of the first month of the "pd" investor's 81st year, months 1 through 11 have a probability of death derived from the year 81 mortality rate. Moving to the end of the "pd" investor's life, months 228 through 239 are the beginnings of the months in year 100, and so the year 100 mortality rate is used. Month 240 is the start of year 101 and the individual is assumed to die with probability 1 at this time.

C. *Law of Motion for Returns*

In order 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. We calibrate real returns since they are more likely to be stationary, investors generally care more about real returns, and there is no money in our model. Also, the dividend yield and stock return series are converted to a continuously compounded basis; hence, R is replaced by $r \equiv \ln(1+R)$ and the dividend yield D is replaced by $d \equiv \ln(1+D)$. Without loss of generality, we normalize the d process to be mean zero with unit variance. 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 riskfree 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%.

C..1 *Approximating the Data assuming Constant Variance*

We assume that $[r, d]'$ follows the vector autoregressive model (VAR):

$$r_{t+1} = a_r + b_r d_t + e_{t+1}, \tag{12}$$

$$d_{t+1} = a_d + b_d d_t + v_{t+1}, \tag{13}$$

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. With this normalization, $a_d = 0$ and $a_r = \mu_r$ where μ_r is the unconditional mean of r. Specification (12)-(13) assumes that d is the only state variable since all information available at t about r_{t+1} is contained in d_t. This characterization of the investment opportunity set is in line with other papers on optimal portfolio selection (e.g., Barberis, 1996; and Campbell and Viceira, 1996). The VAR is estimated using ordinary least squares (OLS).

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:

$$r_{t+1} = a_r + b_r d_t + \eta v_{t+1} + u_{t+1}, \tag{14}$$

where η is the regression coefficient from regressing e on v, and u is i.i.d. normally distributed with variance σ_u^2 and is uncorrelated with v. 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 $v_{t+1} = d_{t+1} - a_d - b_d d_t$.

This approach ensures that d is the only state variable. We chose a specification with 19 quadrature points for the dividend yield and 3 points for the innovations in stock returns. Balduzzi and Lynch (1998) also use this approach and find that the approximation is able to capture important dimensions of the predictability in the data.[14]

C..2 *Varying the Return Predictability*

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 and μ_r constant and varying $b_r\sigma_d$, b_d, and ρ_{ev}. 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, since expected return is a linear function of d. The third parameter, ρ_{ev}, measures the

[14]When we increase the number of quadrature points to 20 for dividend yield or to 6 for stock return innovations, the results for the "si" investor are virtually unchanged using the P1S0 or P1F1 cost combinations.

extent to which next period's shock to r covaries with next period's shock to expected return. It is also an increasing function of r's autocorrelation, taking a value of zero when r's autocorrelation is zero.

When using the Tauchen-Hussey approximation to discretize d, a weighting function must be chosen that, together with the chosen quadrature rule, determines the values taken by d at the 19 quadrature points. We would like the values taken by d to 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. Consequently, we take the weighting function to be the conditional distribution of d_{t+1} implied by the data VAR parameters, given d_t equal to zero. Since Tauchen and Hussey suggest using this weighting function to approximate the data VAR, this weighting function is the natural choice for all the discretizations.

Panel C of Table 1 presents VAR parameter values for both the data and the various quadrature approximations used. Mean return is not reported since this parameter is held constant across the return-generating processes and the quadrature approximation matches this moment exactly. Ten discretizations are considered and are labelled S0 through S9. Discretization S0 attempts to replicate the predictability in the data. Comparing the data VAR parameters (1st row) to the VAR parameters for S0 (2nd row) shows that the approximation is doing a good job, except that the predictable single-period variation in r ($b_r\sigma_d$) is understated. Thus, our results are likely to understate the impact of predictability on portfolio choice.

Panel C orders the discretizations to emphasize how the impact of varying each VAR parameter can be assessed. Within each comparison, parameters held fixed exhibit small variations across generating processes due to the quadrature method not exactly matching return parameters. The first four rows after the "Data" row show two pairs of comparisons: S0 vs S1; and, S2 vs. S3. Each comparison fixes σ_r, b_d, and ρ_{ev}, while allowing $b_r\sigma_d$ to vary. 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. Finally, the effect of varying ρ_{ev} from -0.9 to 0.9 is assessed with two comparisons. The first (S0 vs. S4 vs. S6) fixes b_d and $b_r\sigma_d$ at 0.962 and 0.260 to match the data, while the second (S7 vs. S8 vs. S9) sets b_d equal to zero and halves $b_r\sigma_d$ relative to the data. The parameter, $b_r\sigma_d$, is reduced to increase the number of states with no-trade regions in the interior of the interval [0,1].[15]

C..3 *Approximating the Data assuming Heteroscedasticity*

From a general equilibrium perspective, variation in expected return is typically accompanied by variation in conditional volatility. In particular, equilibrium models like Merton's (1980) dynamic CAPM imply a positive relation between expected return and conditional volatility. 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 conditional return volatility to be a linear function of d. In particular, $[r,d]'$ still satisfies (12)-(13) but $[e_{t+1}, v_{t+1}]'$ is a vector of random variables that satisfies:

$$e_{t+1} = \sigma_{e|d_t} {z^e}_t, \tag{15}$$

$$v_{t+1} = \sigma_{v|d_t} {z^v}_t, \tag{16}$$

where $[z^e, z^v]$ is multivariate normal white noise with z^e and z^v having zero means, unit variances and a correlation of $\rho_{ev|d}$, 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, which implies that it equals the unconditional volatility σ_v. On the other hand, $\sigma_{e|d_t}$ is allowed to be a linear function of d_t:

$$\sigma_{e|d_t} = \zeta_0 + \zeta_1 d_t. \tag{17}$$

[15]Notice that the unconditional volatility of r for the (S7 vs. S8 vs. S9) comparison is about 13% lower than for the other discretizations. When the approximation technique matches a VAR with $b_d = 0$, $b_r\sigma_d = 0.132\%$, $\sigma_r = 5.10\%$, and ρ_{ev} equal to 0.9 or -0.9, the resulting return generating processes, S7 and S9, have $\sigma_r = 4.818\%$. Thus, to obtain a return-generating with $\rho_{ev} = 0$, $b_d = 0$, and $b_r\sigma_d = 0.132\%$, we tried successfully with S8 to match $\sigma_r = 4.818\%$, rather than use S3.

We would like to vary ζ_1 while holding the unconditional volatility of e, σ_e, constant. To do this, ζ_0 is always calculated, for a given ζ_1, using the following equality and σ_e set equal to its value in the data:

$$\zeta_0 = \sqrt{\sigma_e^2 - \zeta_1^2}. \tag{18}$$

The quadrature approximation technique can be adapted to discretize this return-generating process. The procedure is similar except that η in (14) becomes the coefficient, $\eta(d_t)$, from a conditional regression of e_{t+1} on v_{t+1}, and the distribution of u_{t+1} that is discretized depends on d_t. In particular, for given ζ_1, σ_v, $\rho_{ev|d}$, and ζ_1 from (18), $\eta(d_t)$ satisfies:

$$\eta(d_t) = \frac{\rho_{ev|d}\sigma_{e|d_t}}{\sigma_v} = \frac{\rho_{ev|d}(\zeta_0 + \zeta_1 d_t)}{\sigma_v}, \tag{19}$$

and the conditional distribution of u_{t+1} while still normal and mean zero, has a volatility $\sigma_{u|d_t}$ given by

$$\sigma^2_{u|d_t} = \sigma^2_{e|d_t}(1 - \rho^2_{ev|d}) = (\zeta_0 + \zeta_1 d_t)^2(1 - \rho^2_{ev|d}). \tag{20}$$

This return distribution is discretized for three values of ζ_1. 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 1's Panel C reports the data's ζ_1 as 1.358%, which implies a positive relation between conditional volatility and conditional mean. Recent work using more instruments than just dividend yield has found a negative relation (see, for example, Glosten, Jagannathan and Runkle (1993), and Whitelaw (1994)). Also, while intuition suggests that the relation should be positive, a ζ_1 value of 1.358% may be too large to be plausible in an equilibrium setting. Thus, 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% and 0.604%

when the representative agent has a relative risk aversion of 2 and 40 respectively.[16]

The last three rows of Panel C report parameters for the three discretizations, H0,H1 and H2, which attempt to match ζ_1 values of 1.358%, 0.111% and 0.604%. Two points are worth noting. The unconditional volatility for r, σ_r, for each discretization is close to that in data, suggesting that the use of (18) to calculate ζ_0 is allowing us to fix e's unconditional volatility. Second, the ζ_1 value for each discretization is close to the targeted value.

III. Solution and Simulation Details

This section describes how the consumer's optimization problem is solved numerically and explains the difference between conditional and unconditional portfolio choices. We also discuss details of the simulations performed to calculate rebalancing frequency, average cost incurred and average risky-asset holding.

A. *Solution Technique*

The investor's dynamic programming problem is solved by backward recursion. 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: $\hat{\alpha} = 0, 0.02, 0.04, ..., 0.96, 0.98, 1$. In solving the optimization problem in (10), the $(t + 1)$ value function for each D_{t+1} is linearly interpolated between $\hat{\alpha}_{t+1}$ points.

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 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 for speculative purposes.

[16] Whitelaw (1988) also presents a two-regime model that reproduces the negative relation between volatility and expected return. An interesting extension to this paper solves the investor's problem given this two-regime model for returns.

The optimal α_t for a given state is found by first constraining α_t to lie on the same grid as $\hat{\alpha}$: $0, 0.02, 0.04, ..., 0.96, 0.98, 1$. Using $\alpha_t^*(1)$ to denote the first-stage optimal α_t and $g(1)$ to denote the first-stage grid size, the second-stage optimal α_t is obtained from a finer grid ($g(2) = 0.2g(1)$) constructed on the interval $[\alpha_t^*(1)-g(1), \alpha_t^*(1)+g(1)]$.[17] The second stage is repeated several times: the k-stage optimal α_t is obtained from a grid, ($g(k) = 0.2g(k-1)$), constructed on the interval $[\alpha_t^*(k-1) - g(k-1), \alpha_t^*(k-1) + g(k-1)]$.

The resulting α_t does at least as well as the first-stage optimal α_t used by Balduzzi and Lynch as the solution to the investor's problem. Further, for $\alpha_t^*(1)$ at an interior point on the interval $[0, 1]$, this procedure yields an α_t that is at least a local optimum. Finally, the numerical technique yields an approximate solution that converges to the actual solution as the $\hat{\alpha}$ grid becomes finer.[18]

B. *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.

C. *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 at each time $t, t = 1, ..., 239$. For unconditional portfolio

[17]If $\alpha_t^*(1)$ is equal to 1, the second-stage grid is constructed on the interval $[1 - g(1), 1]$. Similarly, if $\alpha_t^*(1)$ is equal to 0, the second-stage grid is constructed on the interval $[0, g(1)]$.

[18]When we increase the number of grid points from 51 to 101 for $\hat{\alpha}$, the results for the "si" investor are virtually unchanged using the P1S0 or P1F1 cost combinations.

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 procedure is repeated 100000 times and average portfolio holdings and rebalancing frequencies at each time t are obtained by averaging over the 100000 replications.

One issue is the starting values assumed for the individual's state variables at time 1. When returns are unpredictable, the only state variable is $\hat{\alpha}_1$ which is set equal to the midpoint of the no-trade region at time 1. When returns are predictable, the two state variables are the value of dividend yield, D_1, and the inherited allocation, $\hat{\alpha}_1$. For each history, the time 1 state is determined by sampling from the unconditional distribution for D_1. Once the time 1 state has been determined, $\hat{\alpha}_1$ is taken to be the midpoint of the no-trade region for that state at time 1. An ex ante concern is the sensitivity of the results to these assumptions about the starting values for the state variables. However, we find that their effect dissipates quickly, since the average holding, rebalancing frequency and average cost incurred all quickly flatten out as a function of t, typically before t is greater than 20.

IV. Results

This section reports several portfolio choice parameters for different combinations of proportional and fixed transaction costs, varying investor profiles and a variety of return-generating processes. Each table and figure reports or plots some subset of the parameters described next.

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, 1998). Both return points lie inside the no-trade region. With a fixed cost and no proportional cost, the upper and lower return points are the same (see Schroder, 1995). With only a proportional cost, optimal portfolio choice still involves a no-trade region, but the return points are now the boundaries of the no-trade region (see Constantinides, 1986, and Davis and Norman, 1990). When the investor uses the conditional return distribution (C) rather than the unconditional (U), the form of the rebalancing rule in each state is determined by the transaction cost structure as described above. However, the position of the no-trade region varies across states (see Balduzzi and Lynch, 1998).

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. 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.[19]

For each t, three additional parameters are obtained from the simulations described in section IV: rebalancing frequency, average cost incurred, and average holding of the risky asset.[20] Thus, there are six parameters all indexed by t that characterize portfolio choice by the investor. The seventh parameter is the first time t that the width of the no-trade region is at least 0.05 larger than at time 1.

A. *Standard Investor facing Returns calibrated to U.S. Data.*

Table 2 reports all seven portfolio choice parameters for the standard "si" investor facing the S0 return-generating process which is calibrated to U.S. equity return. The investor faces one of the nine combinations of proportional (ϕ_1) and fixed (ϕ_2) parameters listed in Panel A of Table 1. Parameter values are reported for both unconditional (U) and conditional (C) portfolio choices. For each parameter, the first three rows apply to unconditional choices while the last three refer to conditional choices.

Table 2 reports time-series averages over two periods in the investor's life for the six

[19]To ensure that the short-sale restriction is not understating the average no-trade width or return distance, the averaging is performed across states for which the short-sale restriction is not binding. On the other hand, since bounds of 0 or 1 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.

[20]Risky-asset holding is determined after any rebalancing has occurred.

parameters indexed by t: rebalancing frequency, average cost incurred, average holding of the risky asset, time-t midpoint, time-t width, and, time-t return distance. The first three columns, labelled "Early", average the relevant parameter value from $t = 96$ to $t = 119$ while the last three columns are labelled "Late" and average from $t = 216$ to 239.

For a given time-series average ("Early" or "Late") and return distribution (U or C), parameter values are contained in a 3 by 3 subtable. For a given subtable, the fixed cost parameter is constant across a row while the proportional cost parameter varies from 0% to 2.5% going from the first to the third column. Conversely, the proportional cost parameter is constant down a column while the fixed cost parameter varies from 0% to 0.01% going from the first to the third row.

The first row of each subtable shows the three combinations with a zero fixed cost: P0F0, P1F0, and P2F0. Consistent with previous literature, the return distance always equals zero for these cost combinations irrespective of whether the investor is early or late in life, or whether the conditional or unconditional distribution is being used. On the other hand, the first column of each subtable shows the three cost combinations with a zero proportional cost: P0F0, P0F1, and P0F2. Consistent with Schroder (1995)'s finding of a single return point, the return distance always equals half the no-trade width for these cost combinations with a zero proportional cost. Finally, the remaining cells of each subtable have nonzero proportional and fixed costs. For these cells, the return distance both early (reported) and late (unreported) is always positive but less than half the no-trade width.[21]

Figures 1 through 4 plot parameters as a function of t for three cost combinations: P1F0, P0F1, and P1F1. Each figure contains three graphs, one for each cost combination while Figure 4 contains a fourth graph for the case of no transaction costs, P0F0. Each graph plots the relevant parameters for the "si" investor using the unconditional as well as the conditional distribution for risky asset return. Figure 1 plots average rebalancing frequency

[21]When one or both no-trade boundaries hit the short-selling constraints in all states at time t, the no-trade width is taken to be 1. However, determining the time-t return distance becomes problematic. The widening of the no-trade region late in life means that the no-trade boundaries start hitting the short-selling constraints in an increasing number of states as t approaches 240. Eventually, the short-selling constraint is hit in every state which explains why the return distance is not reported late in life.

while Figure 2 plots the average transaction cost incurred per month. Each of the three graphs in Figure 3 plots both the no-trade width and the return distance while the no-trade midpoint and average holding of the risky asset are plotted in Figure 4.

A..1 Portfolio Choice over the Lifecycle

The first question that we address is portfolio choice over the lifecyle. Table 2 shows that both average rebalancing frequency and average cost incurred are lower "Late" in the investor's life than "Early". Facing the P1F1 cost combination and the C information set, the investor's rebalancing frequency drops from 8.85% early in life to 6.36% averaged over the final 24 months, while the average cost incurred expressed in basis points of the portfolio's value drops from 0.898 to 0.700. The bottom graphs in Figures 1 and 2 show that these decreases are particularly pronounced just prior to the terminal date when the rebalancing frequency is below 2.5% and the average cost incurred is below 0.35 basis points.

Consistent with this result, we find that the no-trade width is always much larger "Late" than "Early". Table 2 shows that the average no-trade width increases from 0.295 to 0.403 going from "Early" to "Late" for the investor facing P1F1 costs and using C. Figure 3 confirms this result showing that the no-trade width exceeds 0.75 at time $t = 239$ for the investor facing P1F1 costs and using C.

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.

This intuition suggests the widening is not likely to occur until t is close to the terminal date. The last panel of Table 2 confirms this, showing that the time of widening of the no-trade region is always later than $t = 210$ except for $\phi_1 = 2.5\%$. For P1F1, the time of widening is $t = 221$ for U and $t = 230$ for C.

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 a longer period of time between rebalancing is needed to make rebalancing now attractive. This implication is confirmed in Table 2. Rebalancing frequency for a given cost combination is lower for U than C policies, while for a given information set, rebalancing frequency decreases moving across a row (ϕ_1 increasing, ϕ_2 fixed) or down a column (ϕ_1 increasing, ϕ_2 fixed). The last panel of Table 2 reports the time, t, at which the no-trade region becomes considerably wider relative to time 1. As expected, for any given cost combination, the time of widening is earlier for U than C portfolio choices. Also, for a given information set, the time of widening is earlier moving across a row or down a column.

The other dimension of interest is the average holding of the risky asset over the investor's life and its relation to the midpoint of the no-trade region. These parameters are reported in Table 2 and plotted in Figure 4. The behaviour of these parameters over the investor's life depends on whether the investor uses the U or C distribution. If the U distribution is being used, both the midpoint and the average holding are constant over the investor' life. While the U and C distributions have similar no-trade midpoints and average risky-asset holdings at time-240, the C distribution causes the investor to hold more of the risky asset early in life. For example, from Table 2, the investor facing P1F1 has an average holding of approximately 0.59 both "Early" and "Late" when using U. Use of C causes the investor's average "Early" holding to increase to 0.703 while leaving the investor's average "Late" holding largely unaffected at 0.608.

The same result can be seen in Figure 4 where the same patterns are documented for four cost combinations including P0F0. Average holding and midpoint with return predictability (C) start declining toward the U values almost from the start of the investor's life, and by the end ($t = 239$), the two are virtually identical. This finding for the no-transaction-cost case is consistent with evidence in Brennan, Schwartz and Lagnado (1996), Barberis (1997) and Campbell and Viceira (1996). Our paper shows that hedging demand is robust to the presence of transaction costs

Two final observations are in order. First, the decline in the α for C toward the α for U becomes more rapid as the investor becomes older. Second, the average holding is always higher than the midpoint.

A..2 Impact of Return Predictability on Rebalancing

This section assesses the impact of return predictability on portfolio choice by comparing two investors. The comparison is between an investor fully exploiting predictability that is present in return (C investor) and an investor facing an i.i.d. return whose distribution equals the steady-state distribution for the predictable return (U investor). Our focus is on the seven parameters described above.

Table 2 shows that rebalancing frequency both "Early" and "Late" is higher when the investor faces C rather than U. The difference is particularly pronounced when the fixed cost is non-zero. For example, rebalancing frequency for the P1F1 investor "Early" in life increases from 2.32% (or once every 43 months) with i.i.d. returns to 8.85% (or once every 11 months) when returns are predictable. If the investor only faces the proportional cost (P1F0), once every 4.7 months with U becomes once every 3.7 months with C. Figure 1 shows that the rebalancing frequency is higher throughout the investor's life when facing C rather than U for the three cost combinations plotted.

An even stronger result is the effect of return predictability on the actual cost incurred by the investor. Table 2 shows that the cost incurred per month is typically an order of magnitude higher when the investor faces C rather than U, irrespective of the cost combination or whether the investor is young or old. The P1F1 investor early in life spends 0.106 basis points of portfolio value per month when facing U, but spends 0.898 basis points when facing C. Similarly, over the last two years of life, this same investor spends 0.700 rather than 0.030 basis points if facing C rather than U. Thus, return predictability causes the investor to spend considerably more on transaction costs. The increased rebalancing and spending is not surprising since 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 2 shows that the no-trade width early in life for the P1F1 investor increases from 0.18 to 0.295 if returns are predictable. Near the end of the investor's life, the difference narrows, as can be seen in Figure 3. 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 sub-optimal α choice increases in the return volatility. Going from the U to C distribution, the conditional volatility of return decreases. Thus, this same reasoning helps explain why no-trade region widen when returns are predictable. The other reason for wider no-trade regions is cross-state variation in the position of the no-trade region. The current risky-asset holding may be very far from the optimal α today, but be very close to the optimal α next period.

No-trade regions can widen as rebalancing frequencies increase for the following reason. With predictable returns, the state changes through time and an inherited portfolio selection, $\hat{\alpha}$, that is inside the no-trade region for one state, can be outside for another. With i.i.d. returns, there is only one state of the world and so rebalancing does not occur due to changes in the state.

Table 2 also shows that return predictability causes the widening of no-trade regions to start later in life, for all cost combinations. As discussed above, this result is consistent with greater benefits from rebalancing when returns are predictable. Finally, return predictability also increases the return distance, especially when the fixed cost is high.

A..3 Varying Transaction Costs

Table 2 highlights how most of the reported parameters vary systematically with the proportional cost (ϕ_1) and the fixed cost (ϕ_2). Rebalancing frequency is decreasing in both ϕ_1 and in ϕ_2. This result is robust to the age of the investor and the presence of predictability. In particular, within each 3 by 3 subtable for rebalancing frequency in table 2, the frequencies decrease going across any row (ϕ_1 increasing) or down any column (ϕ_2 increasing). This result is consistent with the intuition that a larger cost per trade (ϕ_1 or ϕ_2 bigger) leads to fewer trades.

The presence of a small fixed cost has a particularly large impact. With i.i.d. returns, adding a 0.01% fixed cost to a proportional cost of 0.5% causes rebalancing frequency to drop from once every 4.7 months to once every 43 months (from 21.17% to 2.32%). With predictable returns, the rebalancing frequency drops from once every 3.7 months to once every 11 months (from 27.07% to 8.85%). This fixed cost (0.01%) is an order of magnitude smaller than the smaller proportional cost (0.5%) that we consider. This fixed cost can have such a large effect because it is a percentage of entire portfolio value while the proportional cost is a percentage of the dollar change in risky-asset holding. The entire portfolio value is typically much larger than the dollar change in risky-asset holding.

Turning to cost incurred, its relation to the cost parameters ϕ_1 and ϕ_2 is more complicated. While cost incurred increases monotonically with ϕ_1 early in life, the same is not true at the end of the investor's life. Early in life, the relation between cost incurred and ϕ_2 is monotonically increasing when ϕ_1 is zero or low, but becomes monotonically decreasing when ϕ_1 is high. Late in life, the relation between cost incurred and ϕ_1 is not even monotonic. The convoluted nature of these relations between total cost incurred and cost per trade is understandable. While an increase in the cost per trade may be expected to reduce trading frequency, its impact on total trading cost is ambiguous.

Excessive rebalancing is curbed by widening the no-trade region. Consequently, no-trade width is increasing in both ϕ_1 and in ϕ_2. This result is robust to the age of the investor and the predictability of returns. In particular, within each 3 by 3 subtable for rebalancing frequency in table 2, the widths increase going across any row (ϕ_1 increasing) or down any column (ϕ_2 increasing). On the other hand, return distance is increasing in the fixed cost parameter (ϕ_2), but is insensitive to changes in ϕ_1. Since the fixed cost must be non-zero for the return distance to be positive, this finding is not surprising.

Table 2 and Figure 3 show that, for a given distribution (U or C) and time of life, the no-trade midpoint is unaffected by variation in ϕ_1 or ϕ_2. 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_1 and in ϕ_2. For a given return distribution (U or C) and time of life ("Early" or "Late"), average holding less no-trade midpoint increases going across a row (ϕ_1 increasing) or down a column (ϕ_2 increasing). The likely reason is that rebalancing occurs less frequently when transaction costs are high. Since 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, as described and explained above, the time of widening of the no-trade region is monotonically decreasing in both ϕ_1 and in ϕ_2. So bigger costs cause the no-trade region to start widening earlier.

B. *Varying the Investor's Profile.*

Table 3 reports portfolio choice parameters, for a number of different consumer profiles summarized in Panel B of Table 1 and described in detail in section III. Time-series averages for "Early" and "Late" in life are reported for an investor facing either the unconditional (U) or conditional (C) return distribution. The C distribution is the S0 return-generating process which is calibrated to U.S. returns. The investors face cost combination P1F0 with $\phi_1 = 0.5\%$ and no fixed cost.

B..1 Varying Risk Aversion

The "ha" and "si" investors are identical except that the "ha" investor has a risk aversion coefficient of 7 rather than 4. Parameters for "si" and "ha" are reported in the first two columns of the "Early" and the "Late" sides of Table 3. The higher risk aversion causes rebalancing freqency and cost incurred to be higher for the "ha" investor. On the other hand, average risky-asset holding and the no-trade midpoint are much lower for the "ha" investor than the "si" investor. When facing the C distribution, the average risky-asset holding "Early" in life drops from 0.700 for "si" to 0.472 for "ha".

The higher risk aversion also causes the no-trade width to be smaller. "Early" in life facing the C distribution, the no-trade width is 0.176 for the "si" investor, but drops to 0.136 for the "ha" investor. Consistent with rebalancing frequency being higher for "ha" than "si", the no-trade region starts widening later for "ha" than "si".

Several of these results follow from the high-γ investor caring more about the dispersion of wealth. For example, a narrower no-trade region means a less dispersed portfolio payoff. Similarly, shifting down the no-trade region also reduces the volatility of the portfolio payoff.

B..2 Adding Death Probabilities prior to a Terminal Date of Age 101

The "pa" and "si" investors differ only to the extent that "pa" has a positive death probability prior to the terminal date while "si" does not. The death probability at each t for "pa" is calibrated to match mortality rates for U.S. males with $t = 1$ corresponding to 80 years and 1 month old and $t = 240$ corresponding to 101 years old. The "pa" investor is assumed to die with probability 1 at time $t = 240$. The third column of Table 3 and Figure 5 report parameters for "pa". The main result is that all parameter values are virtually identical to those obtained for the "si" investor. Positive death probabilities are analogous to lower rates of time preference, which cause the investor's time-τ problem (for any given $\tau < 239$) to be more like the time-239 problem. Consequently, we expect the $t = 1$ no-trade width to be larger, and the widening of the no-trade region to start earlier, for "pa" than "si". However, the death probabilities faced by U.S. males in their 80s and 90s are not large enough to materially alter rebalancing rules.

B..3 Allowing the Investor to Live Beyond Age 101

A further question is to examine rebalancing behaviour between the ages of 81 and 1 month and 101 but relax the assumption that death is a probability one event upon reaching 101. To examine this question, the "dl" investor differs from the "pd" investor in that the probability of death between 101 and 120 is also calibrated to U.S. males (see section III for details). Table 3 reports portfolio parameters for these two investors averaged over the two year period starting at age 88 ("Early") and at age 98 ("Late"). Portfolio parameters for the "Early" period are virtually identical for both "pd" and "dl". However, the "Late" parameters for "dl" are much more like the "Early" parameter values than like the "Late" values for "pd". In fact, the last two rows of Table 3 and Figure 5 indicate that the no-trade widening for "dl" remains less than 0.05 until $t = 454$ using U or until $t = 458$ using C. This result suggests that the no-trade region only widens close to the time when death is certain. U.S. investors as old as 100 have annual mortality rates through until 120 that are much lower than 1. Thus, rebalancing behavior of older investors in the U.S. may optimally not involve the rapid widening of the no-trade region that we documented above.

B..4 Eliminating Intermediate Consumption

Another interesting question is the effect of intermediate consumption on portfolio choices. The "nc" and "si" investors whose portfolio paramaters are reported in Table 3 and plotted in Figure 6, only differ to the extent that "nc" does not care about intermediate consumption. The top-right graph of Figure 6 shows that the absence of intermediate consumption causes no-trade regions to widen later in life, irrespective of whether the investor uses the U or C distribution. This result is confirmed by the last two rows of Table 3 which show that the time of no-trade widening is later for "nc" than "si" irrespective of whether the investor uses U or C. Further, when the investor uses the conditional ditribution, the average risky-asset holding and no-trade midpoint rise more rapidly for "nc" than "si" as t goes from 239 back to 1 (see the middle graphs of Figure 6).

The implication is that the presence of intermediate consumption causes the one-period problem to converge more slowly to the infinite-period problem as t declines from 239. This finding follows from intermediate consumption causing end-of-period wealth to be a smaller percentage of current-period wealth, especially as t approaches 239. This reduced importance of future wealth means that more periods are needed for the multi-period problem to converge to the infinite-period problem.

Finally, when the investor uses the conditional distribution, the average risky-asset holding and the no-trade midpoint both rise less in the presence of intermediate consumption.

The two middle graphs in Figure 6 show that these two parameters are lower at time $t = 1$ for "si" than for "nc". To understand why, first note that hedging demand arises because the investor cares about the future. As mentioned above, consumption lowers future wealth relative to current wealth, and hence makes the future less important.

B..5 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 proportional cost parameter. However, this result is likely to be sensitive to the assumption of costless consumption out of either asset. To explore the impact of relaxing this assumption, Table 3 and Figure 6 report choice parameters for the "cl" investor who is the same as the "nc" investor except the "cl" investor pays a proportional cost equal to $\phi_1 = 0.5\%$ to liquidate the risky asset at $t = 239$. Note that liquidation of the riskless asset is still costless so a wedge exists between the liquidation costs of the two assets, that may affect no-trade midpoints as t approaches 239.

Table 3 and especially Figure 6 show that rebalancing behaviour just prior to the terminal date is very different for the "cl" investor than for the "nc" investor. The clearest illustration that differential liquidation costs affect rebalancing rules is contained in the bottom two graphs in Figure 6. These two graphs show the boundaries of the no-trade region for the "nc" and "cl" investors facing the P1F0 cost structure. The left graph shows the boundaries for investories facing the U distribution of S0, while the right graph shows the average boundaries for investors facing the C distribution. Also plotted in each graph is the risky asset holding chosen by the investor facing zero transaction and liquidation costs and facing the relevant distribution. For the "nc" investor, the upper no-trade boundaries moves up and the bottom boundary moves down as the region widens late in life. On the other hand, the upper bound for the "cl" investor drops late in life while the bottom boundary drops lower than that of the "nc" investor. In fact, the upper boundary converges to the no-cost rebalancing point as t approaches 239. The drops in the two boundaries near the terminal date make sense since any holding of the risky asset on that date attracts a liquidation cost.

Thus, reducing the risky asset holding just prior to the terminal date is more attractive and increasing it is less attractive.

Continuing the comparison of "cl" to "nc", the increase in rebalancing due to the dropping upper boundary more than offsets the reduced rebalancing due to the lower bottom boundary. Table 3 reports that both rebalancing frequency and cost incurred are higher "Late" in life for "cl" than "nc". This result holds for both the U and C distribution as the top left graph in Figure 6 confirms.

C. *Varying the Return-generating Process.*

Results for a variety of return-generating processes are contained in Table 4 and Figures 7 and 8. The cost combination is P1F0 and the "si" investor makes the decisions. The return-generating processes are organized, as in Panel C of Table 1, to allow the impact of each return parameter to be assessed in turn. For each return parameter, up to two sets of comparisons are available. Within a comparison, the return parameters are chosen to ensure that the parameters of the unconditional distribution are the same. Table 4 reports choice parameters "Early" in life for an investor facing the conditional distribution together with choice parameters when facing the resulting unconditional distribution.[22] Figures 7 and 8 report $t = 1$ no-trade widths plotted as a function of the dividend state for conditional portfolio choices.[23] The investor's no-trade width when facing the associated unconditional distribution is also graphed as a flat line and labelled U.

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. The first panel of Table 4 contains two comparisons, S0-S1 and S2-S3, for which all return parameters

[22] As expected, unconditional portfolio choice parameters within a comparison are almost always the same across return-generating processes. The only exception is rebalancing frequency which is sometimes higher as a result of the coarser grid for return when the correlation between return and dividend yield is zero. Thus, for a given comparision, unconditional parameters are obtained by averaging across the U parameters for the return-generating processes within the comparison.

[23] If one of the no-trade boundaries in a state equals 0 or 1, then width in that state is not plotted.

except $b_r\sigma_d$ are held fixed. The top left and top right graphs in Figure 7 plot the no-trade width for the S0-S1 and S2-S3 comparison respectively.

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 vs. S1) or have values of zero (as in S2 vs. S3). Table 4 reports an increase in no-trade width from 0.121 to 0.176 going from S1 ($b_r\sigma_d$ =0.13%) to S0 ($b_r\sigma_d$ =0.26%) and from 0.261 to 0.398 going from S3 ($b_r\sigma_d$ =0.26%) to S2 ($b_r\sigma_d = 0.53\%$).

Figure 3 shows that while the other two return parameters affect the shape of the curve that plots no-trade width as a function of dividend state, decreasing $b_r\sigma_d$ shifts the curve down in both comparisons. In particular, the no-trade width for S1 with $b_r\sigma_d = 0.13\%$ drops below the unconditional width in some states.

As $b_r\sigma_d$ increases, conditional single-period volatility decreases because unconditional volatility remains constant within comparisons. As discussed above, lower return volatility leads to wider no-trade regions since the cost of a sub-optimal α are lower. This provides one reason for why no-trade regions are wider when $b_r\sigma_d$ increases. Another reason stems from the positive relation between $b_r\sigma_d$ and the volatility of conditional expected return. This positive relation means that the absolute difference between expected return next period and this period is likely to be bigger. Thus, if $b_r\sigma_d$ increases and no-trade width is held fixed, it is more likely that this period's rebalancing will need to be reversed next period due to a shift in expected return. Instead, no-trade width increases when $b_r\sigma_d$ increases to reduce such rebalancing.

The S0-S1 comparison in the first panel of Table 4 also shows that an increase in $b_r\sigma_d$ when ρ_{ev} is non-zero leads to a higher average risky asset holding and a higher no-trade midpoint. In other words, increasing the magnitude of the single-period predictability magnifies any hedging demand. Finally, going from S2 to S3, $b_r\sigma_d$ decreases and so does the average cost incurred (from 2.495 to 0.156 basis points). The same is not true going from S0 to S1 since $b_r\sigma_d$ decreases but average cost incurred increases slightly (from 0.854 to 0.897 basis points). This last result is a little puzzling but is likely due to the non-zero values for b_d and ρ_{ev} in

S0 and S1 comparison versus the zero values in S2 and S3.

C..2 Varying the Persistence of the Predictive Variable

The parameter b_d controls the persistence of the predictive variable, dividend yield, and we are interested in how this persistence affects rebalancing behaviour. The second panel of Table 4 and the middle graph of Figure 7 report results for the S4-S5-S3 comparison which fixes $b_r\sigma_d$ and ρ_{ev} and varies b_d.

With respect to the no-trade width, the across-state averages reported in Table 4 do not reveal 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 exteme value for dividend yield, and so the no-trade region widens for extreme d states. This disinclination to rebalance is understandable since 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 since expected return reverts more quickly to its unconditional value.

Since the S4-S5-S3 comparison sets ρ_{ev} equal to zero, the investor does not have a hedging demand and so the no-trade midpoint and average risky-asset holding are similar across the three return-generating processes and the unconditional i.i.d. process. However, the persistence of d has a large positive effect on the average cost that the investor incurs. Going from S3 to S5 to S4 in Table 4, b_d increases monotonically from 0 to 0.96 while the average cost incurred also increases monotonically from 0.156 to 0.518 basis points. This result is not surprising and for the following reason. When conditional expected return is persistent, the benefit from rebalancing is more likely to exceed the associated transaction cost and so the investor optimally spends more on transaction costs.

C..3 Varying the Contemporaneous Correlation between Return and the Predictive Variable

For the investor to be able to hedge against future predictability, ρ_{ev} must be non-zero. Thus, ρ_{ev} is referred to as the hedgeability parameter. The third panel of Table 4 and the bottom two graphs in Figure 7 contains choice parameters for two comparisons, S0-S4-S6 and S7-S8-S9. Within each comparison, the only return parameter that is varying is ρ_{ev}.

The average no-trade width does not sytematically vary with ρ_{ev} since the nature of the relation is different across the two comparisons. This result is surprising since the sign of ρ_{ev} determines the sign of the return autocorrelation. Turning to the graphs in Figure 7, we see that the sign of ρ_{ev} impacts the relation between no-trade width and dividend state. The shape of the relation differs across the two graphs due to the different values for b_d in each comparison. However, in both graphs, a large positive ρ_{ev} produces no-trade regions that are wider for high d states and lower for low d states than those produced by a large negative ρ_{ev}.

With respect to these results, a couple of comments are in order. Intuition that negative autocorrelation in $(\hat{\alpha}_t - \alpha_{t-1})$ leads to wider no-trade regions does not apply here. That intuition is built around the autocorrelation in $(\hat{\alpha}_t - \alpha_{t-1})$ determining the direction and magnitude of the relation between $E[\hat{\alpha}_{t+1} - \alpha_t|\hat{\alpha}_t$ and no time-t rebalance] and $\hat{\alpha}_t$. So if the autocorrelation is negative and $\hat{\alpha}_t$ is extreme, then $\hat{\alpha}_{t+1}$ is expected to be less extreme, and the need to rebalance is less urgent. However, here d_t is known to the investor facing the conditional distribution, and d_t determines the expected $t+1$ return at t through b_r. So if d_t is known at t, ρ_{ev} has no impact on $E[r_{t+1}|t]$. Since r_{t+1} determines $\hat{\alpha}_{t+1} - \alpha_t$, it follows that $E[\hat{\alpha}_{t+1} - \alpha_t|d_t]$ is unaffected by α_t or ρ_{ev} and so the earlier intuition fails.

Instead, the mechanism by which ρ_{ev} affects no-trade width seems much more complex. In particular, ρ_{ev} determines the covariance of r_{t+1} with d_{t+1}, conditioning on d_t. Expected r_{t+2} at $t+1$ depends on d_{t+1} and, for a given α_t, r_{t+1} determines $\hat{\alpha}_{t+1}$. Thus, ρ_{ev} is going to determine the extent to which the $t+1$-shock to $\hat{\alpha}_{t+1}$ moves with the $t+1$-shock to expected r_{t+2}. Now the investor is more willing to rebalance at t if, in the absence of such rebalancing, a similar rebalancing is likely to be attractive at $t+1$. However, knowing the correlation of the $t+1$ shocks is not enough to determine whether such a rebalancing is more attractive at $t+1$. Rather, the attractiveness of such a rebalancing at $t+1$ also depends on the expected levels of $\hat{\alpha}_{t+1}$ and $E[r_{t+2}|t+1]$ given d_t.

Further work is needed to better understand exactly how ρ_{ev} affects no-trade region. However, our work suggests that the wider no-trade region for the conditional S0 investor, relative to the U investor, is not being driven by the negative ρ_{ev} for the S0 process. Instead, the magnitude of the single-period predictability $(b_r \sigma_d)$ is the return parameter having the biggest effect on the average no-trade region across states. In fact, S1, which has a low b_r value and a negative ρ_{ev} equal to the data ρ_{ev}, is the only process with no-trade widths that are below the unconditional no-trade width in some states.

A final note concerns the relation between the hedgeability parameter and two choice parameters: average risky-asset holding and no-trade midpoint. Both these parameters are monotonically decreasing in ρ_{ev} for both the S0-S4-S6 and S7-S8-S9 comparisons. Thus, hedging demand is positive when ρ_{ev} is negative, as it is in the data, and negative when ρ_{ev} is positive. Earlier studies without transaction costs also found a positive hedging demand when return is negatively correlated with expected future return and the CRRA investor has a risk aversion coefficient greater than 1.[24] Thus, the presence of transaction costs appears to have little effect on the direction of hedging demand.

C..4 Allowing the Risky-asset Return to Exhibit Heteroscedasticity

The final panel of Table 4 and Figure 8 report results for return-generating processes that allow returns to exhibit heteroscedasticity. As ζ_1 increases going from H2 to H1 to H0, the positive slope for conditional volatility as a linear function of conditional mean is becoming larger. The results in Table 4 which are averaged across states reveal no systematic patterns going from H2 to H1 to H0. This is not surprising since the unconditional volatility of the

[24]See Barberis (1996), Campbell and Viceira (1996), Brennan, Lagnado and Schwartz (1996) and Kim and Omberg (1996)).

return residual (e) is the same across the three processes. In contrast, the plots of no-trade width and no-trade midpoint as functions of the state in Figure 8 vary systematically going from S0 (no heteroscedasticity) to H2 to H1 to H0.[25]

As discussed above, the cost of a sub-optimal 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 since ζ_1 is increasing going across those states. The left graph of Figure 8 confirms this result.

Since ours is the first paper to consider the effect of return heteroscedasticity on portfolio choice in a multiperiod setting, we also report no-trade midpoint by state in the right graph of Figure 8. 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, the right graph of Figure 8 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 heteroscedasticity can potentially have a big impact on rebalancing behavior.

V. Conclusions

Transaction costs and return predictability are realistic features of the environment facing a U.S. investor or mutual fund. How these two 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 frequency change over the lifecycle and are affected by return predictability. Two types of costs are evaluated: *proportional* to the change in the holding of the risky asset and a *fixed* fraction of portfolio value.

[25] Recall that conditional mean return in each state is held constant across return generating processes by construction.

We find that realistic transaction costs can materially affect rebalancing behavior, creating no-trade regions that widen near the investor's terminal date. At the same time, realistic proportional and fixed costs have little effect on the midpoint of the no-trade region, unless liquidation costs differ across assets. Return predictability calibrated to U.S. stock returns is found to have large effects on rebalancing behaviour relative to independent and identically distributed (i.i.d.) returns with the same unconditional distribution. For example, 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. No-trade regions are wider early in life when returns are predictable. Finally, we find that the nature of the return predictability, including the presence of return heteroscedasticity, can have large effects on rebalancing behavior.

The paper's results have a number of implications. First, assessments of the impact of transaction costs on the joint distribution of consumption and asset returns are likely to be sensitive to whether returns are allowed to be predictable. Thus, the ability of transaction costs to explain all or part of the equity premium puzzle may be sensitive to the assumed level of return predictability. More work is needed to assess the importance of return predictability for explanations of the observed equity premium that rely on transaction costs.

Second, our results on rebalancing frequencies have implications for existing studies (He and Modest (1995) and Luttmer (1996a)) documenting how proportional transaction costs weaken Hansen-Jagannathan bounds on intertemporal marginal rates of substitution (IMRSs). When investors face a proportional cost of 50 basis points to trade equity, our results for realistic predictability suggest that these weaker bounds can be informative for data frequencies as high as a third of a year. Third, the short horizon of mutual fund managers has the advantage that an overconfident manager churns less. Finally, Samuelson's (1969) result that optimal portfolio choices remain constant over an investor's lifetime when the opportunity set is constant, breaks down in the presence of transaction costs.

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 porfolio given that information, in a setting with transaction costs.

References

Balduzzi, Pierluigi, and Anthony Lynch, 1998, Transaction costs and predictability: Some utility cost calculations, forthcoming Journal of Financial Economics.

Barberis, Nicholas, 1997, Investing for the long run when returns are predictable, Working Paper, University of Chicago.

Bhardwaj, Ravinder K., and Leroy D. Brooks, 1992, The january anomaly: Effects of low share price, transaction costs and bid-ask bias, *Journal of Finance* 47, 553-574.

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.

Brennan, Michael J., and Eduardo S. Schwartz, 1996, The use of treasury bill futures in strategic asset allocation Programs, Working paper (UCLA, Los Angeles, CA).

Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1996, Strategic asset allocation, forthcoming in the Journal of Economic Dynamics and Control.

Brown, K., W. Harlow and L. Starks, 1996, Of Tournaments and Temptations: An Analysis of Managerial Incentives in the Mutual Fund Industry, *Journal of Finance* 51, 85-110.

Campbell, John Y., 1987, Stock returns and the term structure, Journal of Financial Economics 18, 373-99.

Campbell, J. Y., and L.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, 842-62.

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, G. and A. 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, John Heaton, and Erzo Luttmer, 1995, Econometric evaluation of asset pricing models, *Review of Financial Studies 8*, 237-274.

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.

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

Samuelson, Paul A., 1969, Lifetime Portfolio Selection by Dynamic Stochastic Programming, *The Review of Economics and Statistics* 51, 239-246.

Schroder, Mark, 1995, Optimal portfolio selection with fixed transaction costs, working paper, Northwestern University.

Society of Actuaries Group Annuity Valuation Table Task Force, 1995, 1994 Group Annuity Mortality Table and 1994 Group Annuity Reserving Table, *Society of Actuaries Transactions* 47, 865-913.

Stoll, Hans R., and Robert E. Whaley, 1983, Transaction costs and the small firm effect, *Journal of Financial Economics* 12, 57-79.

Tauchen, George,. and Robert Hussey, 1991, Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models, *Econometrica* 59, 317-96.

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.

Table 1. Return-generating Processes, Transaction Costs, and Investor Profiles

The investor problem described in Section I is solved numerically for a variety of return-generating processes, transaction costs, and investor profiles which are described below. Panel A lists the combinations of proportional (φ_1) and fixed (φ_2) cost parameters that the investor faces. The parameters φ_1 and φ_2 are defined in equation (7). Panel B characterizes the various investor profiles that we consider. For the "pd" and "dl" profiles, t=1 is taken to be age 81 and 1 month when using CAM-94 to obtain death probabilities: details are in section II. Panel C lists statistics for the various return-generating processes considered and for U.S. equity return. The return-generating processes are organized to facilitate comparative-static analysis for each return parameter: the magnitude of the single-period predictability ($b_r\sigma_d$), the persistence of the predictability (b_d), the hedgeability parameter ρ_{ev}, and the heteroscedasticity parameter (ζ_1). Return-generating processes are obtained using the quadrature procedure described in Section II.

Panel A: Transaction Cost Setting

Transaction Cost Setting	P0F0	P0F1	P0F2	P1F0	P1F1	P1F2	P2F0	P2F1	P2F2
Proportional(φ_1)	0%	0%	0%	0.5%	0.5%	0.5%	2.5%	2.5%	2.5%
Fixed(φ_2)	0%	0.01%	0.1%	0%	0.01%	0.1%	0%	0.01%	0.1%

Panel B: Investor Profiles

Profile	Description of Investor	ζ	Prob of Early Death	Time of Death	Intermediate Consump.	Cost to Consume from Risky (φ_T)
si	Standard	4	0	240	Yes	0
ha	High Risk Aversion	7	0	240	Yes	0
pd	Positive Death Prob.	4	CAM-94	240	Yes	0
dl	Die Later	4	CAM-94	468	Yes	0
nc	No Intermediate Consumption	4	0	240	No	0
cl	Costly Risky Liquidation	4	0	240	No	φ_1

Panel C: Return-generating Processes

Parameter Being Varied	Process	σ_r x 100	$b_r\,\sigma_d$ x100	b_d	ρ_{ev}	ζ_1x 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
ρ_{ev}	S0	5.443	0.260	0.962	-0.923	0
	S4	5.443	0.260	0.962	0	0
	S6	5.507	0.260	0.962	0.923	0
	S7	4.819	0.132	0	-0.900	0
	S8	4.819	0.132	0	0	0
	S9	4.818	0.132	0	0.900	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

Table 2. Portfolio Choice Parameters as Transaction Costs Vary.

U.S. equity return is discretized using the quadrature procedure described in Section II to give the return-generating process S0 in Panel C of Table 1. The riskfree rate is assumed constant. The investor profile is "si" which is characterized in Table 1. Nine transaction cost settings are considered, corresponding to the nine settings in Panel A of Table 1. 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. Both the solution technique used to solve the investor's problem and details of the simulation are contained in Section III. Parameters for both unconditional (U) and conditional (C) portfolio choices are reported. The portfolio choice parameters are averaged over two periods in an investor's life: "Early" refers to the 2-year period from t=96 to t=119; and "Late" refers the 2-year period from 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.

		Cost		Proportional					
		Fixed		P 0, φ_1 0%	P 1, φ_1 0.5%	P 2, φ_1 2.5%	P 0, φ_1 0%	P 1, φ_1 0.5%	P 2, φ_1 2.5%
		F	φ_2	Early			Late		
Rebalancing Frequency (%)	U	0	0%	100.00	21.17	14.70	100.00	8.30	0.13
		1	0.01%	3.06	2.32	1.71	2.14	0.69	0.00
		2	0.1%	1.32	1.11	0.72	0.25	0.09	0.00
	C	0	0%	100.00	27.07	16.03	100.00	21.12	5.26
		1	0.01%	15.03	8.85	4.65	13.72	6.36	1.71
		2	0.1%	5.29	3.96	2.31	4.62	2.86	0.93
Average Cost Incurred per Month (basis points of portfolio value)	U	0	0%	0.000	0.096	0.309	0.000	0.037	0.002
		1	0.01%	0.031	0.106	0.300	0.021	0.030	0.000
		2	0.1%	0.132	0.182	0.280	0.025	0.016	0.000
	C	0	0%	0.000	0.854	1.612	0.000	0.702	1.063
		1	0.01%	0.150	0.898	1.576	0.137	0.700	1.036
		2	0.1%	0.529	1.098	1.524	0.462	0.830	0.995
Average Holding	U	0	0%	0.562	0.576	0.604	0.562	0.584	0.646
		1	0.01%	0.574	0.586	0.611	0.571	0.593	0.654
		2	0.1%	0.580	0.593	0.616	0.596	0.617	0.672
	C	0	0%	0.673	0.700	0.742	0.693	0.602	0.660
		1	0.01%	0.685	0.703	0.745	0.587	0.608	0.666
		2	0.1%	0.701	0.712	0.755	0.609	0.627	0.680
Midpoint of No-trade Region (No-trade Midpoint)	U	0	0%	0.562	0.563	0.562	0.562	0.562	0.543
		1	0.01%	0.570	0.570	0.562	0.560	0.559	0.540
		2	0.1%	0.560	0.560	0.559	0.562	0.555	0.532
	C	0	0%	0.673	0.690	0.699	0.693	0.580	0.558
		1	0.01%	0.681	0.684	0.691	0.581	0.575	0.552
		2	0.1%	0.676	0.677	0.678	0.576	0.568	0.540

Table 2. cont.

		Fixed F	φ_2	Early: P 0, φ_1 0%	Early: P 1, φ_1 0.5%	Early: P 2, φ_1 2.5%	Late: P 0, φ_1 0%	Late: P 1, φ_1 0.5%	Late: P 2, φ_1 2.5%
Width of No-trade Region (No-trade Width)	U	0	0%	0.000	0.092	0.188	0.000	0.217	0.690
		1	0.01%	0.140	0.180	0.276	0.172	0.348	0.762
		2	0.1%	0.240	0.280	0.352	0.407	0.562	0.875
	C	0	0%	0.000	0.176	0.349	0.000	0.264	0.726
		1	0.01%	0.197	0.295	0.452	0.204	0.403	0.802
		2	0.1%	0.376	0.409	0.542	0.455	0.603	0.902
Distance between Boundaries and Return Points (Return Distance)	U	0	0%	0.000	0.000	0.000			
		1	0.01%	0.070	0.059	0.058			
		2	0.1%	0.120	0.124	0.114			
	C	0	0%	0.000	0.000	0.000			
		1	0.01%	0.098	0.079	0.070			
		2	0.1%	0.188	0.172	0.165			
Time of Widening of No-trade Region	U	0	0%	239	227	190			
		1	0.01%	235	221	182			
		2	0.1%	217	212	173			
	C	0	0%	239	231	192			
		1	0.01%	238	230	190			
		2	0.1%	229	220	182			

Table 3. Portfolio Choice Parameters as the Investor's Profile Varies.

U.S. equity return is discretized using the quadrature procedure described in Section II to give the return-generating process S0 in Panel C of Table 1. The riskfree rate is assumed constant. The transaction cost setting is P1F0 from Table 1 with the proportional cost parameter (φ_1) set equal to 0.005 and the fixed cost parameter (φ_2) set equal to 0. The various investor profiles that we consider are described in Section II.B and summarized in Panel B of Table I. For the "pd" and "dl" profiles, t=1 is taken tobe age 81 and 1 month when using CAM-94 to obtain death probabilities: details are in Section II.B. 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. Both the solution technique used to solve the investor's problem and details of the simulation are contained in Section III. Parameters for both unconditional (U) and conditional (C) portfolio choices are reported. The portfolio choice parameters are averaged over two periods in an investor's life: "Early" refers to the 2-year period from t=96 to t=119; and "Late" refers the 2-year period from 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.

		Profile											
		si	ha	pd	dl	nc	cl	si	ha	pd	dl	nc	cl
		Early						Late					
Rebalancing Frequency (%)	U	21.17	23.08	21.09	21.16	21.43	21.38	8.30	12.36	8.14	20.90	14.05	26.14
	C	27.07	34.67	27.19	27.09	27.40	27.43	21.12	26.90	20.91	27.78	27.15	33.86
Average Cost per Month (bp)	U	0.096	0.099	0.096	0.096	0.098	0.097	0.037	0.053	0.036	0.095	0.063	0.132
	C	0.854	0.950	0.855	0.852	0.849	0.851	0.702	0.838	0.702	0.849	0.831	0.941
Average Holding	U	0.576	0.329	0.576	0.576	0.576	0.575	0.584	0.334	0.585	0.576	0.578	0.561
	C	0.700	0.472	0.690	0.701	0.738	0.738	0.602	0.370	0.601	0.688	0.622	0.603
Midpoint of No-trade Region	U	0.563	0.321	0.562	0.563	0.563	0.563	0.562	0.321	0.562	0.562	0.562	0.510
	C	0.690	0.468	0.681	0.692	0.729	0.729	0.580	0.356	0.579	0.679	0.600	0.559
Width of No-trade Region	U	0.092	0.070	0.092	0.092	0.090	0.090	0.217	0.133	0.219	0.093	0.159	0.172
	C	0.176	0.136	0.176	0.176	0.172	0.172	0.264	0.173	0.265	0.176	0.220	0.235
		Time of Widening of No-trade Region											
	U	227	232	226	454	233	232						
	C	231	235	230	458	235	234						

Table 4. Portfolio Choice Parameters as the Return-generating Process Varies.
U.S. equity return is discretized using the quadrature procedure described in Section II and the riskfree rate is assumed constant. The investor profile is "si" and the transaction cost setting is P1F0, and both are characterized in Table 1. Twelve return-generating processes for the risky asset are considered, corresponding to the twelve processes in Panel C of Table 1. The return-generating processes are organized to facilitate comparative-static analysis for each return parameter: the magnitude of the single-period predictability ($b_r\ \sigma_d$), the persistence of the predictability (b_d), the hedgeability parameter ρ_{ev}, and the heteroscedasticity parameter (ζ_1). The return-generating processes within each comparison have the same unconditional means and volatilities. Average cost incurred per month and average holding of the risky asset are obtained by simulation of the investor's portfolio choices. Both the solution technique used to solve the investor's problem and details of the simulation are contained in Section III. Parameters for conditional and unconditional (U) portfolio choices are reported. The portfolio choice parameters are averaged over the 2-year period from t=96 to t=119 which is denoted "Early" in the other tables. Average cost incurred per month is expressed in basis points of portfolio value.

Return Parameter being Varied		Process			
$b_r\sigma_d$	b_d =0.96; ρ_{ev}=-0.92; ζ_1=0	S0 ($b_r\sigma_d$ =0.26%)		S1 ($b_r\sigma_d$=0.13%)	U (σ_r=5.5%)
	Av. Cost/mth (bp)	0.854		0.897	0.096
	Av. Risky Holding	0.700		0.638	0.576
	No-trade Midpoint	0.690		0.625	0.562
	No-trade Width	0.176		0.121	0.092
	b_d =0; ρ_{ev}=0; ζ_1=0	S2 ($b_r\sigma_d$ =0.53%)		S3 ($b_r\sigma_d$=0.26%)	U (σ_r=5.5%)
	Av. Cost/mth (bp)	2.495		0.156	0.099
	Av. Risky Holding	0.555		0.560	0.569
	No-trade Midpoint	0.546		0.553	0.557
	No-trade Width	0.398		0.261	0.087
b_d	$b_r\sigma_d$ =0.26%; ρ_{ev}=0; ζ_1=0	S4 (b_d=0.96)	S5 (b_d=0.32)	S3 (b_d=0)	U (σ_r=5.5%)
	Av. Cost/mth (bp)	0.518	0.371	0.156	0.100
	Av. Risky Holding	0.553	0.555	0.560	0.569
	No-trade Midpoint	0.549	0.552	0.553	0.557
	No-trade Width	0.180	0.270	0.261	0.086
ρ_{ev}	$b_r\sigma_d$ =0.26%; b_d=0.96; ζ_1=0	S0 (ρ_{ev}=-0.92)	S4 (ρ_{ev} =0)	S6 (ρ_{ev}=0.92)	U (σ_r=5.5%)
	Av. Cost/mth (bp)	0.854	0.518	1.749	0.097
	Av. Risky Holding	0.700	0.553	0.478	0.574
	No-trade Midpoint	0.690	0.549	0.472	0.561
	No-trade Width	0.176	0.180	0.197	0.090
	$b_r\sigma_d$ =0.13%; b_d=0; ζ_1=0	S7 (ρ_{ev}=-0.90)	S8 (ρ_{ev} =0)	S9 (ρ_{ev}=0.90)	U (σ_r=4.8%)
	Av. Cost/mth (bp)	0.068	0.077	0.075	0.066
	Av. Risky Holding	0.724	0.695	0.674	0.699
	No-trade Midpoint	0.704	0.683	0.664	0.682
	No-trade Width	0.202	0.187	0.176	0.090
ζ_1	$b_r\sigma_d$ =0.26%; b_d=0.96; ρ_{ev}=-0.92	H0 (ζ_1=1.28%)	H1 (ζ_1=0.57%)	H2 (ζ_1=0.11%)	
	Av. Cost/mth (bp)	0.836	0.808	0.837	
	Av. Risky Holding	0.684	0.682	0.695	
	No-trade Midpoint	0.673	0.668	0.685	
	No-trade Width	0.141	0.141	0.171	

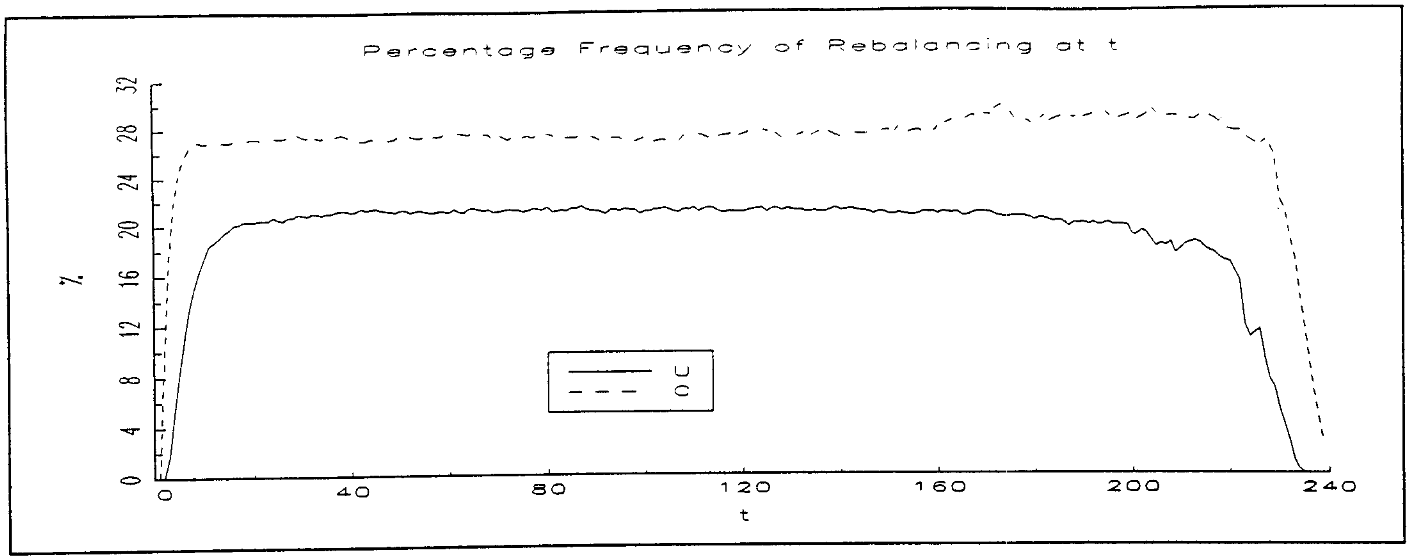

Setting P1F0 (Proportional Cost Only): φ_1=0.005, φ_2=0.

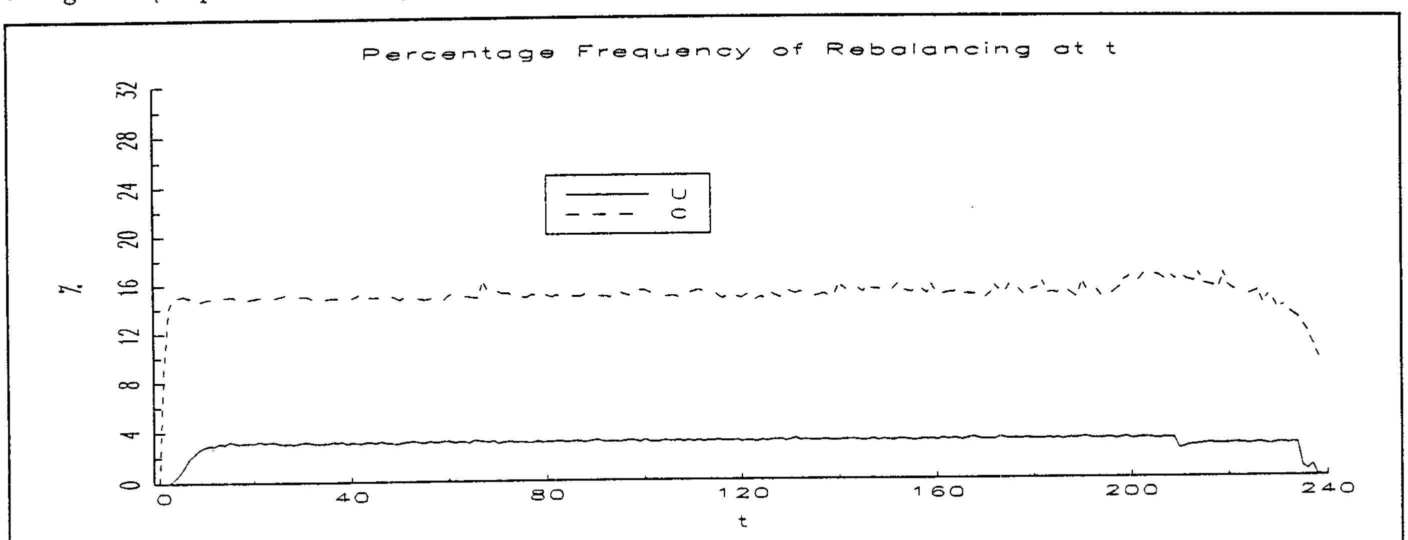

Setting P0F1 (Fixed Cost Only): φ_1=0, φ_2=0.0001.

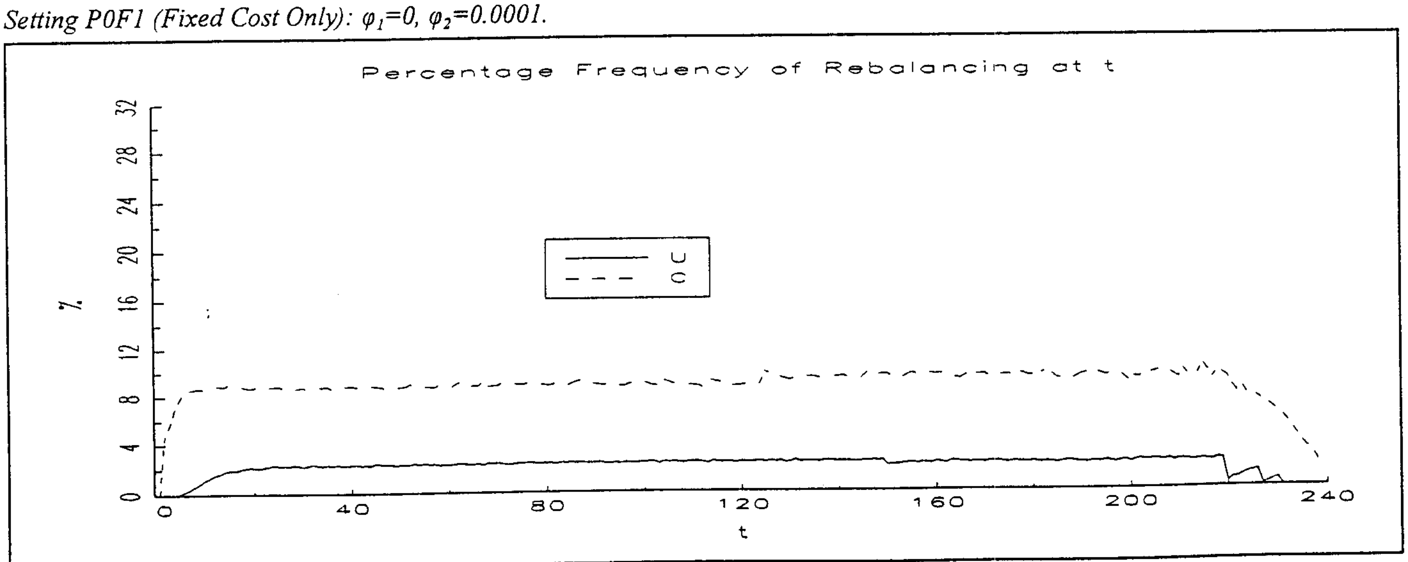

Setting P1F1 (Proportional and Fixed Costs): φ_1=0.005, φ_2=0.0001.

Figure 1. Average Rebalancing Frequency: Unconditional vs Conditional Portfolio Choices.

Average rebalancing frequency is plotted for the standard investor, si, (CRRA utility, $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption) facing return-generating process S0 (quadrature approximation to U.S. data) and making unconditional (U) and conditional (C) portfolio choices. The investor problem is described in Section I while section II describes the quadrature approximation.

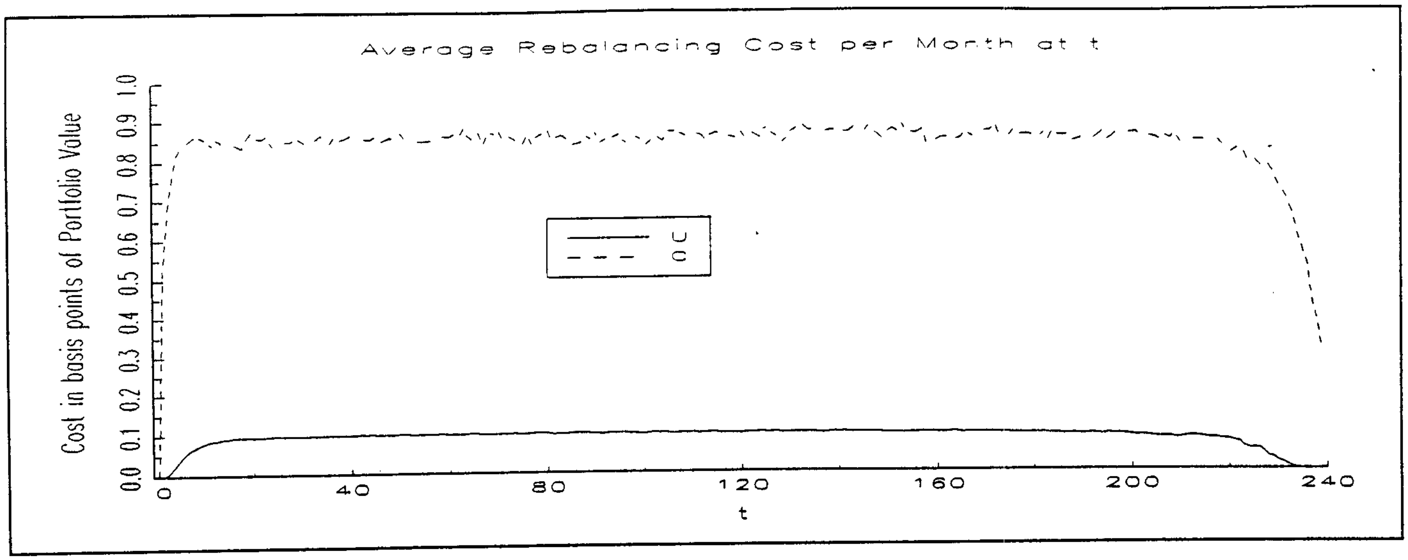

Setting P1F0 (Proportional Cost Only): $\varphi_1=0.005$, $\varphi_2=0$.

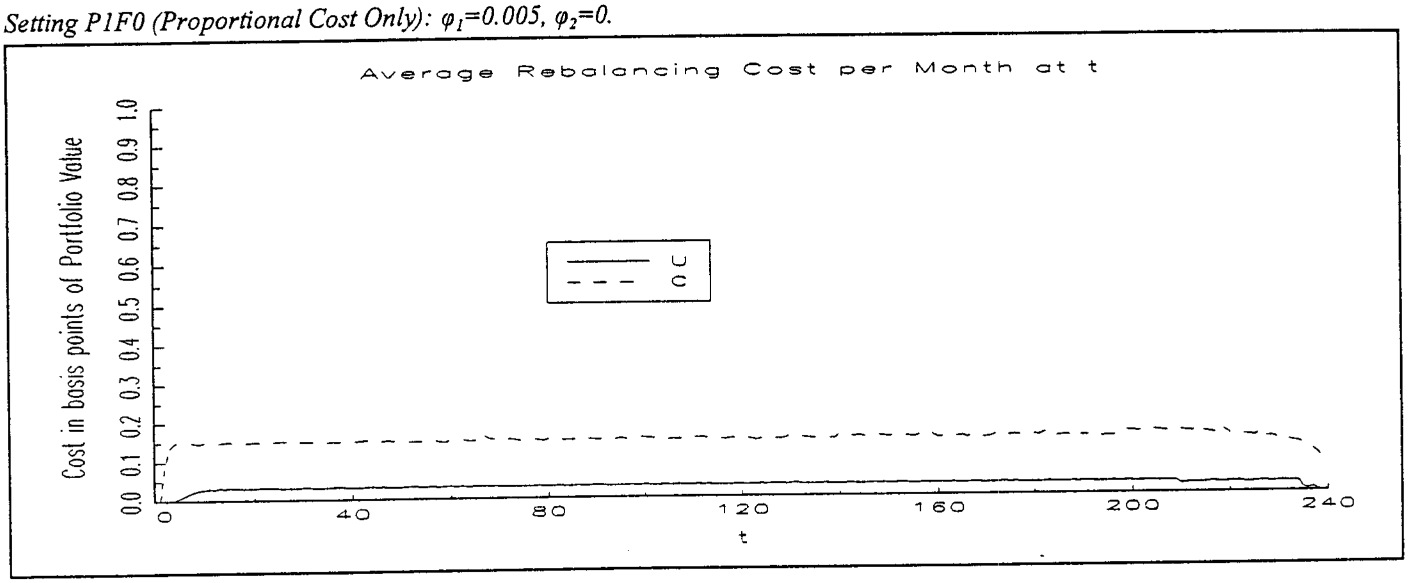

Setting P0F1 (Fixed Cost Only): $\varphi_1=0$, $\varphi_2=0.0001$.

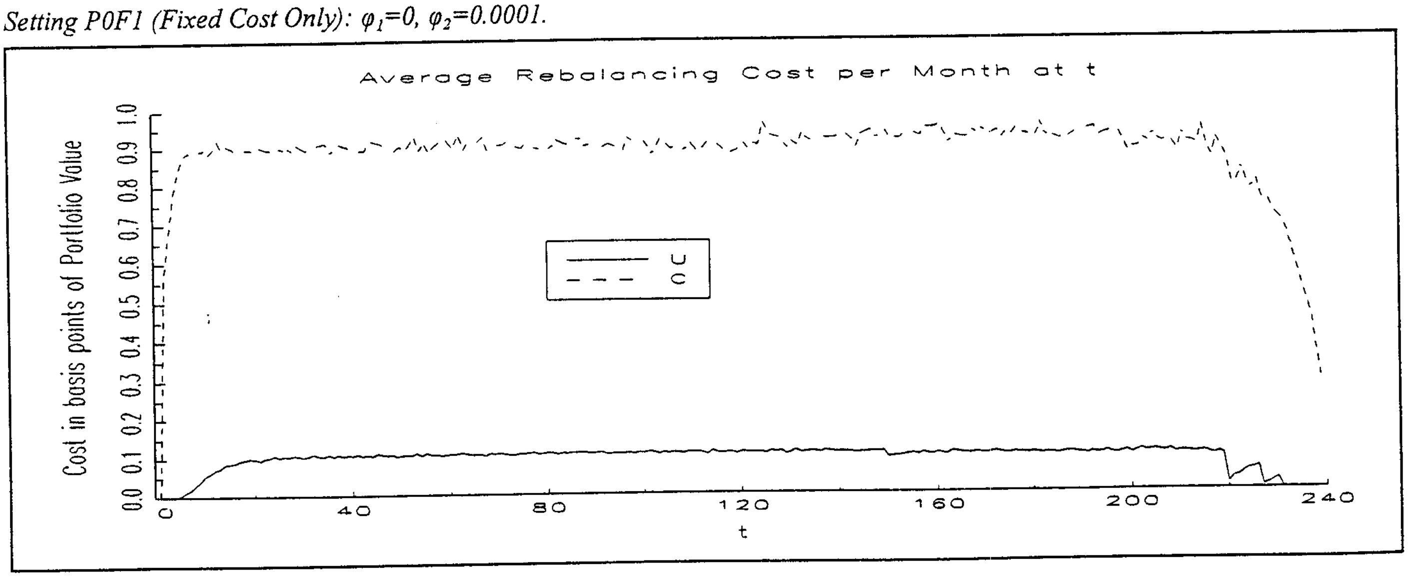

Setting P1F1 (Proportional and Fixed Costs): $\varphi_1=0.005$, $\varphi_2=0.0001$.

Figure 2. Average Monthly Cost per \$ of Portfolio Value: Unconditional vs Conditional Choices.

The average incurred cost per month is plotted for the standard investor, si, (CRRA utility, $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption) facing return-generating process S0 (quadrature approximation to U.S. data) and making unconditional (U) and conditional (C) portfolio choices. The investor problem is described in Section I while section II describes the quadrature approximation.

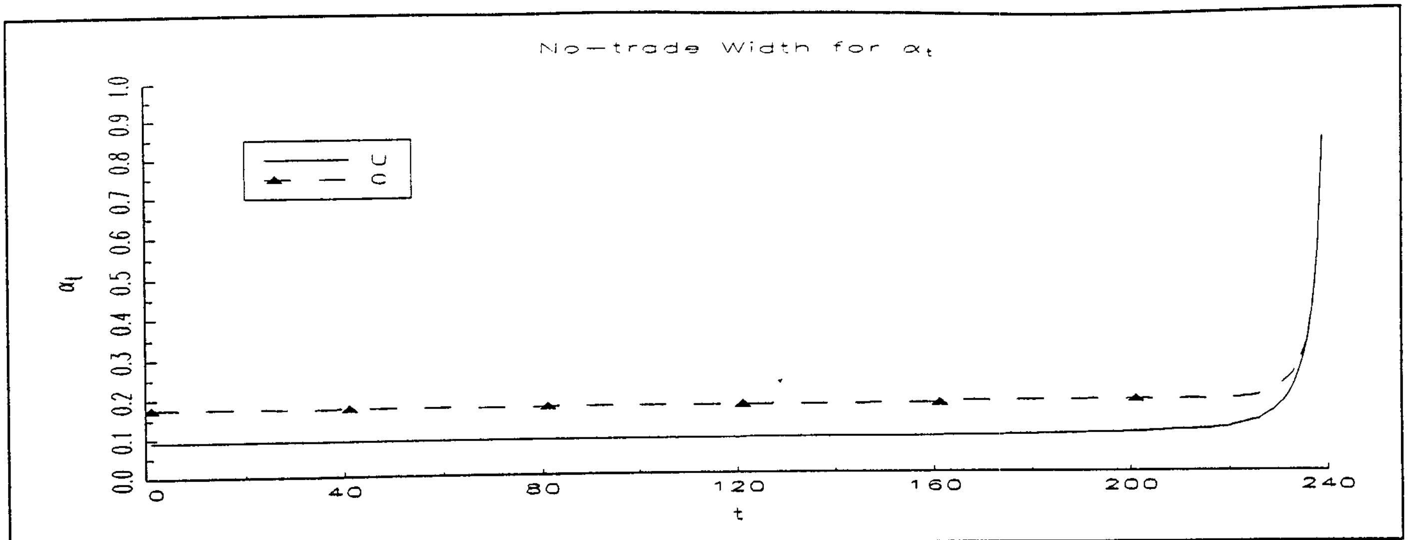

Setting P1F0 (Proportional Cost Only): φ_1=0.005, φ_2=0.

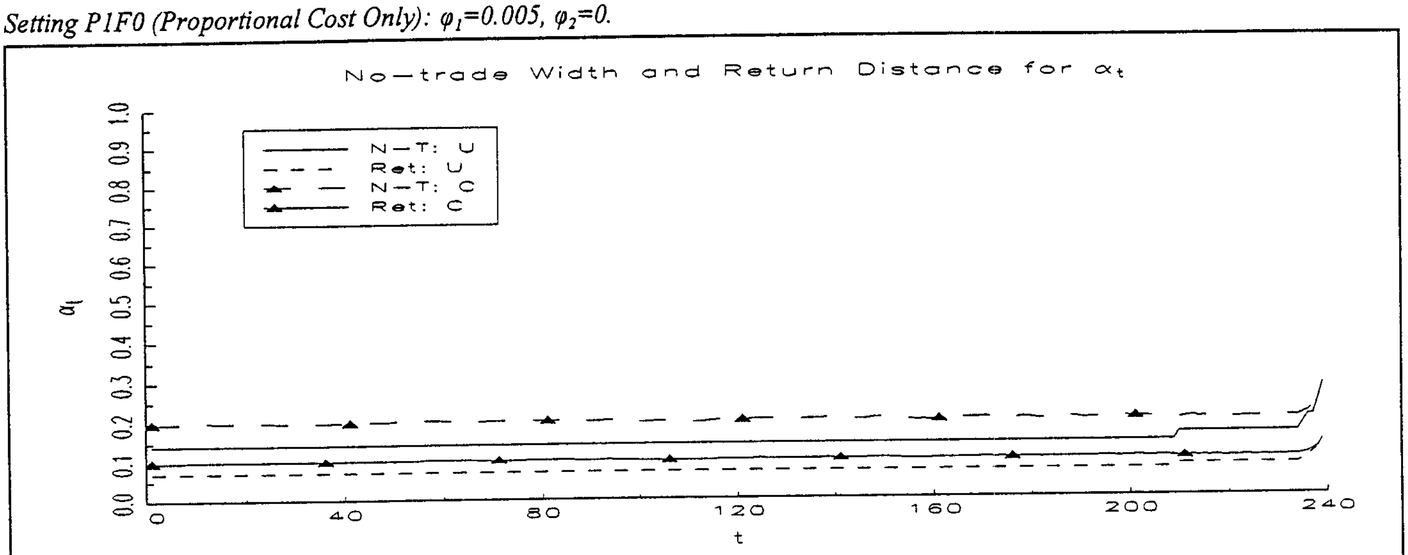

Setting P0F1 (Fixed Cost Only): φ_1=0, φ_2=0.0001.

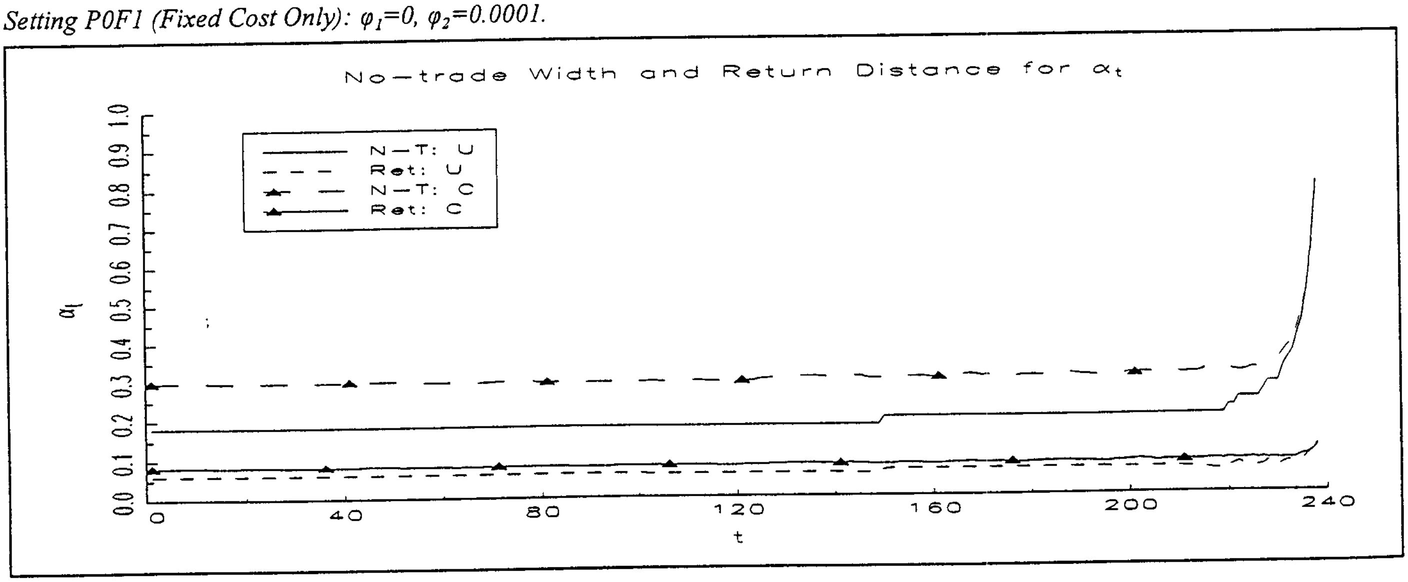

Setting P1F1 (Proportional and Fixed Costs): φ_1=0.005, φ_2=0.0001.

Figure 3. No-trade Width and Return Distance: Unconditional vs Conditional Portfolio Choices. The width of the no-trade region (N-T) and the average return distance (Ret) are plotted for the standard investor, si, (CRRA utility, $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption) facing return-generating process S0 (quadrature approximation to U.S. data) and making unconditional (U) and conditional (C) portfolio choices. The investor problem is described in Section I while section II describes the quadrature approximation.

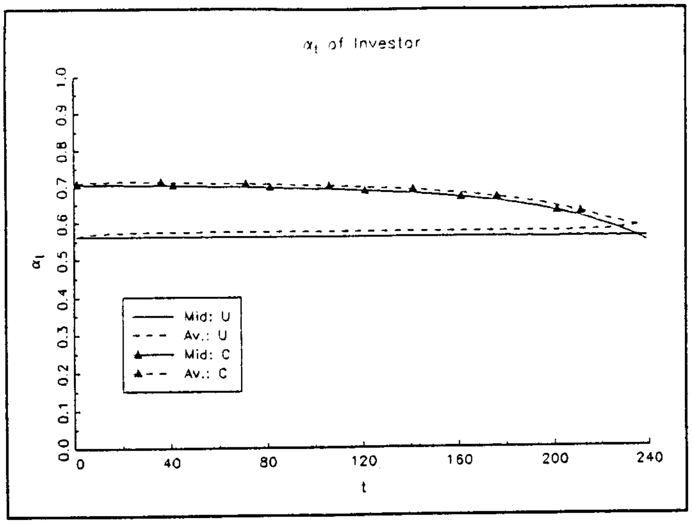

Setting P1F0 (Proportional Cost Only): $\varphi_1=0.005$, $\varphi_2=0$.

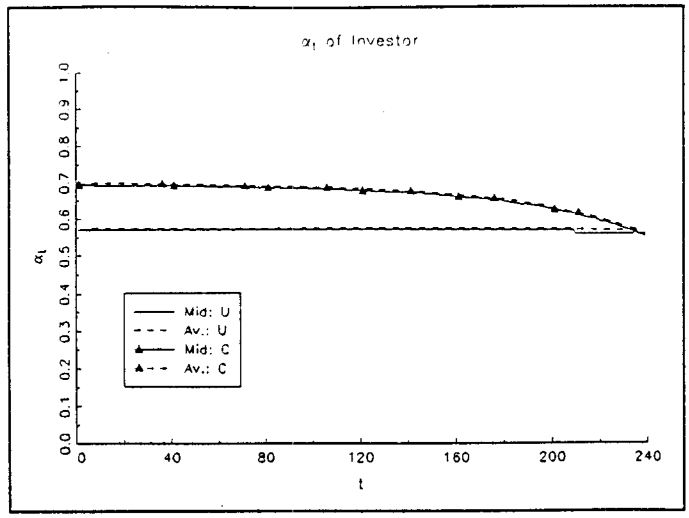

Setting P0F1 (Fixed Cost Only): $\varphi_1=0$, $\varphi_2=0.0001$.

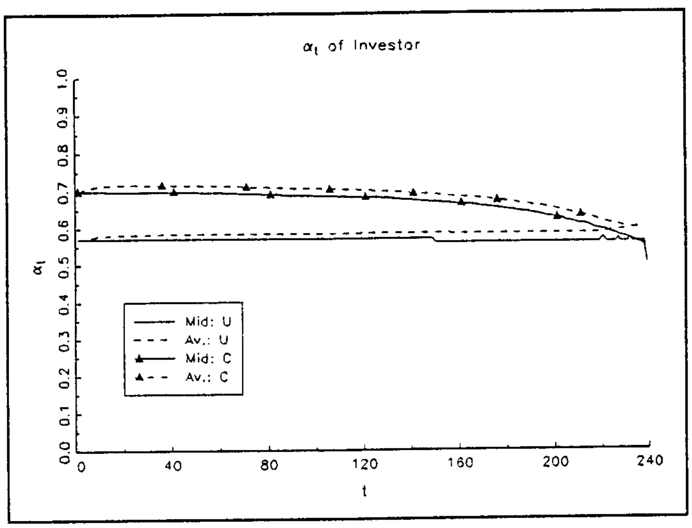

Setting P1F1 (Proportional and Fixed Costs): $\varphi_1=0.005$, $\varphi_2=0.0001$.

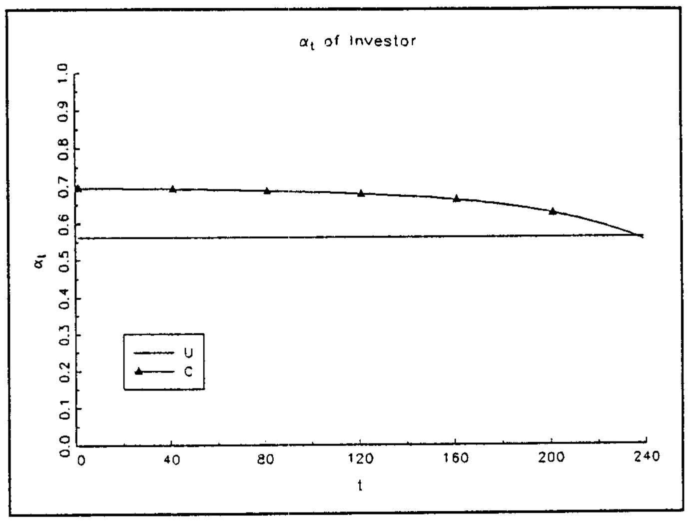

Setting P0F0 (No Costs): $\varphi_1=0$, $\varphi_2=0$.

Figure 4. Average Risky Asset Holding and the No-trade Midpoint: Unconditional vs Conditional Choices. The average holding of the risky asset and the midpoint of the no-trade region are plotted for the standard investor, si, (CRRA utility, $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption) facing return-generating process S0 (quadrature approximation to U.S. data) and making unconditional (U) and conditional (C) portfolio choices. The investor problem is described in Section I while section II describes the quadrature approximation.

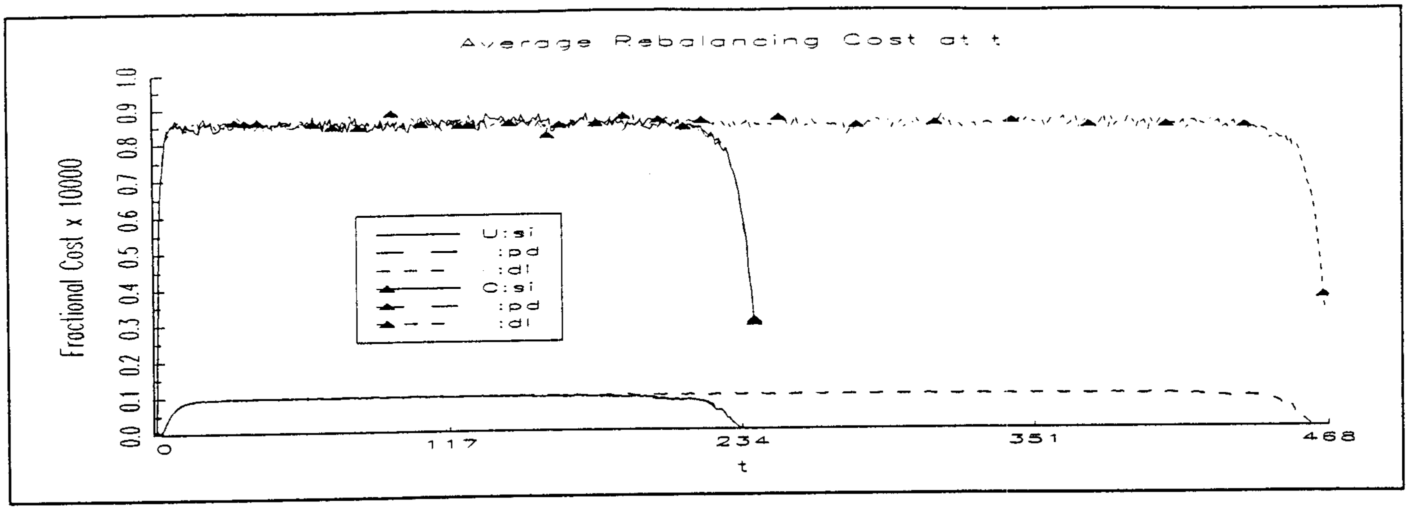

Average Rebalancing Cost at t

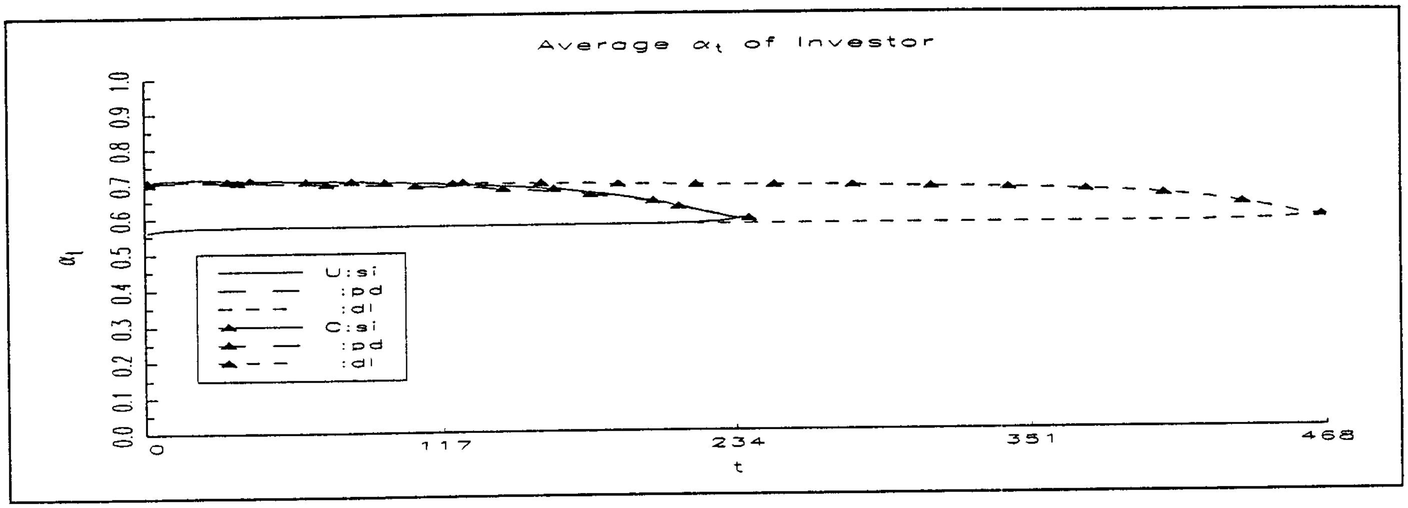

Average α_t

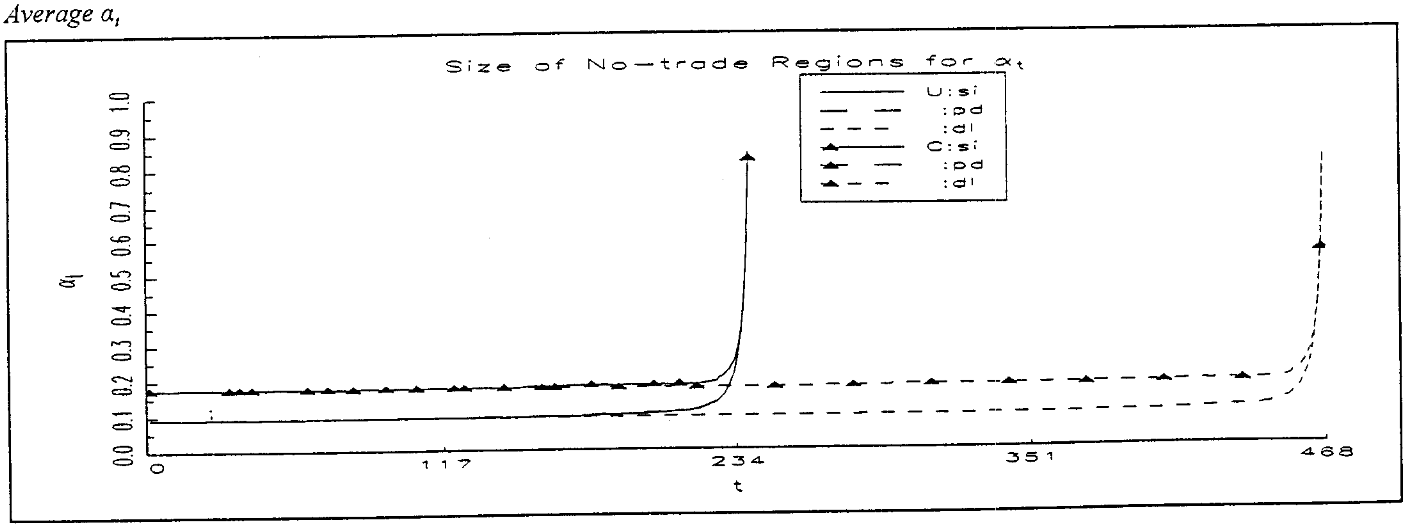

Width of No-trade Region for α_t

Figure 5. Portfolio Choice Parameters: Impact of a Positive Death Probability.

Portfolio choice parameters are plotted for three investor profiles described in Table 1: standard "si" (CRRA utility, $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption); "pd"; and, "dl". The transaction cost setting is P1F0 from Table 1: no fixed cost and proportional cost (φ_1) set equal to 0.5%. The "pd" and "si" investors are the same except "pd" faces death probabilities comparable to those for a U.S. man who is 81 years and 1 month old at t=1. The "dl" and "pd" investors only differ in one respect: "pd" dies upon reaching 101 years or t=240, while "dl" faces death probabilities comparable to those for a U.S. man through until 120 years or t=468. The investors face the return-generating process S0 (quadrature approximation to U.S. data) and make unconditional (U) and conditional (C) portfolio choices. The investor problem is described in Section I while section II describes the quadrature approximation.

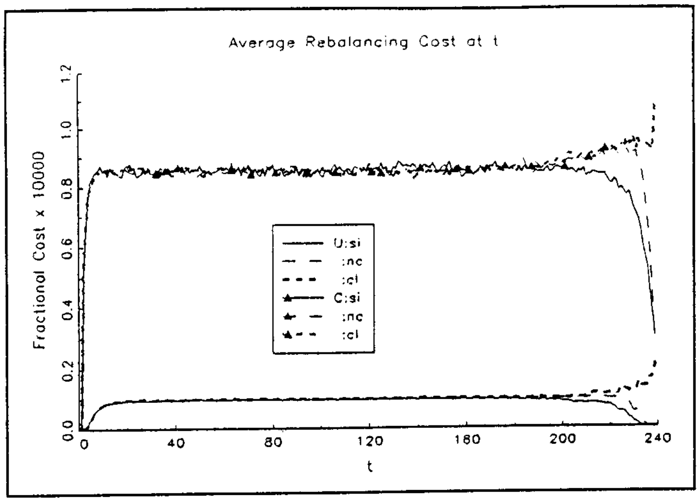

Average Rebalancing Cost at t

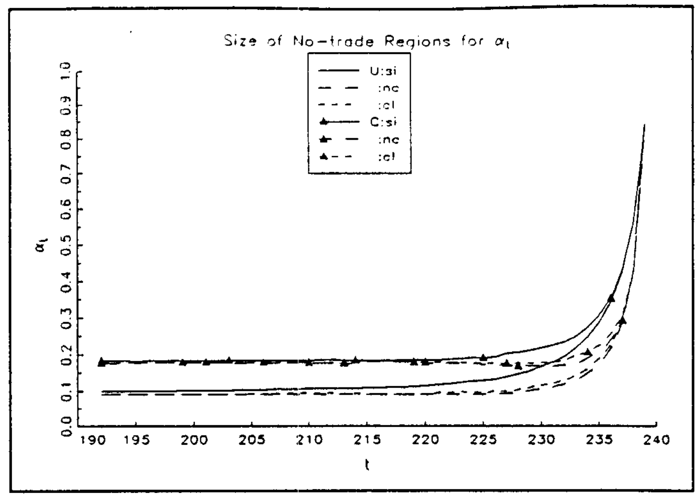

Width of No-trade Region for α_t

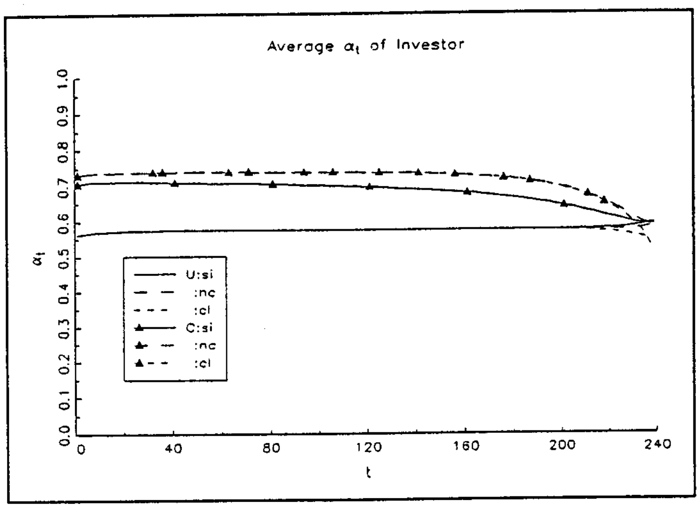

Average α_t

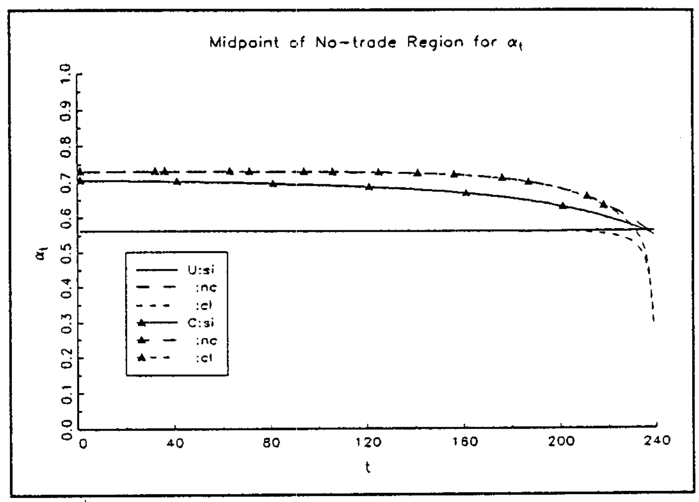

Midpoint of No-trade Region at t

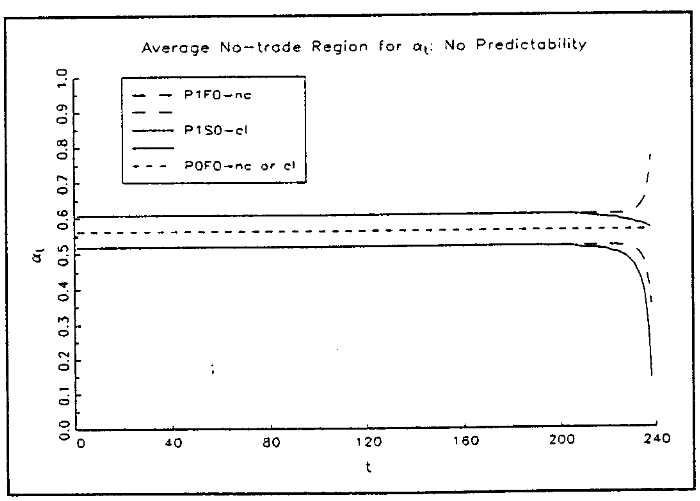

No-trade Region at t: Unconditional Portfolio Choice

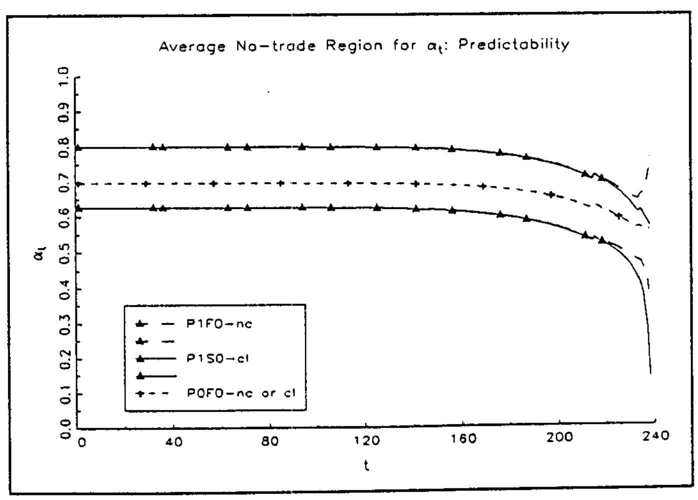

No-trade Region at t: Conditional Portfolio Choice

Figure 6. Portfolio Choice Parameters: Impact of Intermediate Consumption and Liquidation Costs. Portfolio choice parameters are plotted for three investor profiles described in Table 1: standard "si" (CRRA utility $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption); "nc"; and, "cl". The transaction cost setting is P1F0 from Table 1: no fixed cost and proportional cost (φ_1) set equal to 0.5%. The "nc" and "si" investors are the same except "nc" has intermediate consumption, while "cl" and "nc" only differ in one respect: the cost of liquidating the risky asset is 0.5% for "cl" vs. 0% for "nc". The investors face the return-generating process S0 (quadrature approximation to U.S. data) and make unconditional (U) and conditional (C) portfolio choices. The investor problem is described in Section I while section II describes the quadrature approximation.

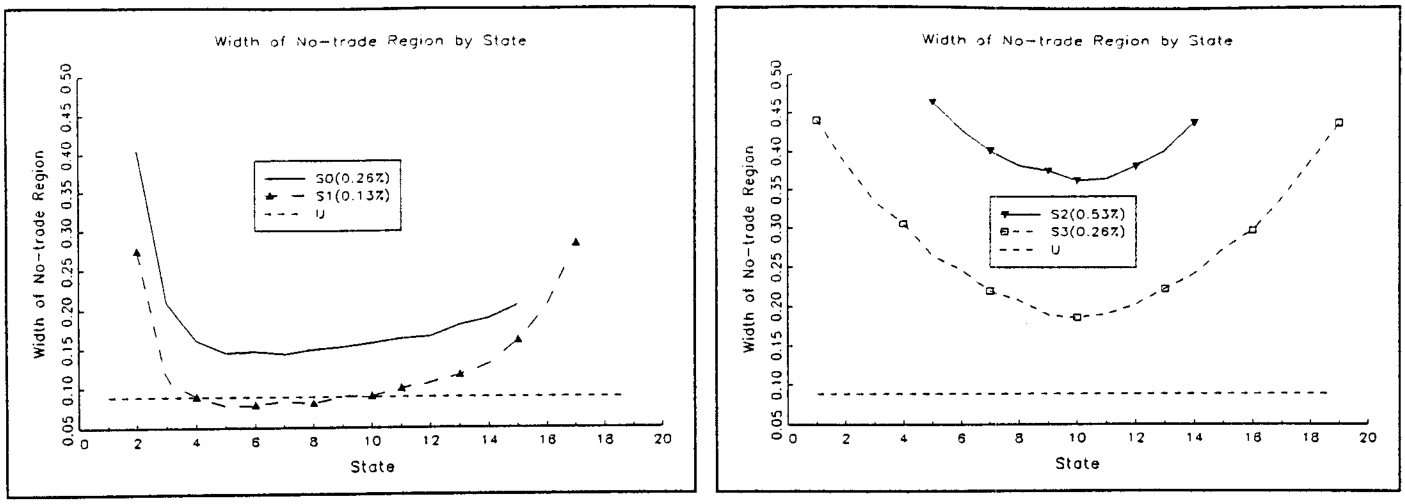

Varying $b_1\, \sigma_d$ with b_2 and ρ_{ev} fixed.

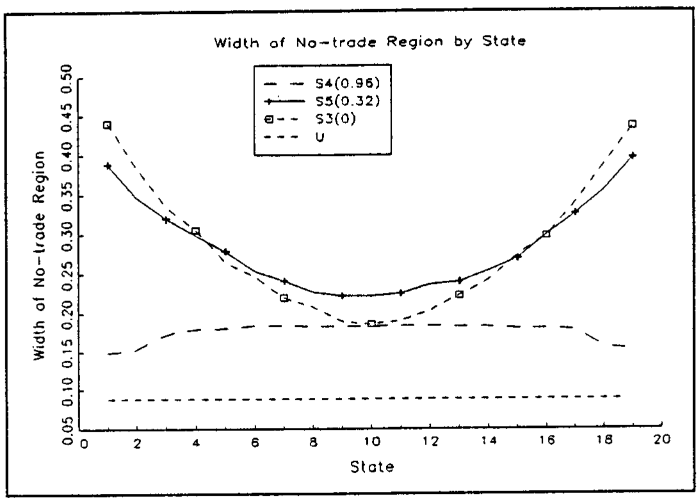

Varying b_2 with $b_1\, \sigma_d$ and ρ_{ev} fixed

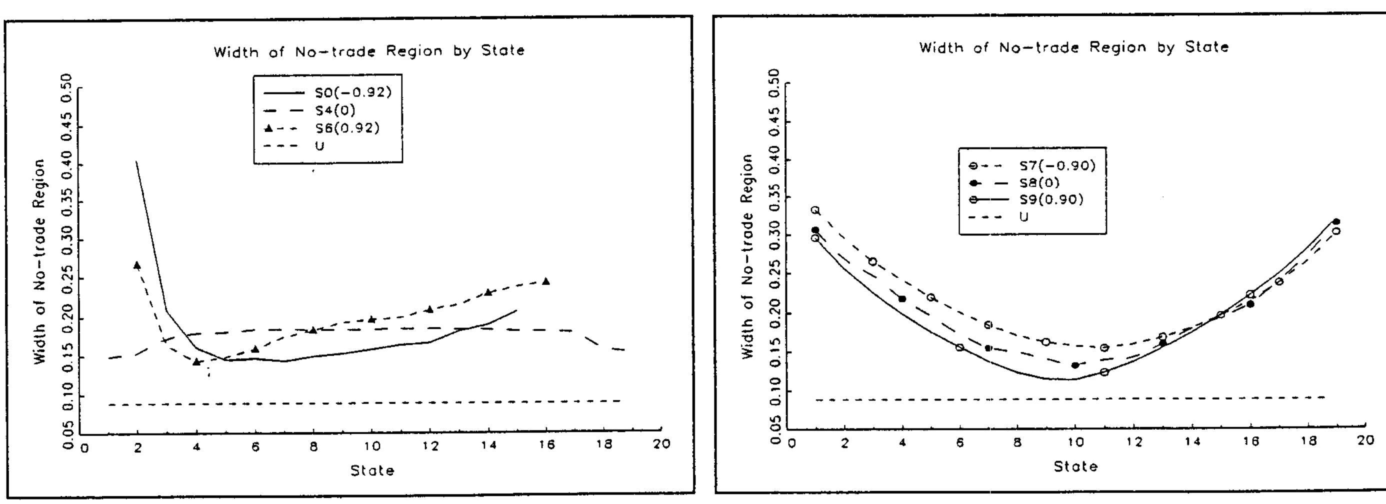

Varying ρ_{ev} with $b_1\, \sigma_d$ and b_2 fixed.

Figure 7. Width of No-trade Region by State at t=1: Varying Return Parameters.

The width of the no-trade region is plotted by state at t=1for the standard "si" investor (CRRA utility $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption) making conditional and unconditional (U) portfolio choices. The transaction cost setting is P1F0 from Table 1: no fixed cost and proportional cost (φ_1) set equal to 0.5%. Parameters of the return-generating processes are summarized in Panel C of Table 1 while section II describes the quadrature approximation. The graphs are organized to facilitate comparative-static analysis for each return parameter: the magnitude of the single-period predictability ($b_1\, \sigma_d$), the persistence of the predictability (b_2), the hedgeability parameter ρ_{ev}, and the heteroscedasticity parameter (ζ_1). The value of the parameter being varied is in brackets for each process. The investor's problem is described in section I.

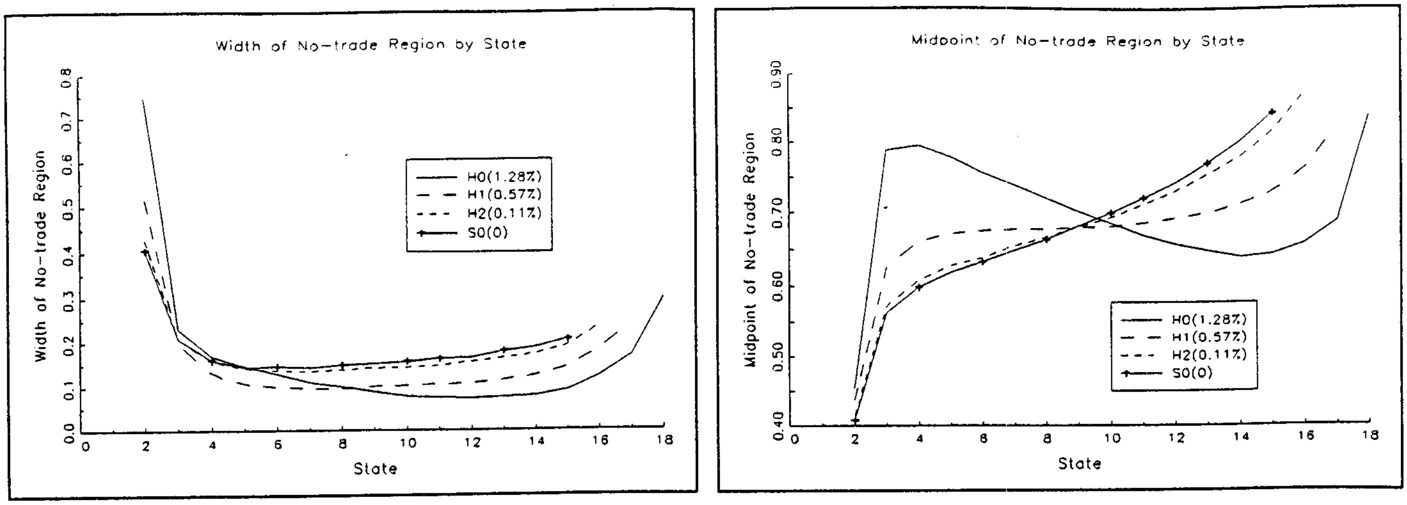

Width of No-trade Region by State *Midpoint of No-trade Region by State*

Figure 8. Width and Midpoint of No-trade Region at t=1 by State with Conditional Choices: Introducing Heteroscedasticity.

The width and midpoint of the no-trade region for t=1 are plotted by state for the standard "si" investor (CRRA utility, $\beta = 1/R_f$, $\gamma = 4$, intermediate consumption) making conditional portfolio choices. The transaction cost setting is P1F0 from Table 1: no fixed cost and proportional cost (φ_1) set equal to 0.5%. Parameters of the four return-generating processes are summarized in Panel C of Table 1 while section III describes the quadrature approximation. The heteroscedasticity parameter, ζ_1, which measures the extent to which conditional volatility varies with condition mean, ranges from 0 to high moving across return-generating processes: S0 to H2 to H1 to H0. The value of ζ_1 is in brackets for each process. Details of the calculation of ζ_1 are contained in section II while the investor's problem is described in section I.