Two Birds, One Stone: Joint Timing of Returns and Capital Gains Taxes

Yaoting Lei [*] Ya Li [†] Jing Xu [‡]

November 2, 2017

Abstract

Since realized returns and future expected returns of financial assets tend to move along opposite directions, tax-timing incentive can be in sharp conflict with market-timing incentive. This paper studies a portfolio choice model with return predictability, and characterizes the optimal policy that jointly times market returns and capital gains taxes. The calibrated model suggests that return predictability significantly increases both the utility loss due to capital gains taxes and the value of tax-deferral, and these findings shed new light on the "asset location puzzle." We further measure the costs of separately conducting market- and tax-timing, and examine the welfare implication of the optimal policy through an empirical simulation.

Keywords: Portfolio Choice, Return Predictability, Capital Gains Tax.

JEL Classification: G11, H24, K34.

[*] Yaoting Lei, School of Economics and Management, Nanchang University, 999 Xuefu Road, Nanchang, China 330031, Email: leiyaoting@ncu.edu.cn.

[†] Ya Li, Department of Economics and Finance, Hang Seng Management College, Hang Shin Link, Siu Lek Yuen, Shatin, New Territories, Hong Kong. Email: leah.y.l.imperial@gmail.com.

[‡] Jing Xu (Corresponding author), School of Finance, Renmin University of China, 59 Zhongguancun Street, Beijing, China 100872, Email: jing.xu@ruc.edu.cn. This draft is preliminary and comments are welcome.

Two Birds, One Stone: Joint Timing of Returns and Capital Gains Taxes

Abstract

Since realized returns and future expected returns of financial assets tend to move along opposite directions, tax-timing incentive can be in sharp conflict with market-timing incentive. This paper studies a portfolio choice model with return predictability, and characterizes the optimal policy that jointly times market returns and capital gains taxes. The calibrated model suggests that return predictability significantly increases both the utility loss due to capital gains taxes and the value of tax-deferral, and these findings shed new light on the "asset location puzzle." We further measure the costs of separately conducting market- and tax-timing, and examine the welfare implication of the optimal policy through an empirical simulation.

1 Introduction

Numerous empirical studies have documented predictable variations in equity returns.[1] The literature finds that exploiting these time-variations through "market-timing," i.e., choosing the appropriate time to tilt-up or reduce market exposure, can generate handsome profits for stock investors.[2] However, selling or purchasing asset shares as responses to market-timing incentives can have immediate or future tax consequences.[3] As capital gains tax rate can be as high as 35-40%, these tax payments can take big bites on investors' profits. Therefore, "tax-timing," i.e., choosing the right time to realize gains, has also proved to be crucial for successful investments.

Unfortunately, market- and tax-timing incentives are generally not in line with each other. A conflict between them exists due to the *empirically* estimated negative correlation between the shocks on asset returns and on economic variables that predict future returns.[4] Due to this negative correlation, positive shocks on asset prices are likely to be accompanied by negative shocks on future expected returns. If an asset's price experiences a series of positive shocks, the investor is likely to reach capital gains status, and the tax-timing incentive encourages her to defer the realization of capital gains to save the time value of taxes. Meanwhile, the expected return is likely to be reduced by a series of negative shocks, and the market-timing incentive will encourage her to reduce position as a response to the reduced future expected return. This conflict makes the joint-timing of returns and capital gains

[1] See, for example, Keim and Stambaugh (1986), Campbell (1987), Campbell and Shiller (1988a, 1988b), Fama and French (1988, 1989), Campbell (1991), Lettau and Nieuwerburgh (2008), and Cochrane (2008), among many others.

[2] See Kandel and Stambaugh (1996), Barberis (2000), Xia (2001), Diether et al. (2009), Wachter and Warusawitharana (2009), Lynch and Tan (2010).

[3] Our study is focusing on household investors who are directly subject to capital gains taxes, rather than large institutional investors who are exempt from these taxes.

[4] For example, the empirical estimates of Campbell and Viceira (1999), Stambaugh (1999), Barberis (2000), Xia (2001) and Huang and Liu (2007) all document negative correlations between the changes in stock market returns and the changes in the predictive variables considered. These estimates typically range from -0.62 to -0.93, depending on the predictive variable and sample period.

taxes highly nontrivial.

Although portfolio selection problems with either return predictability or capital gains taxes have received a great deal of attention, the interaction between these two timing incentives, and the resultant economic implications, are rarely examined in the literature. To our best knowledge, this study is the first one to examine the joint timing of market returns and capital gains taxes. We consider the problem of a small investor (i.e., a price-taker whose trading activities have no significant price impact) whose objective is to maximize the expected utility she derives from her net wealth level at some finite horizon. The investor allocates her financial wealth between a risky asset (the stock) and a risk-free money market account (the bond). The stock has time-varying expected returns that are predicted by an economic variable.[5] Thus, the investor has a market-timing incentive to take advantage of the time-varying investment opportunity. Meanwhile, in accordance with tax codes commonly enforced in most countries, the investor is subject to capital gains taxes when she sells stock shares. Consequently, the investor also has a tax-timing incentive to defer the realization of gains and save the time value of taxes.

In order to examine its economic implications, we calibrate our model to the moments of U.S. stock market data during the post-war period. We follow Xia (2001) and use the dividend yield as the major predictive variable of stock market returns. With this empirical calibration, we explore the answers to the following questions: How does the optimal trading policy respond to the investor's market- and tax-timing incentives? How does it mitigate the potential conflict between market- and tax-timing? How does the return predictability affect the welfare loss due to capital gains taxes and the value of tax-deferral? How costly is it for the investor to conduct market- and tax-timing in separated manners?

We show that the optimal trading policy includes a tax-timing component, a market-

[5] In reality, multiple variables, including past returns, dividend yield, earnings-price ratio, nominal interest rates and expected inflation, may exhibit predictive power on future equity returns. In this paper, for tractability reason, we consider the case with a unique return predictor.

timing component, and a hedging demand component simultaneously. The tax-timing component is characterized by a tax-deferral region in the domain of gains. The investor should not trade the stock when her portfolio lies in this region so that she can defer the realizations of capital gains to save the time value of taxes. The market-timing component is characterized by the changes in the location of the tax-deferral region with respect to the return predictor. The tax-deferral region shifts upward when the value of the return predictor increases, indicating increased stock exposure. The hedging demand component is characterized by an overall increase in stock exposure, compared with that implied by a myopic trading policy.[6] Moreover, to mitigate the conflict between market and tax-timing so as to control the amount of taxes, the optimal tax-deferral region is much wider than the tax-deferral region implied by a model without return predictability.

We find that in the presence of return predictability, the investor pays substantially more taxes during her course of trading. In the baseline case, the expected tax payments can be 90% higher than those in the absence of return predictability. This is because, firstly, the hedging demand induced by return predictability increases the investor's allocation in the stock. As it weakens over time, the investor gradually reduces her stock exposure and such selling scheme can generate heavy tax bills. Secondly, the negative correlation between return shocks and predictor shocks increases the propensity of gains realization, hence increases the amount of tax payments. In addition to the heavy tax bills, another source of costs arises in the presence of return predictability: the incentives to defer capital gains realizations can deter the efficiency of market-timing, and impose opportunity costs on the investor. Due to these reasons, capital gains taxes should be more costly to investors when return predictability is taken into account. We find that with return predictability, the certainty equivalent wealth loss (CEWL) due to capital gains taxes can be about 100% higher than

[6] In other words, the investor on average allocates more wealth in the market. This is because the correlation between the shocks on market returns and the shocks on future expected returns is negative.

that without return predictability. This result suggests that models that do not incorporate return predictability may underestimate the true welfare costs of capital gains taxes.

When facing capital gains taxes, the investor has an option to defer the realization of capital gains and save the time value of taxes. An interesting question is whether the conflict between market- and tax-timing incentives reduces the value of tax-deferral. In order to answer this question, we calculate the value of tax-deferral by comparing the value function associated with the optimal trading policy and the value function associated with a suboptimal policy that never defers capital gains realizations. We find that return predictability in fact increases the value of tax-deferral. This seemingly surprising result can be attributed to two reasons. First, the hedging demand induced by return predictability makes the investor allocate more wealth in the stock, and hence the tax-deferral option enables her to save more time value of taxes. Second, and more importantly, the strategy that never defers the realizations of capital gains incurs excessive tax bills, especially when the investor responds strongly to her market-timing incentives.

One of the central messages of our study is that market- and tax-timing should be jointly solved in an investor's portfolio choice problem. To lend support to this argument, we calculate the investor's utility loss if she adopts a suboptimal trading policy which responds to market- and tax-timing incentives in separated manner. We find that adopting this policy can result in a CEWL of 3%. These utility costs arise because this policy fails to mitigate the conflict between market- and tax-timing incentives. As a result, the investor pays more taxes than she would pay if she adopts the optimal policy.

Furthermore, we examine the potential welfare implications of the optimal trading policy through an empirical simulation. We run a horserace during the entire sample period between four trading strategies: the buy-and-hold strategy; Merton's constant ratio strategy; a market-timing strategy that does not time taxes; and the optimal trading strategy implied by our model. The buy-and-hold strategy and Merton's strategy do not exploit return predictability and lead to substantially lower returns. Although the market-timing strategy successfully times the variations in market returns, it fails to control the tax payments which significantly reduces net returns. In comparison, the optimal trading strategy implied by our model outperforms the other three trading strategies by generating a much higher net return. Therefore, conducting joint timing of returns and taxes can also be empirically beneficial for investors.

Our analyses have important empirical implications. The heavy tax burden and the increased value of tax-deferral associated with return predictability may shed new light on the well-known "asset location puzzle," i.e., the phenomenon that a significant portion of investors hold mixed portfolios of bonds and stocks in their taxable and tax-deferred accounts (cf. Amromin (2003)). The literature argues that such an asset location rule is tax-inefficient, as the interests of bonds are usually taxed heavier than the capital gains earned from stocks (cf. Shoven and Sialm (2003), Dammon et al. (2004) and Fischer and Gallmeyer (2017)). However, our analyses imply that if the benefits of market-timing in a tax-deferred account exceed the costs of holding bonds in a taxable account, it can be rational for the investor to locate stocks in the tax-deferred account and bonds in the taxable account. As a result, the incentives behind the asset location puzzle deserve further investigation.[7]

The remainder of this paper is organized as follows. In the next section, we review the literature related to our study. In Section 3, we present our theoretical framework. In Section 4, we perform an empirical calibration and conduct a comprehensive set of numerical analyses to examine the major economic implications of our model. We conclude the paper in Section 5. All technical issues are relegated to the Appendix.

[7] Garlappi and Huang (2006) argue that such allocation can be optimal if the investor faces portfolio constraints.

2 Related Literature

Our paper connects two strands of literature: the literature on portfolio selection with return predictability, and the literature on portfolio selection with capital gains taxes.

Portfolio Selection with Return Predictability. Campbell and Vicera (1999), Barberis (2000), and Lynch (2001) examine the implications of return predictability for portfolio selection in discrete-time settings. Xia (2001) considers the continuous-time case and examines the effect of uncertainty on the slope of predictive regression. These papers, however, generally assume frictionless markets. Lynch and Tan (2010) extend Lynch (2001) to a case with two stocks and transaction costs. Recent advances in this field include Branger et al. (2013), Garleanu and Pedersen (2013), Tsai and Wu (2015) and Moallemi and Saglam (2017).

Portfolio Selection with Capital Gains Taxes. Under an exact cost basis system, Constantinides (1983, 1984), Dybvig and Koo (1996), and DeMiguel and Uppal (2005) study portfolio selection problems with capital gains taxes in discrete-time frameworks. However, they can only solve the problems for a few time steps, due to the strong path dependency introduced by the exact cost basis system. To overcome this difficulty, Dammon et al. (2001, 2004) introduce an average cost basis system, which proves to be a good approximation of the exact tax basis system (cf. DeMiguel and Uppal (2005), and Dai et al. (2015)).[8] Marekwica (2012) and Ehling et al. (2013) further account for the limited use of losses to claim a rebate, as stipulated by the U.S. tax code. Ben Tahar et al. (2007, 2010) formulate a continuous-time version of the model proposed in Dammon et al. (2001, 2004). Dai et al. (2015) extend Ben Tahar et al. (2010) to incorporate the important asymmetry between tax rates for long-term and short-term investments.

[8] For an example of the exact cost basis system and the average cost basis system, the interested readers are referred to footnote 10 and 16 in Dai et al. (2015).

The aforementioned papers assume constant investment opportunity sets. Cai et al. (2017) extend Ben Tahar et al. (2007, 2010) to a two-state regime switching market. Nonetheless, they assume a *null* correlation between the changes in realized returns and in investment opportunity set, despite that most predictive regressions have documented negative correlations between these changes.

This Paper. Our study differs from the above studies by *simultaneously* considering the effect of return predictability and capital gains taxes. Importantly, our model incorporates the negative correlation between changes in price and in investment opportunity set, which may cause a crucial conflict between market- and tax-timing incentives.

3 The Model

This section presents a framework for examining how an investor makes portfolio choice when she is facing return predictability and capital gains taxes. This framework is primarily based on Dammon et al. (2001, 2004) and Ben Tahar et al. (2007, 2010).

3.1 Basic Setup

Assets Market. We assume that time is continuous and indexed by $t \geq 0$. The investment opportunity set includes a risk-free money market account that grows at a constant after-tax interest rate of $r \geq 0$, and a risky asset[9] ("stock" hereinafter) whose cum-dividend value S_t is assumed to evolve according to the following stochastic process

$$\frac{dS_t}{S_t} = (\mu_0 + \mu_1(z_t - \bar{z}))dt + \sigma_S dB_t^S, \tag{1}$$

[9] In this paper, we interpret the risky asset as an exchange traded fund (ETF) that represents a diversified portfolio of the stock market, like Dai et al. (2015).

where z_t is a variable that predicts the stock's expected returns, $\bar{z}$ is the stationary long-term mean value of z_t, μ_1 is the loading of the expected returns on this variable, and μ_0 is the stock's long-term average return. We assume z_t follows an Ornstein-Uhlenbeck process with a mean-reversion level of $\bar{z}$ and speed of $g_1 > 0$

$$dz_t = g_1(\bar{z} - z_t)dt + \sigma_z dB_t^z. \tag{2}$$

In Equation (1) and (2), (B_t^S, B_t^z) is a two-dimensional Brownian motion, defined on a complete probability space $(\Omega, \mathcal{F}, P)$, with constant correlation coefficient $\rho \in [-1, 1]$; other parameters are all assumed to be constant.

According to current tax codes, selling stock shares is subject to capital gains taxes. In the U.S., the amount of realized capital gains or losses is determined by the exact price at which shares are sold and the exact *past* price at which each sold shares were purchased. To maintain tractability, we approximate this exact cost basis by the average cost basis of the current stock position, as does most of the existing literature.[10] In other words, the amount of capital gains or losses is calculated by subtracting the weighted average purchase price of the current stock holding from the actual sale price. We assume the tax rate on capital gains is a constant $\tau \geq 0$.

When selling stock with a loss, the investor is eligible for a capital loss rebate. We follow Dai et al. (2015) and consider two simplified cases: the full rebate (FR) case and the full carry-forward (FC) case. In the FR case, the investor can claim full rebates on her capital losses; in the FC case, losses can only be carried forward indefinitely to offset future capital gains. Therefore, capital gains and losses are treated symmetrically in the FR case, and asymmetrically in the FC case.

[10] See, for example, Dammon et al. (2001, 2004), Gallmeyer et al. (2006), Dai et al. (2015), and Cai et al. (2017). Also note that the average cost basis system is used in Canada.

State Variables. Let x_t be the dollar amount invested in the money market account, y_t be the dollar amount invested in the stock, and k_t be the total cost basis of the stock holding. Then, (x_t, y_t, k_t, z_t, t) is the vector of state variables in our model. We have the following dynamic budget constraints

$$
\begin{aligned}
dx_t &= rx_t dt - dL_t + f(0, y_{t-}, k_{t-}; l) dM_t, \quad (3) \\
dy_t &= (\mu_0 + \mu_1(z_t - \bar{z})) y_t dt + \sigma_S y_t dB_t^S + dL_t - y_{t-} dM_t, \quad (4) \\
dk_t &= dL_t - k_{t-} dM_t + l \left(k_{t-} - y_{t-} \right)^+ dM_t, \quad (5)
\end{aligned}
$$

where M_t and I_t are non-decreasing processes, with $0 \leq dM_t \leq 1$ representing the fraction of the current stock position that is sold, and $dI_t \geq 0$ representing the dollar amount of new stock shares purchased. Besides,

$$
f(x_t, y_t, k_t; l) = x_t + y_t - \tau \left[(1 - l) (y_t - k_t) + l (y_t - k_t)^+ \right] \quad (6)
$$

is the investor's net wealth at time t. l is a parameter indicating the capital losses rebate rule: $l = 0$ or 1 corresponds to the FR case or the FC case, respectively.[11] Note that Equation (5) implies that in the FR case, the cost basis of each stock share remains unchanged, and the total cost basis of the stock position is proportionally reduced at the time of selling.

The Investor's Problem. The investor's objective is to choose the optimal trading strategy (L_t^*, M_t^*) to maximize the expected constant relative risk aversion (CRRA) utility she

[11] When $l = 0$, the expression for net wealth reads $f(x_t, y_t, k_t; l) = x_t + y_t - \tau (y_t - k_t)$, indicating that if $y_t < k_t$, then the capital losses $(y_t - k_t)$ are rebated at rate τ. When $l = 1$, the expression for net wealth reads $f(x_t, y_t, k_t; l) = x_t + y_t - \tau (y_t - k_t)^+$, indicating that only capital gains are taxed, while capital losses are not rebated to the investor. Instead, capital losses can only be used to increase the total cost basis and hence reduce future capital gains (as implied by Equation (5)).

derives from the net wealth level at finite horizon $T > 0$; that is,

$$\max_{L_t, M_t} E\left[u(f(x_T, y_T, k_T; l))\right], \tag{7}$$

subject to the solvency constraint $f(x_t, y_t, k_t; l) \geq 0$ and the short-sale restriction $y_t \geq 0$ for all t, where

$$u(W) = \frac{1}{1-\gamma} W^{1-\gamma} \tag{8}$$

is the CRRA utility function, and $\gamma > 0$ with $\gamma \neq 1$ is the investor's relative risk-aversion coefficient.[12]

3.2 Discussion of the Assumptions

Before proceeding to solve the investor's problem, we first briefly discuss some simplifying assumptions made in our model and their potential implications.

Symmetric Tax Rates for Long-Term/Short-Term Investments. In our model, we ignore the differential taxation of short- and long-term capital gains/losses.[13] Incorporating asymmetric tax rates requires us to introduce another state variable, which is the holding period of the current stock position. Unfortunately, such an increase in dimensionality would make the problem computationally infeasible.[14]

Our paper focuses on the effect of return predictability on the investor's optimal tax-

[12] The short-sale restriction is imposed in most of the literature on portfolio choice with capital gains taxes. Maximizing the objective function (7) is equivalent to maximizing the final net wealth deflating by a constant inflation rate, as assumed in Dammon et al. (2001). Moreover, we do not directly incorporate intertemporal consumption in our model. This allows us to focus on how joint-timing of market returns and capital gains taxes improve the investor's investment performance. We strongly believe that incorporating intertemporal consumption would not change our results significantly, as the optimal trading strategy is likely to remain qualitatively the same in the presence of intertemporal consumption.

[13] In the U.S., short-term gains (holding time less than one year) are taxed at the marginal ordinary income tax rate, while long-term gains (holding time longer than one year) are taxed at a lower rate.

[14] Dammon and Spatt (1996) and Dai et al. (2015) examine the tax-timing problem with asymmetric short-/long-term capital gains tax rates. However, their models assume a constant expected stock return to maintain tractability.

timing policy. As return predictability is unlikely to change with the investor's stock holding period, we strongly believe that most of our results can be extrapolated to the case with asymmetric long-term/short-term tax rates.

Losses Rebate Method. We assume the capital losses are either fully rebated or fully carried forward. The actual U.S. tax code stipulates that an investor is eligible for rebates up to capital losses of $3,000 annually, and the remaining capital losses are carried forward indefinitely to offset future capital gains. However, incorporating this limited use of capital losses would destroy the homogeneity of the problem and make it much more difficult to solve. On the other hand, since the actual capital losses rebate rule interpolates between the two extreme scenarios examined in our paper, we believe our model can provide reasonable upper and lower bounds of the economic effects we are interested in.[15]

Constant Level of Return Predictability. We assume that the loading of the stock's expected returns on the predictor, i.e., μ_1, is constant. However, as argued by Xia (2001), *"The controversy surrounding stock return predictability is symptomatic of the fact that the predictive relation is quite uncertain."* When the loading parameter is taking a large (small, resp.) value, the investor is likely to respond more (less, resp.) actively to the changes in the predictor. Therefore, the effect of return predictability is likely to increase with the value of the loading parameter μ_1. As we will calibrate the model using a full sample of historical data during the post-war era in our numerical analyses, the results reported in this paper may well represent an average effect of the presumably time-varying return predictability.

[15] For the effects of limited use of capital losses, interested readers are referred to Marekwica (2012) and Ehling et al. (2013). Dai et al. (2015) argue that the FR case is a more suitable model for less wealthy investors, while the FC case is more suitable for wealthy investors.

3.3 The HJB Equation

We now proceed to characterize the solution to the investor's problem. As the presence of capital gains taxes renders the market incomplete, we utilize the direct method of dynamic programming. For this purpose, we define the investor's value function as

$$J(x, y, k, z, t) = \max_{(L_s, M_s): s \geq t} E_t \left[u[f(x_T, y_T, k_T; l)] \right], \tag{9}$$

where the expectation is taken under filtration $\mathcal{F}_t$. With the assumption of sufficient regularity of the value function, it can be shown that $J(x, y, k, z, t)$ must satisfy the following Hamilton-Jacobi-Bellman (HJB) equation:

$$\max \left\{ \mathcal{L}J + \partial_t J, \;\; \mathcal{B}J, \;\; \mathcal{S}J \right\} = 0 \tag{10}$$

on the domain $\Omega = \{(x, y, k, z, t) : x \in R, y \in R^+, k \in R^+, z \in R, t \in [0, T], f(x, y, k; l) \geq 0\}$, with terminal condition

$$J(x, y, k, z, T) = u(f(x, y, k; l)). \tag{11}$$

The differential operators in Equation (10) are given by

$$\mathcal{L}J = rx\partial_x J + (\mu_0 + \mu_1(z - \bar{z}))y\partial_y J + \frac{1}{2}\sigma_S^2 y^2 \partial_{yy} J + g_1(\bar{z} - z)\partial_z J + \frac{1}{2}\sigma_z^2 \partial_{zz} J + \rho\sigma_S\sigma_z y \partial_{yz} J, \tag{12}$$

$$\mathcal{B}J = -\partial_x J + \partial_y J + \partial_k J, \tag{13}$$

and

$$\mathcal{S}J = f(0, y, k; l)\partial_x J - y\partial_y J + (l(k - y)^+ - k)\partial_k J, \tag{14}$$

respectively, where ∂_t denotes partial derivative with respect to t, etc.

A heuristic derivation of Equation (10) is presented in Appendix A. Equation (10) splits the solution domain Ω into three parts: a buy region in which $\mathcal{B}J = 0$ holds; a sell region in which $\mathcal{S}J = 0$ holds; and a non-trading region in which $\mathcal{L}J + \partial_t J = 0$ holds. The optimal trading strategy is fully characterized by these regions. Detailed characterizations of these regions are presented in the next section.

In addition to the time variable t, the value function J involves four spacial variables and is difficult to solve even numerically. Fortunately, we can reduce the dimensionality of this problem by exploiting the homogeneity of the utility function (8) and the linearity of dynamics (3)-(5). For the sake of brevity, we relegate these analyses to Appendix A.

4 Model Implications

4.1 Data and Model Calibration

Our major purposes are to identify the optimal trading strategy that jointly times market returns and capital gains taxes, and to examine the economic implications associated with such joint timing. For these purposes, we first describe an empirical calibration of the model.

We calibrate our model to the moments of U.S. market returns. First, the risk-free rate r is approximated by the average historical yields of treasury bills matured in exactly one year (i.e., the "constant maturity yield rate") after being adjusted for taxes.[16] We exploit the documented predictability from dividend yield to future stock returns to calibrate the parameters in Equations (1) and (2).[17] Specifically, we obtain monthly returns of the U.S. weighted-average market index inclusive or exclusive of dividend, spanning from January 1950 to December 2016, from the CRSP database. The dividend yields are then factored out

[16] The treasury bill yield data is obtained from the Federal Reserve Bank, and is available from 1960 to 2016. Following Dammon et al. (2001), we assume a tax rate of 0.36 on the interest to adjust the treasury bill yield.

[17] Research has found that other macro-economic variables, such as the earnings-price ratio, nominal interest rates, and expected inflation, can also predict future market returns. Motivated by Lettau and Ludvigson (2001), we also implement a calibration based on cay_t and find quantitatively similar results.

from these two return series.[18] The average value and standard deviation of the historical market returns are used to estimate μ_0 and σ_S, respectively. Other parameter values are estimated through a VAR regression, with details presented in Appendix D.

We obtain the following baseline parameter values: risk-free rate $r = 0.031$, long-term expected return $\mu_0 = 0.116$, slope parameter $\mu_1 = 4.397$, volatility of stock returns $\sigma_S = 0.147$, average dividend yield $\bar{z} = 0.026$, mean reversion speed $g_1 = 0.141$, volatility of dividend yield $\sigma_z = 0.005$, and correlation between return shocks and dividend yield shocks $\rho = -0.895$.[19] As we show in later sections, this negative correlation between return and predictor shocks has important implications.

Moreover, like Xia (2001), we set the investor's relative risk aversion level to five, i.e., $\gamma = 5$, and fix her investment horizon at ten years, i.e., $T = 10$. In the baseline case, we consider a capital gains tax rate of $\tau = 0.25$, which is close to the average capital gains tax rate of long-term and short-term investment for middle-class investors in the U.S.[20] Table 1 summarizes the baseline parameter values described above.

[Table 1 About Here]

The model is solved numerically by fully implicit finite difference method. Since the major purpose of this paper is to examine how the presence of return predictability affects tax-timing and its resultant economic implications, we will frequently compare our model with a benchmark case in which return predictability is eliminated by setting $\mu_1 = 0$.

---

[18] The complete CRSP data sample begins from January 1926. However, in the subsample of 1926-1949, there is little evidence of return predictability. Therefore, like Xia (2001), we choose the post-war sample to calibrate our baseline model, which explicitly specifies return predictability. Different from Xia (2001), we do not deflate the return series by CPI growth. According to the U.S. tax code, capital gains taxes are levied on nominal quantities rather than real quantities.

[19] We also implement a maximal likelihood estimator as in Huang and Liu (2007), and the estimating results are close to what we obtain using the VAR regression.

[20] We also examine the high tax rate case in which $\tau = 0.35$ and the low tax rate case in which $\tau = 0.15$ to check the robustness of our results. The qualitative results are found to be the same.

4.2 Return Predictability and Hedging Demand

The literature documents that return predictability generates hedging demand in the investor's investment policy and leads to deviation from the myopic investment policy (cf. Campbell and Viceira (1999), Lynch (2001), and Xia (2001)). We first reexamine such hedging demand for completeness. We assume the absence of capital gains taxes for now.

[Figure 1 About Here]

Panel A of Figure 1 plots two stock allocation rules as functions of (z, t). The first rule is the optimal allocation rule, and the second rule is the myopic allocation rule

$$\pi_0(z,t) = \left(\frac{\mu_0 + \mu_1 z - r}{\gamma \sigma_S^2}\right)^+,$$

where we take the positive part because the short-sale restriction is imposed. Generally, these two allocation rules do not coincide, and the difference between them, as depicted in Panel C of Figure 1, reflects the hedging demand induced by return predictability. A positive hedging demand arises in our baseline case due to the negative correlation between return and predictor shocks. Intuitively, when this correlation is negative, a negative shock on the stock price is likely to be compensated by an increase in the expected return, and a negative shock on the expected return is likely to be compensated by a positive shock on the price. This "compensation effect" reduces the riskiness of the stock and induces the investor to allocate more wealth to the stock.[21]

Panels B and D of Figure 1 show the optimal and myopic allocation rules, and the hedging demand, for fixed $z = 0$. It is easy to see that the myopic policy is optimal for the investor when there is no return predictability, i.e., $\mu_1 = 0$ (cf. Merton (1969, 1971)). Therefore, consistent with Campbell and Viceira (1999), return predictability sizably increases the

[21] As can be expected, when the correlation between return and predictor shocks is positive, the hedging demand is negative. However, such a positive correlation is rarely found in empirical studies.

investor's optimal stock exposure. As time approaches the final time, such hedging demand weakens and the optimal allocation converges to the myopic allocation rule.

4.3 Optimal Trading Policies

In this section, we examine the optimal trading strategy in the presence of both return predictability and capital gains taxes. We examine both the static profiles and dynamic properties of the optimal trading strategy.

4.3.1 Static Profiles

[Figure 2 About Here]

We begin our analyses by briefly reviewing the optimal tax-timing policy in the absence of return predictability.[22] We first examine the FR case. Figure 2 plots the selling region (SR), buying region (BR), and non-trading region (NTR) at time $t = 5$. The x-axis and y-axis are the basis-price ratio $b \equiv k/y$ and the fraction of net wealth invested in stock, i.e., $y/f(x, y, k, 0)$, respectively. When the investor has capital gains ($b < 1$) and the weight of the stock in her portfolio is high (low, resp.) enough to enter the interior of SR (BR, resp.), the investor sells (buys, resp.) a minimal amount of stock so that the stock position is pushed to the boundary of NTR, as signified by the arrow from point A to point B (point C to point D, resp.). When the weight of the stock lies in NTR, the investor is better off not trading the stock, indicating an incentive to defer the realization of capital gains. This is consistent with many papers on portfolio selection with capital gains taxes such as Dammon et al. (2001) and Dai et al. (2015). When the investor has capital losses, i.e., $b > 1$, similar to Dammon et al. (2001) and Dai et al. (2015), the investor first sells the entire stock position (as signified by the arrow from point E to point F) to obtain loss rebates, and then

[22] The characterizations of optimal tax-timing strategy are well documented in the literature; see, for example, Dai et al. (2015) and Cai et al. (2017). We present these results in our paper only for completeness.

buys back some stock shares (as signified by the arrow from point F to point G) to rebuild an optimal stock risk exposure.[23]

[Figure 3 About Here]

Next, we examine the optimal tax-timing policy in the presence of return predictability. Figure 3 shows two cross-sections of the optimal trading policy. Panels A-C plot the optimal trading policy at time $t = 5$ for three different values of z: $z = -0.02$, $z = 0$, and $z = 0.02$. When the value of z increases, the stock's instantaneous expected return also increases, and hence the investor desires larger exposure to the stock to extract its higher expected return. As a result, the NTR shifts upward when the value of z increases. These shifts of the tax-deferring region reflect the investor's market-timing incentive.

In Panel B of Figure 3, the instantaneous expected return of the stock equals its long-term average return μ_0. However, comparing Panel B with Figure 2, we find that the NTR clearly shifts upward, indicating a higher desirable exposure to the stock. This pattern results from the hedging demand induced by return predictability, and is consistent with that observed in Figure 1 when capital gains taxes are absent.

Another interesting observation is that the NTR in Panel B of Figure 3 is considerably wider than that in Figure 2. The maximal width of the NTR in Figure 2 is about 0.23, while it increases to about 0.7 in Panel B of Figure 3. The widened NTR indicates that the optimal trading policy mitigates the conflict between market- and tax-timing incentives. As aggressive market-timing can generate excessive tax bills, the investor chooses a wider non-trading region to control the amount of tax payments.

Panels D-F of Figure 3 show the time evolution of the optimal trading policy for a fixed value $z = 0$. Specifically, we plot the snapshots of the NTR at $t = 1$, 5, and 9. Consistent with the decreasing stock exposure observed in Figure 1, the location of the NTR

[23] The investor repurchases because the stock has a positive return in excess of the risk-free rate.

moves downward as time proceeds. This implies diminishing hedging demand over time. Meanwhile, the NTR becomes narrower as time approaches the final horizon, indicating weakening tax-deferral incentives. This is because our model assumes the investor derives utility from the after-tax net wealth at the final horizon.

[Figure 4 About Here]

We now examine the FC case. Figure 4 shows the optimal trading policy when return predictability is absent. Similar to the FR case, when the investor has capital gains, there is a non-trading region inside which the investor should not trade stock to defer taxes. The major difference between the FC and FR cases is the way in which capital losses are treated. In the FC case, the investor does not conduct a wash sale because she cannot claim rebates on losses. Instead, she trades the stock continuously so that the weight of the stock position lies exactly on the black line. In other words, the investor may either buy or sell stock when she has capital losses, depending on whether the stock's weight is higher or lower than a target level.

[Figure 5 About Here]

Finally, Figure 5 shows the optimal policy in the FC case when returns are predictable. The shape of the non-trading region is qualitatively similar to that in the constant return case. However, the overall allocation in the stock is higher due to the hedging demand. As expected, the NTR moves upward when the predictor's value increases, and moves downward when time approaches the final horizon. These patterns are generally consistent with those obtained in the FR case.

Overall, the above results indicate that the optimal trading policy implied by our model is a nontrivial combination of a tax-timing component, a market-timing component, and a component of hedging demand.

4.3.2 Dynamic Statistics

The results presented in the previous section indicate that for each fixed value of return predictor z, the shapes of the optimal trading regions are qualitatively similar to those obtained in the case without return predictability. However, as the locations of these trading regions also vary with z, these static profiles cannot fully reveal the dynamic properties of our model. In this section, we run Monte Carlo simulations to explore some dynamic properties of our model. In particular, we compute the average tax payments discounted by the risk-free rate (as a fraction of the initial wealth) and the average duration between realizations of capital gains implied by the optimal trading strategies. We run the simulations for both the cases with and without return predictability.

[Table 2 About Here]

Table 2 shows the results obtained from 10,000 paths of Monte Carlo simulations. We find that the investor incurs much heavier tax bills when stock returns are predictable. Panel A shows that in the FR case, the average amount of tax payments is 34.3% of the investor's initial wealth when return predictability is absent, and this amount rises sharply to 71.8% when return predictability is present. The average duration between gain realizations reduces from 0.50 years to 0.34 years when return predictability is present. Qualitatively similar results are obtained in the FC case, and are presented in Panel B.

The increased capital gains taxes can be primarily attributed to two sources: (1) the weakening hedging demand makes the investor gradually reduce her stock position over time, and such a selling scheme can generate heavy tax bills; (2) the negative correlation between the shocks on the price and on the return predictor induces the investor to sell, as a response to market-timing incentives, when the stock price experiences positive shocks.[24]

[24] To confirm that the decreasing hedging demand is not the only source of the heavy tax bills, we also consider the logarithm utility case in which the investor exhibits no hedging demand. The results remain qualitatively unchanged.

In summary, when returns are predictable, the incentive to time the market can make the investor realize capital gains more frequently, and result in heavier tax bills. This is the first evidence indicating that return predictability may amplify the costs of capital gains taxes.

4.4 How Costly are Capital Gains Taxes to Investors?

The literature shows that capital gains taxes can significantly reduce the investor's welfare (cf. Poterba (1987)). One obvious reason is that paying capital gains taxes reduces the investor's net return. In the presence of return predictability, there is another source of utility costs due to capital gains taxes: the incentive to defer capital gains realizations can reduce the efficiency of market-timing. Due to a negative correlation between return and predictor shocks, the future expected return is likely to be decreased by a series of negative shocks after the stock's price experiences a series of positive shocks. In such cases, if the investor has capital gains, the incentive to defer capital gains realization may deter the investor's response to the change in expected return. This source of costs reflects an inherent conflict between tax- and market-timing incentives.

In this section, we quantitatively examine the welfare implication of capital gains taxes in the presence of return predictability. We measure the utility costs due to capital gains taxes through the certainty equivalent wealth loss (CEWL) δ_1, which solves

$$J(1-\delta_1, 0, 0, z_0, 0; 0) = J(1, 0, 0, z_0, 0; \tau) \tag{15}$$

In Equation (15), $J(x, y, k, z, t; \tau)$ is the investor's value function when the capital gains tax rate is τ. Therefore, δ_1 can be interpreted as the fraction of initial wealth that the investor would like to forgo in exchange for a zero-capital-gains tax rate.

[Table 3 About Here]

Table 3 presents the results on the CEWL due to capital gains taxes for cases with or without return predictability. Consistent with previous intuition, we find that in the presence of return predictability, the utility loss due to capital gains taxes is substantially higher than that in the absence of return predictability. For example, in the FR case, the CEWL can be as high as 10.2% when return predictability is present. In comparison, the CEWL is only 5.4% when return predictability is absent. Similar results are found in the FC case. These results suggest that capital gains taxes are more burdensome to investors when they can exploit the variations in expected returns through market-timing.

Table 3 also shows the comparative statics of the utility costs with respect to changes in some key model parameters. Perhaps the most interesting parameters for our paper are those related to return predictability. First, the utility costs increase with the loading parameter μ_1. The larger this loading parameter, the more the expected return varies with the predictor; as a result, the investor needs to respond to changes in the return predictor more frequently, and such trading will incur heavier tax payments. Second, the utility costs increase with the volatility of the shocks on the predictor, i.e. σ_z. This is because the variability of the predictor increases with its volatility, and the investor needs to respond to such variability at the costs of paying heavier tax bills. Third, the utility costs decrease with the correlation between the return shocks and the predictor shocks, i.e. ρ. The more negative the correlation, the stronger the tax-timing incentive contradicts with the market-timing incentive, and the investor needs to mitigate such conflict at higher utility costs.

4.5 The Value of Tax-Deferral

In the presence of capital gains tax, the investor has an option to defer the realization of capital gains and earn the time value of the tax payments.[25] However, as argued previously, the negative correlation between return and predictor shocks generates conflict between market- and tax-timing incentives: with positive shocks on the stock price, the investor

[25] See, for example, Dammon et al. (1989) and Chay et al. (2006).

is likely to reach capital gains status, and is motivated to hold the stock to defer taxes. Meanwhile, the expected return of the stock is more likely to decrease, accordingly, the investor is also motivated to sell stock to reduce exposure. In this section, we examine how this conflict affects the value of the tax-deferral option.

To evaluate this tax-deferral option, we calculate the maximal utility level generated by a family of suboptimal trading strategies that never defer the realizations of capital gains. Let the value function associated with the "best" suboptimal trading strategy be $J^{ND}(x, y, k, z, t)$.[26] We then define the value of tax-deferral as the quantity δ_2, which solves

$$J^{ND}(1 + \delta_2, 0, 0, z_0, 0) = J(1, 0, 0, z_0, 0) \tag{16}$$

Put differently, δ_2 is the fraction of initial wealth that the investor gains by optimally exercising her tax-deferral option.

[Table 4 About Here]

Using the baseline calibration, we find that the tax-deferral option is more valuable when the stock's expected returns are predicted by the dividend yield. Table 4 shows the results. In the FR case, we obtain $\delta_2 = 0.021$ when μ_1 equals 0 (i.e., there is no return predictability) and $\delta_2 = 0.043$ when μ_1 equals our empirical estimate of 4.397. This increased value of the tax-deferral option can be attributed to two sources. First, the hedging demand induced by return predictability makes the investor allocate more wealth in the stock, and therefore the tax-deferral option enables the investor to save more time value of tax payments. Second, the strategy that never defers realizations of capital gains would incur excessive tax bills, especially when the investor is strongly driven by her market-timing incentive.

In the FC case, the values of tax-deferral are higher than those in the FR case across all scenarios (with or without return predictability). This is consistent with Dai et al. (2015),

[26] The characterizations of such value functions are provided in Appendix C.

who argue that there is an additional source of the value of tax-deferral in the FC case: it can make some of the future losses effectively rebatable. As the return predictor is negatively correlated with stock price, future losses are likely to be accompanied by good investment opportunities. Under such circumstances, making future losses rebatable can be valuable to the investor.

Table 4 also shows the values of tax-deferral for some perturbed parameter values. Generally, we find that these patterns are consistent with those displayed in Table 3. For instance, the value of tax deferral increases with the loading parameter μ_1 and with the volatility parameter σ_z, and decreases with the shock correlation ρ. This is due to the fact that the amount of tax payments incurred by the suboptimal strategy, which never defers capital gains realizations, increases with μ_1 and σ_z, and decreases with ρ. Tax deferral option would allow the investor to save more time value of taxes in these cases and hence becomes more valuable.

Importantly, these comparative static results suggest that for many plausible parameter values, the value of tax deferral is substantially higher in the presence of return predictability. This increased value of tax-deferral, together with the heavier utility cost due to capital gains taxes shown previously, may shed new light on the famous "asset location puzzle," i.e., the phenomenon that a large portion of investors hold both bonds and stocks in their taxable and tax-deferred accounts. The literature argues that such asset location decisions are tax inefficient, as the stream of interest generated by bonds is usually taxed at a higher rate than the capital gains generated by stock investments (e.g. Dammon et al. (2004)). However, our results suggest that if the benefits of conducting market-timing in a tax-deferred account exceed the costs of holding bonds in a taxable account, it may be optimal for the investor to locate stocks in the tax-deferred account and bonds in the taxable account. Therefore, the incentives behind "asset location puzzle" deserve further investigation. Nonetheless, a formal analysis of a portfolio choice problem with taxable and tax-deferred accounts in the presence of return predictability is beyond the scope of the current paper, and we leave it for future studies.

4.6 Cost of Separating Market- and Tax-Timing

One of the central messages of our paper is that if stock returns are predictable and if realized capital gains are taxed, then market- and tax-timing strategy should be jointly determined. This naturally raises a question: how costly is it for the investor to separately respond to market- and tax-timing incentives in her trading strategy? In this section, we calculate the utility cost of adopting an heuristic trading strategy that times market returns and capital gains taxes in a separate manner.

For this purpose, we first construct such a trading strategy. In Section 4.2, we show that when capital gains taxes are absent but return predictability is present, the investor's optimal portfolio consists of three distinct components: (C1) the classic Merton's optimal portfolio allocation with constant expected return μ_0, i.e., $\theta_0 = \frac{\mu_0 - r}{\gamma\sigma_S^2}$; (C2) a state-dependent component that responds to the changes in the value of return predictor, i.e., $\theta_1(z,t) = \frac{\mu_1 z}{\gamma\sigma_S^2}$; and (C3) the hedging demand induced by return predictability. Motivated by this observation, when both return predictability and capital gains taxes are incorporated in the model, we construct a trading strategy Π, which tries to separate the market-timing incentive from the tax-timing incentive as follows. First, the investor constructs a tax-timing strategy by assuming the stock's expected return is a constant μ_0. Second, at each time $t \in [0,T]$, conditional on the realized value of return predictor z_t, she shifts this tax-timing strategy (as characterized by the buying and selling boundaries) vertically by a certain amount as her response to the market-timing incentive. This amount of shifting equals the sum of the state-dependent component (C2) and the hedging demand component (C3) she would obtain in the model without capital gains taxes. By construction, the market- and tax-timing incentives are *separately* implemented in policy Π.[27]

[Figure 6 About Here]

Figure 6 illustrates the construction of policy Π (in the left subfigure) and the difference between policy Π and the optimal policy (in the right subfigure) for $t = 5$ and $z_5 = \bar{z} = 0.026$, assuming capital losses are fully rebated.[28] As can be observed, the policy Π clearly deviates from the optimal policy. In particular, the NTR of the optimal policy is significantly wider because the investor needs to control the amount of tax bills so that the benefit provided by market-timing is not overwhelmed by excessive tax payments.

Given the discrepancy between policy Π and the optimal policy, it is costly for the investor to adopt policy Π. We calculate the utility costs due to separating tax- and market-timing incentives through the CEWL resulting from adopting policy Π. This CEWL equals the quantity δ_3, which solves the following equation

$$J(1 - \delta_3, 0, 0, z_0, 0) = J^{\Pi}(1, 0, 0, z_0, 0) \tag{17}$$

where J^{Π} is the value function associated with policy Π, and J is the value function associated with the optimal trading policy. We use Monte Carlo simulation to compute the right-hand side of Equation (17).

Using our baseline calibration, we find that the CEWL resulting from adopting policy Π equals 3.0% in the FR case and 2.4% in the FC case. These utility costs are due to the fact that although policy Π separately responds to the investor's market and tax-timing incentives, it does not mitigate the conflict between market- and tax-timing as previously

[27] We have also considered the following policy Π' that directly shifts a set of myopic tax-timing policies by the hedging demand. First, at each time t, conditional on the realization of z_t, the investor constructs a myopic tax-timing policy by assuming that all future expected returns are constantly equal to $\mu_0 + \mu_1 z_t$. Second, the investor shifts the trading boundaries she obtained in the first step by the amount of hedging demand she would obtain in the model with return predictability but without capital gains taxes. We find that the utility costs of adopting policy Π' are quantitatively similar to that of adopting policy Π.

[28] This result for the FC case is qualitatively similar.

argued. By actively responding to changes in predicted market returns, the investor in fact pays excessive taxes. This result suggests that it is economically important to consider market- and tax-timing *jointly*.

4.7 Empirical Simulations

In this section, we use historical U.S. stock market data to perform empirical simulations for various popular trading strategies. These simulations are not intended to provide exact performance measures for these strategies, as doing so requires knowing the capital gains taxes codes enforced each year, the tax brackets to which the investor belongs, the returns obtained from holding treasury bills, and the stock trading costs. These factors are either changing over time or are difficult to measure accurately. Instead, the primary purpose of our empirical simulations is to illustrate that the optimal trading strategy implied by our model can have advantages when tested using real stock market data.

Specifically, we run a horserace over the period spanning from January 1950 to December 2016 between four investment strategies: (1) the "buy-and-hold" strategy (BH); (2) the Merton's constant ratio strategy (Merton), assuming a constant expected return of μ_0 ; (3) the market-timing strategy (MT) that exploits the variations in dividend yield; and (4) the optimal trading strategy implied by our model (OPT).[29] To implement the MT and OPT strategies, we compute respective trading strategies spanning over 67 years (1950-2016). As the optimal policy in the FC case is much more time-consuming to compute, we focus on the FR case in the empirical simulations. Moreover, we make the following assumptions to simplify the simulations: (1) the interest rate r is constant (0.031); (2) the tax code does not change over time and the tax rate τ is constant (25%); (3) all dividend distributions are

[29] When constructing these strategies, we constantly assume the investor has a relative risk-aversion coefficient of five. When examining Merton's constant ratio strategy and the market-timing strategy, we assume the investor rebalances her portfolio monthly so that the weight of the stock in her portfolio reaches the target level at the beginning of each new month.

re-invested.

[Figure 7 About Here.]

We calculate the growth of one dollar over the entire testing period. The top panel of Figure 7 shows the evolution of the investor's risk exposure to the stock market (as its weight in the portfolio) and her portfolio value net of capital gains taxes.[30] By construction, the BH and Merton strategies allocate constant fractions of wealth in the stock market. In contrast, the MT and OPT strategies adjust the stock exposures actively in response to changes in the dividend yield. As can be observed, the equity allocations of the MT and OPT strategies are highly correlated. The bottom panel of Figure 7 demonstrates how market-timing can generate extraordinary returns for the investor. In fact, the BH and Merton strategies underperform significantly, and their wealth paths are dominated by those generated by the MT strategy and the OPT strategy.

[Figure 8 About Here.]

The bottom panel of Figure 7 clearly shows that the OPT strategy outperforms the MT strategy. The MT strategy does not respond to tax-timing incentives; therefore, a significant portion of the capital gains are consumed by tax bills *too early*. To confirm this intuition, we depict in Figure 8 the cumulative tax bills generated by the MT and OPT strategies. As can be observed, the OPT strategy significantly defers capital gains taxes.

5 Conclusion

In this paper, we propose a continuous-time model to examine the joint timing of market returns and capital gains taxes. We find that the optimal trading policy exhibits several

[30] Note that we compare the net wealth levels rather than gross wealth levels. The OPT strategy defers capital gains taxes and can result in a higher gross wealth level.

salient features. First, it contains tax-deferral regions as a response to the tax-timing incentive. Second, the locations of these tax-deferral regions vary with respect to the expected return of the stock as a response to the market-timing incentive. The tax-deferral regions are much wider than those obtained in the absence of return predictability, which reflects an incentive to mitigate the conflict between market- and tax-timing incentives caused by the negative correlation between return and predictor shocks. We also show that it is economically important to adopt the optimal joint-timing policy. An investor with an investment horizon of 10 years and a relative risk aversion coefficient of 5 effectively loses 2.4-3% of her initial wealth if she responds to market- and tax-timing incentives separately.

Our model suggests that with an empirically plausible amount of return predictability, capital gains taxes are more burdensome to equity investors than previously thought. This is also reflected in the increased value of tax-deferral. Such an observation can have implications for the famous "asset location puzzle," i.e., the investors' tendency to hold mixed equity-bond portfolios in their tax-deferred investment accounts. Our analysis implies that holding stocks in the tax-deferred accounts may benefit the investors by allowing them to adjust equity portfolios in a more tax-efficient way as responses to varying equity returns.

References

Amromin, G., 2003, Household portfolio choices in taxable and tax-deferred accounts: Another puzzle? *Review of Finance*, 7(3), 547–582.

Barberis, N., 2000, Investing for the long run when returns are predictable, *Journal of Finance*, 55, 225–264.

Ben Tahar, I., H. M. Soner, and N. Touzi, 2007, The dynamic programming equation for the problem of optimal investment under capital gains taxes, *SIAM Journal on Control and Optimization*, 46(5), 1779–1801.

Ben Tahar, I., H. M. Soner, and N. Touzi, 2010, Merton problem with taxes: Characterization, computation, and approximation, *SIAM Journal on Financial Mathematics*, 1(1), 366–395.

Branger, N., L. S. Larsen, and C. Munkc, 2013, Robust portfolio choice with ambiguity and learning about return predictability, *Journal of Banking and Finance*, 37, 1397-1411.

Cai, J., X. Chen, and M. Dai, 2017, Portfolio selection with capital gains tax, recursive utility, and regime switching, *Management Science*, Forthcoming.

Chay, J. B., D. Choi, and J. Pontiff, 2006, Market valuation of tax-timing options: Evidence from capital gains distributions, *Journal of Finance*, 61, 837–865.

Campbell, J. Y., 1991, A variance decomposition for stock returns, *The Economic Journal*, 101(405), 157–179.

Campbell, J. Y., and R. J. Shiller, 1988a, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies*, 1, 195-228.

Campbell, J. Y., and R. J. Shiller, 1988b, Stock prices, earnings, and expected dividends, *Journal of Finance*, 43, 661-676.

Campbell, J. Y., and L. M. Vicera, 1999, Consumption and portfolio decisions when expected returns are time varying, *The Quarterly Journal of Economics*, 114, 433–495.

Campbell, J. Y., and L. M. Viceira, 2002, Strategic Asset Allocation Portfolio Choice for

Long-Term Investors (Oxford University Press, Oxford, UK).

Cochrane J. H., 2008, The dog that did not bark: A defense of return predictability, *Review of Financial Studies*, 21(4), 1533–1575.

Constantinides, G. M., 1983, Capital-market equilibrium with personal tax, *Econometrica*, 51, 611–636.

Constantinides, G. M., 1984, Optimal stock trading with personal taxes: Implications for prices and the abnormal January returns, *Journal of Financial Economics*, 13(1), 6589.

Dai, M., H. Liu, C. Yang, and Y. Zhong, 2015, Optimal tax timing with asymmetric long-term/short-term capital gains tax, *Review of Financial Studies*, 28, 2687–2721.

Dammon, R. M., K. B. Dunn, and C. S. Spatt, 1989, A reexamination of the value of tax options, *Review of Financial Studies*, 2, 341-372.

Dammon, R. M., and C. S. Spatt, 1996, The optimal trading and pricing of securities with asymmetric capital gains taxes and transaction costs, *Review of Financial Studies*, 9, 921–952.

Dammon, R. M., C. S. Spatt, and H. H. Zhang, 2001, Optimal consumption and investment with capital gains taxes, *Review of Financial Studies*, 14, 583–616.

Dammon, R. M., C. S. Spatt, and H. H. Zhang, 2004, Optimal asset location and allocation with taxable and tax-deferred investing, *The Journal of Finance*, 59(3), 999–1037.

DeMiguel, V., and R. Uppal, 2005, Portfolio investment with the exact tax basis via nonlinear programming, *Management Science*, 51(2), 277–290.

Diether, K. B., Kuan-Hui Lee, and I. M. Werner, 2009, Short-sale strategies and return predictability, *Review of Financial Studies*, 22(2), 575–607.

Ehling, P., M. F. Gallmeyer, S. Srivastava, S. Tompaidis, and C. Yang, 2013, Portfolio choice with capital gain taxation and the limited use of losses, EFA 2008 Athens Meetings Paper.

Fama, E., and K. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3-27.

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

Fischer, M., and M. Gallmeyer, 2017, Taxable and tax-deferred investing with the limited use of losses, *Review of Finance* 21(5), 1847-1873.

Garlappi, L., and J. Huang, 2006, Are stocks desirable in tax-deferred accounts? *Journal of Public Economics*, 90, 2257-2283.

Garleanu, N. and L. H. Pedersen, 2013, Dynamic trading with predictable returns and transaction costs, *Journal of Finance*, 68, 2309–2340.

Huang, L., and H. Liu, 2007, Rational inattention and portfolio selection, *Journal of Finance*, 62(4), 1999–2040.

Kandel, S., and R. Stambaugh, 1996, On the predictability of stock returns: An asset allocation perspective, *Journal of Finance*, 51, 385-424.

Keim, D. B., and R. F. Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics*, 17(2), 357–390.

Lettau, M., and S. Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance*, 56, 815-849.

Lettau, M. and S. V. Nieuwerburgh, 2008, Reconciling the return predictability evidence, *Review of Financial Studies*, 21(4), 1607–1652.

Lynch, A. W., 2001, Portfolio choice and equity characteristics: Characterizing the hedging demands induced by return predictability, *Journal of Financial Economics*, 62, 67-130.

Lynch, A. W., and S. Tan, 2010, Multiple risky assets, transaction costs and return predictability: Allocation rules and implications for U.S. investors, *Journal of Financial and Quantitative Analysis*, 45(4), 1015–1053.

Marekwica, M., 2012, Optimal tax-timing and asset allocation when tax rebates on capital losses are limited, *Journal of Banking and Finance*, 36(7), 2048–2063.

Merton, R. C., 1969, Lifetime portfolio selection under uncertainty: The continuous-time case, *The Review of Economics and Statistics*, 3, 247–257.

Merton, R. C., 1971, Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory*, 3, 373–413.

Moallemi, C. C., and M. Saglam, 2017, Dynamic portfolio choice with linear rebalancing rules, *Journal of Financial and Quantitative Analysis*, 52(3), 1247–1278.

Poterba, J. M., 1987, How burdensome are capital gains taxes?: Evidence from the United States, *Journal of Public Economics*, 33(2), 157–172.

Shoven, J. B., and C. Sialm, 2003, Asset location in tax-deferred and conventional savings accounts, *Journal of Public Economics* 88, 23–38.

Stambaugh, R. F., 1999, Predictive regressions, *Journal of Financial Economics*, 54(3), 375–421.

Tsai, H., and Y. Wu, 2015, Optimal portfolio choice with asset return predictability and nontradable labor income, *Review of Quantitative Finance and Accounting*, 45 (1), 215–249.

Wachter, J. A., and M. Warusawitharana, 2009, Predictable returns and asset allocation: Should a skeptical investor time the market? *Journal of Econometrics*, 148, 162–178.

Xia, Y., 2001, Learning about predictability: The effects of parameter uncertainty on dynamic asset allocation, *Journal of Finance*, 56, 205–246.

Appendix

The content of this appendix is as follows. In Appendix A, we present a set of theoretical analysis of our model. In Appendix B, we briefly describe our parameter estimation method. In Appendix C, we characterize the value functions associated with the suboptimal policy under which the investor never defers the realizations of capital gains.

A Theoretical Analysis

A.1.1 A Heuristic Derivation of the HJB Equation

First, the terminal condition is easy to verify. To derive equation (10), we begin by approximating the processes by $dI_t = i_t dt$, $dM_t = m_t dt$ with $0 \le i_t, m_t \le N$. As such, the dynamics of state variables read

$$
\begin{aligned}
dx_t &= rx_t dt - i_t dt + f(0, y_{t-}, k_{t-}; l) m_t dt, \\
dy_t &= (\mu_0 + \mu_1(z_t - \bar{z})) y_t dt + \sigma_S y_t dB_t^S + i_t dt - y_{t-} m_t dt, \\
dk_t &= i_t dt - k_{t-} m_t dt + l\,(k_{t-} - y_{t-})^+ m_t dt
\end{aligned}
$$

We denote the approximated value function by $J^N(x, y, k, z, t)$. Then, the Bellman principle of optimality suggests that

$$\sup_{m,i} \partial_t J^N + \mathcal{L}^{m,i} J^N = 0$$

where

$$
\begin{aligned}
\mathcal{L}^{m,i} J^N = \; & rx\partial_x J^N + (\mu_0 + \mu_1(z - \bar{z}))y\partial_y J^N + \frac{1}{2}\sigma_S^2 y^2 \partial_{yy} J^N + g_1(\bar{z} - z)\partial_z J^N \\
& + \frac{1}{2}\sigma_z^2 \partial_{zz} J^N + \rho \sigma_S \sigma_z y \partial_{yz} J^N + (-\partial_x J^N + \partial_y J^N + \partial_k J^N) i \\
& + (f(0, y, k; l)\partial_x J^N - y\partial_y J^N + (l\,(k - y)^+ - k)\partial_k J^N) m
\end{aligned}
$$

Taking the superimum over m and i, one obtains

$$\begin{aligned}\partial_t J^N &+ rx\partial_x J^N + (\mu_0 + \mu_1(z-\bar{z}))y\partial_y J^N + \frac{1}{2}\sigma_S^2 y^2 \partial_{yy} J^N + g_1(\bar{z}-z)\partial_z J^N \\ &+ \frac{1}{2}\sigma_z^2 \partial_{zz} J^N + \rho\sigma_S\sigma_z y\partial_{yz} J^N + (-\partial_x J^N + \partial_y J^N + \partial_k J^N)^+ N \\ &+ (f(0,y,k;l)\partial_x J^N - y\partial_y J^N + (l\,(k-y)^+ - k)\partial_k J^N)^+ N = 0\end{aligned}$$

Therefore,

$$\begin{aligned}\partial_t J^N &+ rx\partial_x J^N + (\mu_0 + \mu_1(z-\bar{z}))y\partial_y J^N + \frac{1}{2}\sigma_S^2 y^2 \partial_{yy} J^N + g_1(\bar{z}-z)\partial_z J^N \\ &+ \frac{1}{2}\sigma_z^2 \partial_{zz} J^N + \rho\sigma_S\sigma_z y\partial_{yz} J^N \le 0,\end{aligned}$$

$$-\partial_x J^N + \partial_y J^N + \partial_k J^N \le 0,$$

and

$$f(0,y,k;l)\partial_x J^N - y\partial_y J^N + (l\,(k-y)^+ - k)\partial_k J^N \le 0.$$

Otherwise, one can derive contradiction by sending $N \to \infty$. Moreover, if

$$-\partial_x J^N + \partial_y J^N + \partial_k J^N < 0,$$

and

$$f(0,y,k;l)\partial_x J^N - y\partial_y J^N + (l\,(k-y)^+ - k)\partial_k J^N < 0,$$

one must have

$$\begin{aligned}\partial_t J^N &+ rx\partial_x J^N + (\mu_0 + \mu_1(z-\bar{z}))y\partial_y J^N + \frac{1}{2}\sigma_S^2 y^2 \partial_{yy} J^N + g_1(\bar{z}-z)\partial_z J^N \\ &+ \frac{1}{2}\sigma_z^2 \partial_{zz} J^N + \rho\sigma_S\sigma_z y\partial_{yz} J^N = 0.\end{aligned}$$

As $J^N \to J$ when $N \to \infty$, the results follow.

A.1.2 Dimensional Reduction and Verification

In this Appendix, we illustrate the dimensional reduction of the HJB equation and the verification argument. Since different changes of variables are used in the FR and FC cases, we analyze these two cases separately.

The FR Case. Due to the homogeneity of the utility function, there exists a function $\varphi(b, \pi, z, t)$ such that

$$J(x, y, k, z, t) = \frac{1}{1-\gamma} f(x, y, k; 0)^{1-\gamma} e^{(1-\gamma)\varphi(b,\pi,z,t)},$$

where $b = \frac{k}{y}$, $\pi = \frac{(1-\tau)y}{x+(1-\tau)y+\tau k}$. Then, it can be shown that $\varphi(b, \pi, z, t)$ solves the following equation:

$$\max\{\mathcal{L}_1\varphi + \varphi_t, \;\; \mathcal{B}_1\varphi, \;\; \mathcal{S}_1\varphi\} = 0, \tag{A-1}$$

on the domain $\Sigma = \{(b, \pi, z, t) \in R^+ \times R^+ \times R \times [0, T]\}$, with terminal condition

$$\varphi(b, \pi, z, T) = 0 \tag{A-2}$$

where

$$\mathcal{B}_1\varphi = \pi\varphi_\pi + (1 - b)\varphi_b$$

$$\mathcal{S}_1\varphi = -\varphi_\pi$$

and

$$
\begin{aligned}
\mathcal{L}_1\varphi &= \frac{1}{2}\sigma_S^2\pi^2(1-\pi)^2[\varphi_{\pi\pi}+(1-\gamma)\varphi_\pi^2]+\frac{1}{2}\sigma_S^2 b^2[\varphi_{bb}+(1-\gamma)\varphi_b^2]+\frac{1}{2}\sigma_z^2[\varphi_{zz}+(1-\gamma)\varphi_z^2] \\
&\quad -\sigma_S^2\pi(1-\pi)b[\varphi_{\pi b}+(1-\gamma)\varphi_\pi\varphi_b]+\rho\sigma_S\sigma_z\pi(1-\pi)[\varphi_{\pi z}+(1-\gamma)\varphi_\pi\varphi_z] \\
&\quad -\rho\sigma_S\sigma_z b[\varphi_{bz}+(1-\gamma)\varphi_b\varphi_z]+[(\mu_0+\mu_1(z-\bar{z})-r)\pi(1-\pi)+\frac{r\tau}{1-\tau}\pi^2 b-\gamma\sigma_S^2\pi^2(1-\pi)]\varphi_\pi \\
&\quad +[\sigma_S^2(b-(1-\gamma)\pi b)-(\mu_0+\mu_1(z-\bar{z}))b]\varphi_b+[g_1(\bar{z}-z)+(1-\gamma)\rho\sigma_S\sigma_z\pi]\varphi_z \\
&\quad +r(1-\pi-\frac{\tau}{1-\tau}\pi b)+(\mu_0+\mu_1(z-\bar{z}))\pi-\frac{1}{2}\sigma_S^2\gamma\pi^2.
\end{aligned}
$$

Equation (A-1) splits the solution domain Σ into three regions: a buy region

$$BR \equiv \{(b,\pi,z,t) : \mathcal{B}_1\varphi = 0\}; \tag{A-3}$$

a sell region

$$SR \equiv \{(b,\pi,z,t) : \mathcal{S}_1\varphi = 0\}; \tag{A-4}$$

and a non-trading region

$$NTR \equiv \{(b,\pi,z,t) : \mathcal{B}_1\varphi < 0,\ \mathcal{S}_1\varphi < 0\}. \tag{A-5}$$

Given these regions, the following verification theorem characterizes the value function and optimal trading strategy:

Proposition 1: *(VT for the FR case) Let $\varphi(b,\pi,z,t)$ be a solution to equation (A-1) with terminal condition (A-2) satisfying certain regularity conditions, with the trading regions defined as above. Then, the optimal trading policy (L_t^*, M_t^*) is given by*

$$L_t^* = \int_0^t 1_{\{(b,\pi,z,s)\in\partial BR\bigcap\partial NTR\}} dL_s^*,$$

and

$$M_t^* = \int_0^t 1_{\{(b,\pi,z,s)\in\partial SR\bigcap\partial NTR\}} dM_s^*.$$

In addition, we define

$$V(x,y,k,z,t) = \frac{1}{1-\gamma} f(x,y,k;0)^{1-\gamma} e^{(1-\gamma)\varphi\left(\frac{k}{y},\frac{(1-\tau)y}{f(x,y,k;0)},z,t\right)},$$

then $V(x,y,k,z,t)$ coincides with the value function $J(x,y,k,z,t)$.

The proof of the above verification theorem is similar to that in Dai et al. (2015), we omit it for brevity.

The FC Case. In the FC case, we make the following transformation:

$$J(x,y,k,z,t) = \frac{(x+y)^{1-\gamma}}{1-\gamma} e^{(1-\gamma)\varphi(b,\pi,z,t)},$$

where $\pi = \frac{y}{x+y}$ and $b = \frac{k}{y}$.[31] Then, φ satisfies

$$\max\{\mathcal{L}_2\varphi + \varphi_t,\ \ \mathcal{B}_2\varphi,\ \ \mathcal{S}_2\varphi\} = 0, \tag{A-6}$$

on the domain $\{(b,\pi,z,t) \in R^+ \times R^+ \times R \times [0,T]\}$, with terminal condition

$$\varphi(b,\pi,z,T) = \log[1 - \tau\pi(1-b)^+],$$

where

$$\mathcal{B}_2\varphi = \pi\varphi_\pi + (1-b)\varphi_b,$$

$$\mathcal{S}_2\varphi = -\tau\pi(1-b)^+ + [-1+\tau\pi(1-b)^+]\pi\varphi_\pi + (b-1)^+\varphi_b,$$

[31] Here, we slightly abuse the same notation π to denote another change of variable.

and

$$
\begin{aligned}
\mathcal{L}_2\varphi &= \frac{1}{2}\sigma_S^2\pi^2(1-\pi)^2[\varphi_{\pi\pi}+(1-\gamma)\varphi_\pi^2]+\frac{1}{2}\sigma_S^2b^2[\varphi_{bb}+(1-\gamma)\varphi_b^2]+\frac{1}{2}\sigma_z^2[\varphi_{zz}+(1-\gamma)\varphi_z^2] \\
&\quad -\sigma_S^2\pi(1-\pi)b[\varphi_{\pi b}+(1-\gamma)\varphi_\pi\varphi_b]+\rho\sigma_S\sigma_z\pi(1-\pi)[\varphi_{\pi z}+(1-\gamma)\varphi_\pi\varphi_z] \\
&\quad -\rho\sigma_S\sigma_z b[\varphi_{bz}+(1-\gamma)\varphi_b\varphi_z]+[(\mu_0+\mu_1(z-\bar{z})-\gamma\sigma_S^2\pi-r)\pi(1-\pi)]\varphi_\pi \\
&\quad +[-(\mu_0+\mu_1(z-\bar{z}))+\sigma_S^2(1-(1-\gamma)\pi)]b\varphi_b+[g_1(\bar{z}-z)+(1-\gamma)\rho\sigma_S\sigma_z\pi]\varphi_z \\
&\quad -\frac{1}{2}\gamma\sigma_S^2\pi^2+(\mu_0+\mu_1(z-\bar{z})-r)\pi+r.
\end{aligned}
$$

As

$$\pi\varphi_\pi+(1-b)\varphi_b=0, \quad b>1,$$

continuous trading is optimal for $b>1$. Then, in this region, we can rewrite φ as follows:

$$\varphi(b,\pi,z,t)=\psi(\zeta,z,t),$$

where $\zeta=\frac{1}{1+\pi(b-1)}$. Then, ψ satisfies

$$\max_\pi\{\mathcal{L}_3\psi+\psi_t\}=0, \tag{A-7}$$

where

$$
\begin{aligned}
\mathcal{L}_3\psi &= \frac{1}{2}\sigma_S^2\pi^2\zeta^2[\psi_{\zeta\zeta}+(1-\gamma)\psi_\zeta^2]+\frac{1}{2}\sigma_z^2[\psi_{zz}+(1-\gamma)\psi_z^2]+\rho\sigma_S\sigma_z\pi\zeta[\psi_{\zeta z}+(1-\gamma)\psi_\zeta\psi_z] \\
&\quad +[(1-\gamma)\sigma_S^2\pi^2+(\mu_0+\mu_1(z-\bar{z})-r)\pi+r-r(1-\pi)\zeta]\zeta\psi_\zeta \\
&\quad +[g_1(\bar{z}-z)+(1-\gamma)\rho\sigma_S\sigma_z\pi]\psi_z-\frac{1}{2}\gamma\sigma_S^2\pi^2+(\mu_0+\mu_1(z-\bar{z})-r)\pi+r
\end{aligned} \tag{A-8}
$$

and the optimal strategy is given by

$$\pi^* = \frac{(\mu_0 + \mu_1(z - \bar{z}) - r)(1 + \zeta\psi_\zeta) + r\zeta^2\psi_\zeta + \rho\sigma_S\sigma_z[\zeta(\psi_{\zeta z} + (1-\gamma)\psi_\zeta\psi_z) + (1-\gamma)\psi_z]}{-\sigma_S^2[-\gamma + 2(1-\gamma)\zeta\psi_\zeta + \zeta^2(\psi_{\zeta\zeta} + (1-\gamma)\psi_\zeta^2)]}.$$

The structure of the optimal trading strategy and the verification theorem for the FC case are similar to those for the FR case, so we do not restate them for the sake of brevity.

B Parameter Estimation Method

We recall that the dynamics for the stock price S_t and the dividend yield z_t are as follows

$$\frac{dS_t}{S_t} = (\mu_0 + \mu_1(z_t - \bar{z}))dt + \sigma_S dB_t^S, \tag{A-9}$$

$$dz_t = g_1(\bar{z} - z_t)dt + \sigma_z dB_t^z, \tag{A-10}$$

where $\mathbb{E}[dB_t^z dB_t^S] = \rho dt$. We therefore estimate a VAR regression for the annualized stock return r_i and the dividend yield z_i:

$$\begin{aligned} r_i &= \mu_0 - \mu_1\bar{z} + \mu_1 z_{i-1} + \varepsilon_{1i} := a_1 + B_1 z_{i-1} + \varepsilon_{1i} \\ z_i &= \bar{z}(1 - e^{-g_1/12}) + e^{-g_1/12} z_{i-1} + \varepsilon_{2i} := a_2 + B_2 z_{i-1} + \varepsilon_{2i} \end{aligned}$$

for $i = 2, 3, \ldots, n$. We can rewrite the preceding equations in matrix form:

$$\begin{pmatrix} r_2 & z_2 \\ r_3 & z_3 \\ \vdots & \vdots \\ r_n & z_n \end{pmatrix} = \begin{pmatrix} 1 & z_1 \\ 1 & z_2 \\ \vdots & \vdots \\ 1 & z_{n-1} \end{pmatrix} \begin{pmatrix} a_1 & a_2 \\ B_1 & B_2 \end{pmatrix} + \begin{pmatrix} \varepsilon_{12} & \varepsilon_{22} \\ \varepsilon_{13} & \varepsilon_{23} \\ \vdots & \vdots \\ \varepsilon_{1n} & \varepsilon_{2n} \end{pmatrix}$$

or

$$Z = XC + E$$

Then, the estimated coefficients are $\hat{C} = (X'X)^{-1}X'Z$ and the residual is $\hat{E} = Z - X\hat{C}$. We denote by σ_{ε_1} and σ_{ε_2} the standard deviation for the two column vectors of $\hat{E}$. Then, the volatility of the dividend yield is $\sigma_z = \sigma_{\varepsilon_2}\sqrt{2g_1/(1-e^{-g_1/6})}$, and the correlation coefficient between the two column vectors of $\hat{E}$ is ρ.

C Characterizations of Value Functions When Capital Gains Realizations are Never Deferred

The FR Case. In the FR case, if the investor never defers realizations of capital gains, then, similar to Dai et al. (2015), it can be shown that she effectively solves the following problem:

$$\max_{y_t} E[u(W_T)], \tag{A-11}$$

subject to the budget constraint

$$dW_t = rx_t dt + (1-\tau)y_t(\mu_0 + \mu_1(z_t - \bar{z}))dt + (1-\tau)\sigma_S y_t dB_t^S \tag{A-12}$$

and the short-sale constraint $y_t \geq 0$. We solve this problem using a dynamic programming method.[32]

The FC Case. Suppose the investor does not defer realizing any gains. When a capital loss occurs, the investor sells the entire stock holdings, and buys back to make the fraction

[32] The details are available upon request.

of wealth in stock equal to π_*; that is,

$$(x, y, k) \to (x + y, 0, k - y) \to ((1 - \pi_*)(x + y), \pi_*(x + y), k - y + \pi_*(x + y)).$$

After rebalancing, the basis-to-price ratio b_* is still greater than one.[33] Accordingly, the function $\psi_1(\zeta, z, t)$ defined by the transformation

$$J^{ND}(x, y, k, z, t) = \frac{(x + y)^{1-\gamma}}{1 - \gamma} e^{(1-\gamma)\psi_1(\zeta,z,t)}, \quad \text{for } k \geq y,$$

where $\zeta = \frac{x+y}{x+k} \in [0, 1]$, satisfies the following equation

$$\max_{\nu} \left\{ \frac{\partial \psi_1}{\partial t} + \mathcal{L}_3 \psi_1 \right\} = 0 \text{ in } 0 \leq \zeta < 1, \text{ with}$$
$$-\frac{\partial \psi_1}{\partial \zeta} = \tau \text{ at } \zeta = 1,$$

where the operator $\mathcal{L}_3$ is given by (A-8) and the boundary condition is implied by equation (A-13). In addition, the optimal strategy satisfies the first-order condition of the above equation.

When there is a capital gain, the optimal strategy is

$$(x, y, k) \to (W, 0, 0) \to ((1 - \pi_*)W, \pi_*W, \pi_*W), \text{ where } W = x + y - \tau(y - k),$$

and the value function satisfies

$$J^{ND}(x, y, k, z, t) = J^{ND}((1 - \pi_*)W, \pi_*W, \pi_*W, z, t).$$

[33] Note that $b_* = \frac{k-y+\pi_*(x+y)}{\pi_*(x+y)} = \frac{\pi}{\pi_*}(b-1)+1 > 1$. If $\pi_* = \pi$, then $b_* = b$. This suggests that the investor does not adjust his stock holdings at all.

We rewrite the value function as

$$J^{ND}(x, y, k, z, t) = \frac{(x+y)^{1-\gamma}}{1-\gamma} e^{(1-\gamma)\varphi_1(b,\pi,z,t)}, \quad \text{for } k < y.$$

Then, we have

$$\frac{\varphi_1(1, \pi, z, t) - \varphi_1(b, \pi, z, t)}{1-b} = \frac{-\log[1 - \tau\pi(1-b)]}{1-b}$$

As $b \to 1$,

$$\frac{\partial \varphi_1(1, \pi, z, t)}{\partial b} = \tau\pi. \tag{A-13}$$

Table 1: Baseline Calibration

This table summarizes our baseline calibration. The model is calibrated to the value-weighted market returns in the U.S. from January 1950 to December 2016. We use dividend yield as the unique return predictor in our calibration. The risk-free rate is approximated by the average value of the after-tax yield of treasury bills with a constant maturity of one year.

Parameter	Symbol	Baseline value
Investment horizon (years)	T	10
Relative risk-aversion coefficient	γ	5
Tax-adjusted risk-free rate	r	0.031
Long-term average return of the stock	μ_0	0.116
Loading on the predictive variable	μ_1	4.397
Volatility of stock returns	σ_S	0.147
Average dividend yield	$\bar{z}$	0.026
Mean reverting speed of the predictor	g_1	0.141
Volatility of the predictor	σ_z	0.005
Correlation between return and predictor shocks	ρ	-0.895
Capital gains tax rate	τ	0.25

Table 2: Simulation Results

This table shows the expected discounted tax bill paid during the entire investment horizon as a fraction of the investor's initial wealth and the average duration between realizations of capital gains. These results are obtained from 10,000 paths of Monte Carlo simulation of the optimal investment policy. We report the results for both the case without return predictability (under the column "Without R. P.") and the case with return predictability (under the column "With R. P."). Parameter values: $T = 10$, $\gamma = 5$, $r = 0.031$, $\mu_0 = 0.116$, $\mu_1 = 4.397$, $\sigma_S = 0.147$, $g_1 = 0.141$, $\bar{z} = 0.026$, $\sigma_z = 0.005$, $\rho = -0.895$, $\tau = 0.25$, and $z_0 = \bar{z} = 0.026$.

	Panel A: The FR Case		Panel B: The FC Case	
	Without R. P.	With R. P.	Without R.P.	With R.P.
Time btw Capital Gains Realizations	0.497	0.338	1.045	0.510
Discounted Capital Gains Tax	0.343	0.718	0.310	0.615

Table 3: CEWL due to Capital Gains Taxes

This table shows the utility costs accrued to the investor due to capital gains tax liability. We report the results for both the case without return predictability (under the column "Without Ret. Pred.") and the case with return predictability (under the column "With Ret. Pred."). Baseline parameter values: $T = 10$, $\gamma = 5$, $r = 0.031$, $\mu_0 = 0.116$, $\mu_1 = 4.397$, $\sigma_S = 0.147$, $g_1 = 0.141$, $\bar{z} = 0.026$, $\sigma_z = 0.005$, $\rho = -0.895$, $\tau = 0.25$, and $z_0 = \bar{z} = 0.026$.

	Panel A: The FR Case		Panel B: The FC Case	
	Without Ret. Pred.	With Ret. Pred.	Without Ret. Pred.	With Ret. Pred.
Baseline case	0.054	0.102	0.070	0.146
$\mu_1 \times 0.9$	-	0.091	-	0.137
$\mu_1 \times 1.1$	-	0.116	-	0.154
$\sigma_z = 0.004$	-	0.084	-	0.126
$\sigma_z = 0.006$	-	0.105	-	0.160
$\rho = -0.85$	-	0.098	-	0.139
$\rho = -0.95$	-	0.104	-	0.156
$r = 0.02$	0.042	0.079	0.067	0.137
$r = 0.04$	0.059	0.118	0.071	0.149
$\gamma = 4$	0.067	0.124	0.087	0.170
$\gamma = 6$	0.046	0.086	0.059	0.129
$\tau = 0.15$	0.030	0.050	0.041	0.086
$\tau = 0.35$	0.083	0.161	0.103	0.209

Table 4: Value of Tax Deferral

This table shows the value of deferring capital gains realizations, i.e., the quantity δ_2 that solves Equation (16). We report the results for both the case without return predictability (under the column "Without Ret. Pred.") and the case with return predictability (under the column "With Ret. Pred."). Baseline parameter values: $T = 10$, $\gamma = 5$, $r = 0.031$, $\mu_0 = 0.116$, $\mu_1 = 4.397$, $\sigma_S = 0.147$, $g_1 = 0.141$, $\bar{z} = 0.026$, $\sigma_z = 0.005$, $\rho = -0.895$, $\tau = 0.25$, and $z_0 = \bar{z} = 0.026$.

	Panel A: The FR Case		Panel B: The FC Case	
	Without Ret. Pred.	With Ret. Pred.	Without Ret. Pred.	With Ret. Pred.
Baseline case	0.021	0.043	0.055	0.090
$\mu_1 \times 0.9$	-	0.036	-	0.088
$\mu_1 \times 1.1$	-	0.053	-	0.090
$\sigma_z = 0.004$	-	0.040	-	0.082
$\sigma_z = 0.006$	-	0.051	-	0.091
$\rho = -0.85$	-	0.032	-	0.082
$\rho = -0.95$	-	0.056	-	0.102
$r = 0.02$	0.014	0.045	0.062	0.107
$r = 0.04$	0.024	0.042	0.050	0.077
$\gamma = 4$	0.027	0.045	0.071	0.116
$\gamma = 6$	0.017	0.040	0.045	0.072
$\tau = 0.15$	0.011	0.031	0.034	0.053
$\tau = 0.35$	0.033	0.058	0.074	0.122

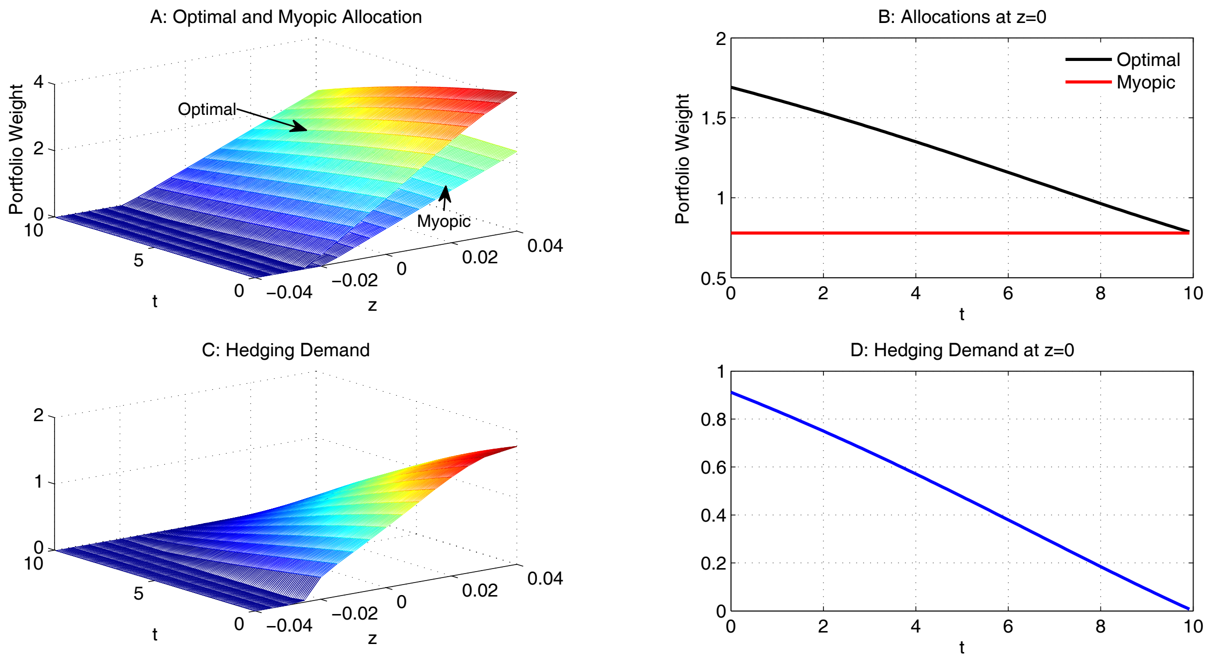

Figure 1: Optimal Allocation without Capital Gains Taxes

This figure shows the optimal allocation in the stock without capital gains taxes. Note that in the figure, z denotes $z - \bar{z}$. Parameter values: $T = 10$, $\gamma = 5$, $r = 0.031$, $\mu_0 = 0.116$, $\mu_1 = 4.397$, $\sigma_S = 0.147$, $g_1 = 0.141$, $\bar{z} = 0.026$, $\sigma_z = 0.005$, and $\rho = -0.895$. Capital gains taxes are not incorporated, i.e., $\tau = 0$.

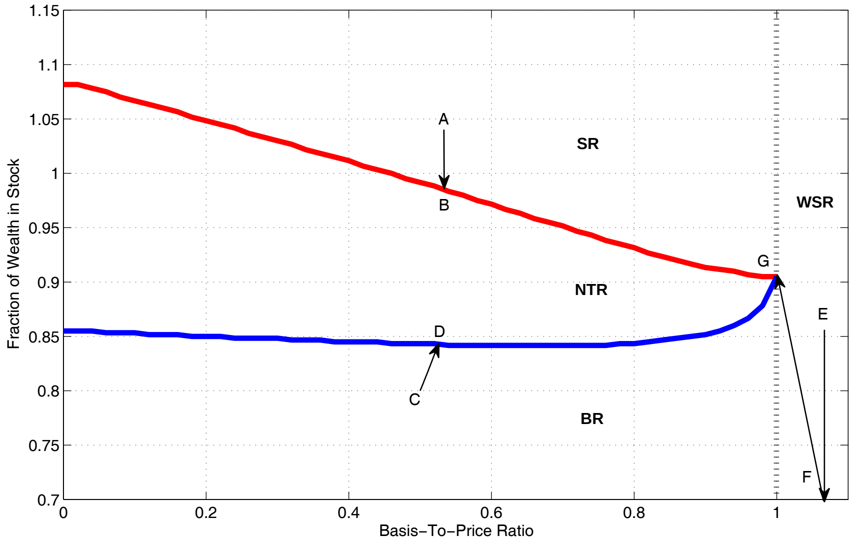

Figure 2: Optimal Policy without Return Predictability: The FR Case
This figure shows the optimal sell and buy boundaries at $t = 5$ in the presence of capital gains taxes, without return predictability. Parameter values: $T = 10$, $\gamma = 5$, $r = 0.031$, $\mu_0 = 0.116$, $\mu_1 = 0$, $\sigma_S = 0.147$, $g_1 = 0.141$, $\bar{z} = 0.026$, $\sigma_z = 0.005$, $\rho = -0.895$, and $\tau = 0.25$.

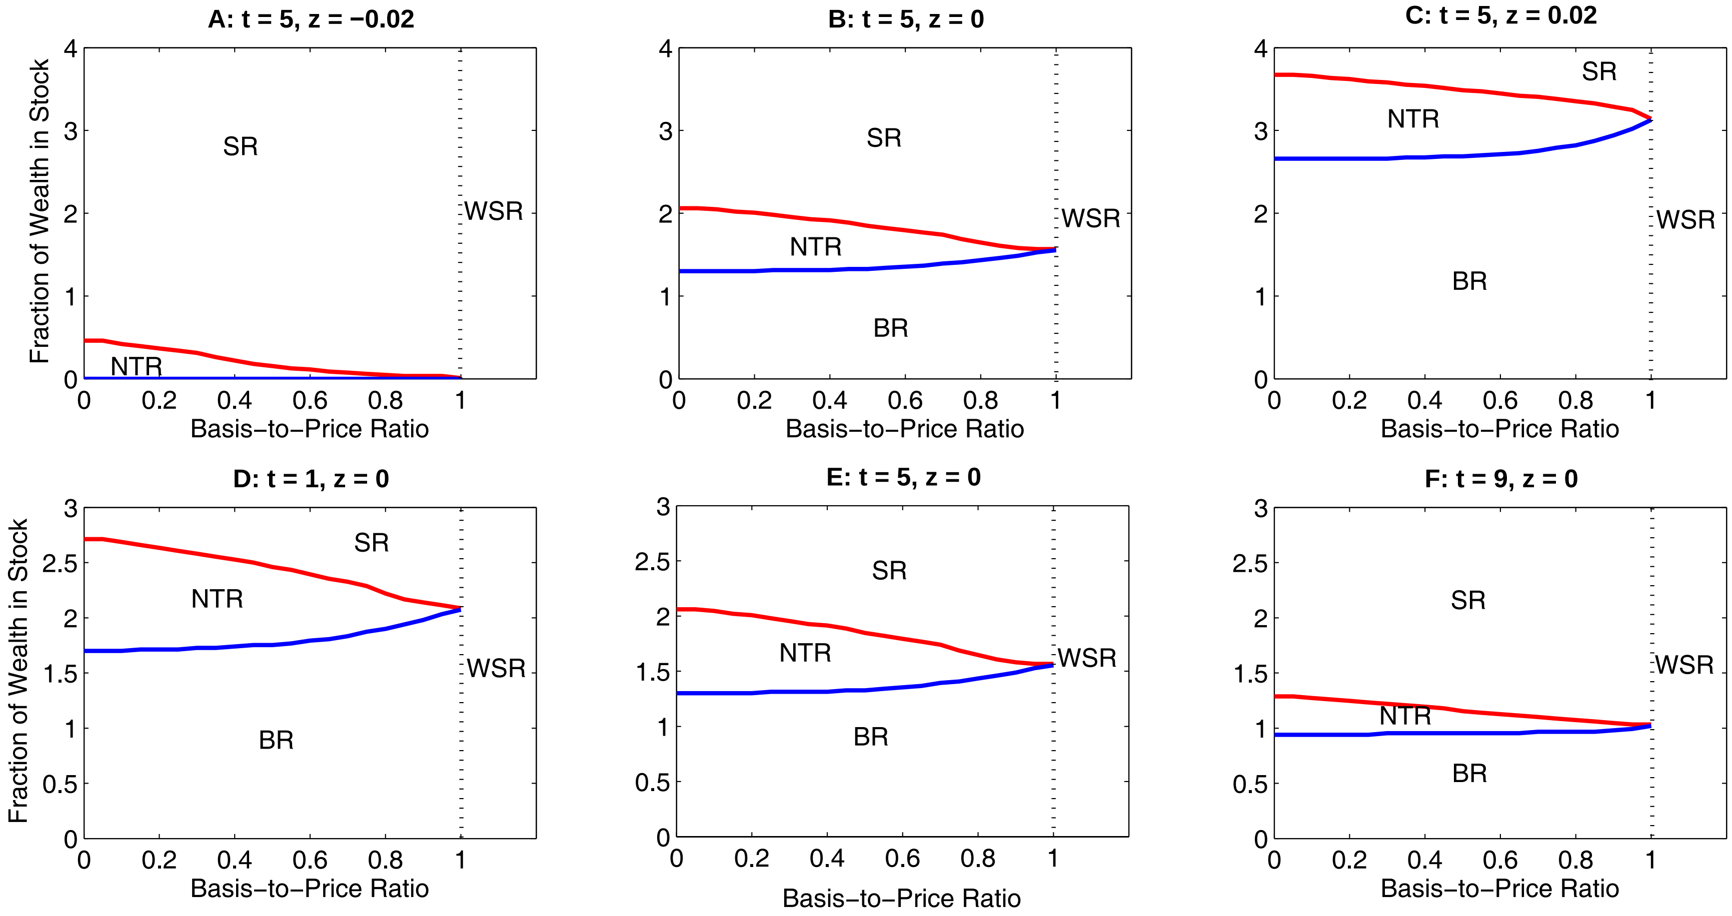

Figure 3: Optimal Policy with Return Predictability: The FR Case

This figure shows the optimal trading boundaries at $t = 5$ for three different values of the predictive variable $z = z_t - \bar{z}$ (in the three subfigures at the top), and the optimal trading boundaries at $z = z_t - \bar{z} = 0$ for three different points in time (in the three subfigures at the bottom), for the FR case. Note that in the figure, z denotes $z - \bar{z}$. Parameter values: $T = 10$, $\gamma = 5$, $r = 0.031$, $\mu_0 = 0.116$, $\mu_1 = 4.397$, $\sigma_S = 0.147$, $g_1 = 0.141$, $\bar{z} = 0.026$, $\sigma_z = 0.005$, $\rho = -0.895$, and $\tau = 0.25$.

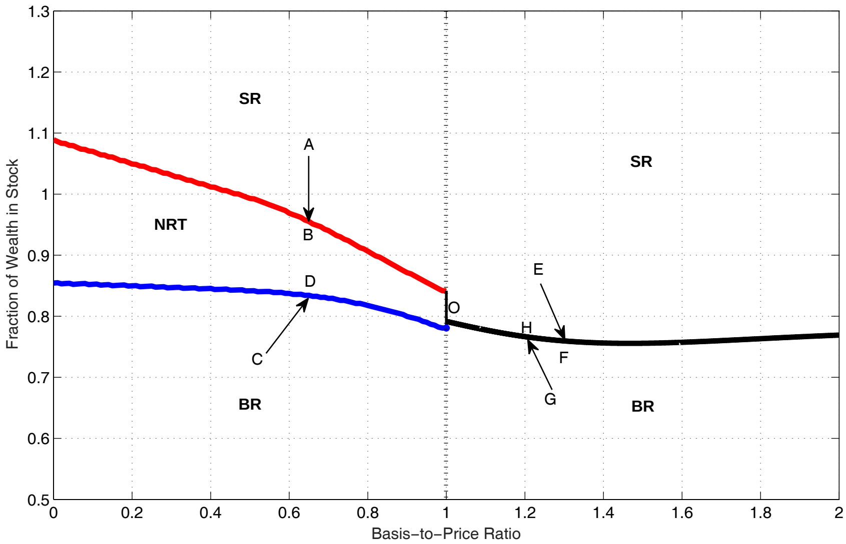

Figure 4: Optimal Policy without Return Predictability: The FC Case

This figure shows the optimal sell and buy boundaries at $t = 5$ in the presence of capital gains taxes, without return predictability. Parameter values: $T = 10$, $\gamma = 5$, $r = 0.031$, $\mu_0 = 0.116$, $\mu_1 = 0$, $\sigma_S = 0.147$, $g_1 = 0.141$, $\bar{z} = 0.026$, $\sigma_z = 0.005$, $\rho = -0.895$, and $\tau = 0.25$.

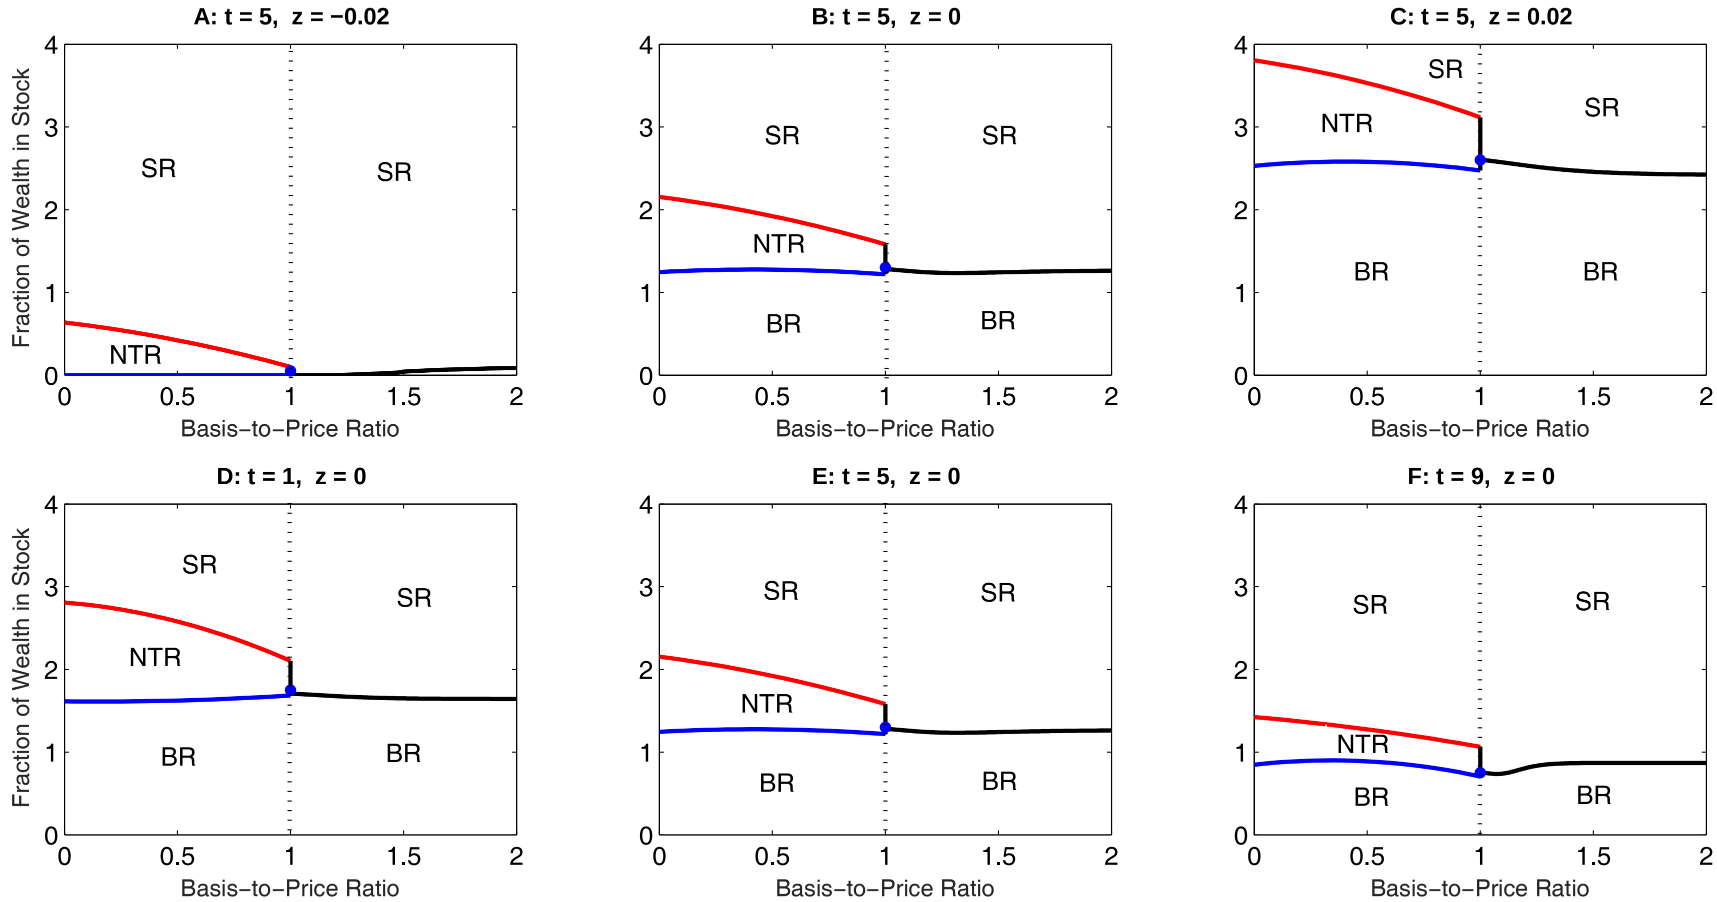

Figure 5: Optimal Policy with Return Predictability: The FC Case

This figure shows the optimal trading boundaries at $t = 5$ for three different values of the predictive variable $z = z_t - \bar{z}$ (in the three subfigures at the top), and the optimal trading boundaries at $z = z_t - \bar{z} = 0$ for three different points in time (in the three subfigures at the bottom), for the FC case. Note that in the figure, z denotes $z - \bar{z}$. Parameter values: $T = 10$, $\gamma = 5$, $r = 0.031$, $\mu_0 = 0.116$, $\mu_1 = 4.397$, $\sigma_S = 0.147$, $g_1 = 0.141$, $\bar{z} = 0.026$, $\sigma_z = 0.005$, $\rho = -0.895$, and $\tau = 0.25$.

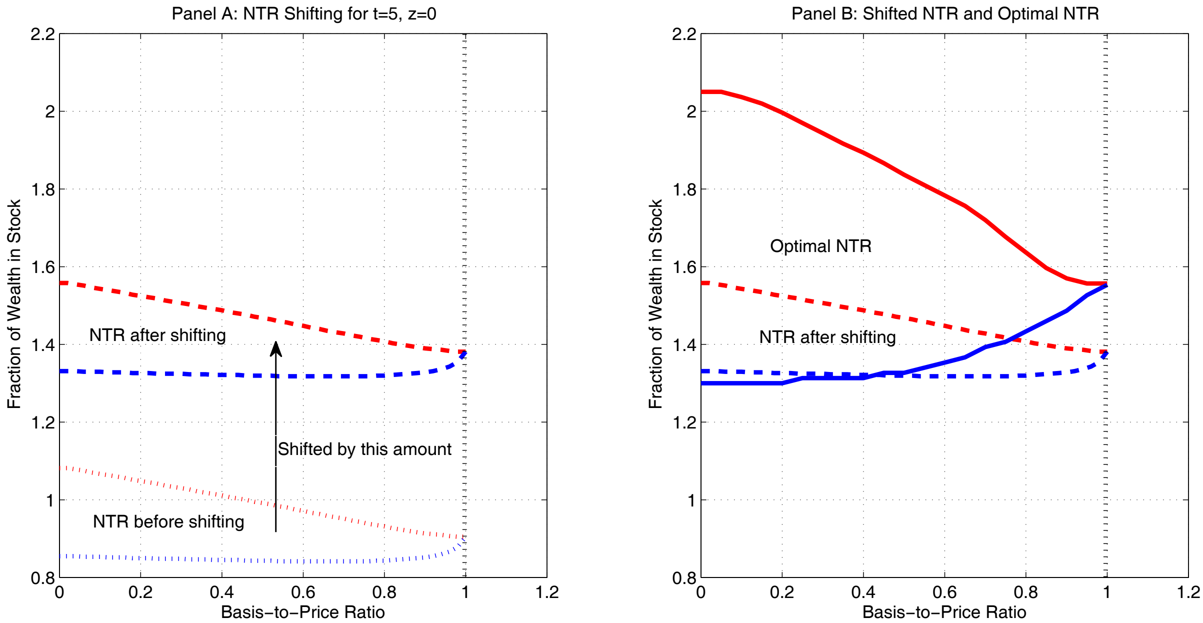

Figure 6: Shifted Policy and Optimal Policy: The FR Case

This figure shows the suboptimal trading strategy and the optimal trading strategy implied by our model. The suboptimal strategy is constructed by shifting the NTR with constant expected return by the market-timing and hedging demand induced by return predictability. Note that in the figure, z denotes $z - \bar{z}$. Parameter values: $T = 10$, $\gamma = 5$, $r = 0.031$, $\mu_0 = 0.116$, $\mu_1 = 4.397$, $\sigma_S = 0.147$, $g_1 = 0.141$, $\bar{z} = 0.026$, $\sigma_z = 0.005$, $\rho = -0.895$, and $\tau = 0.25$.

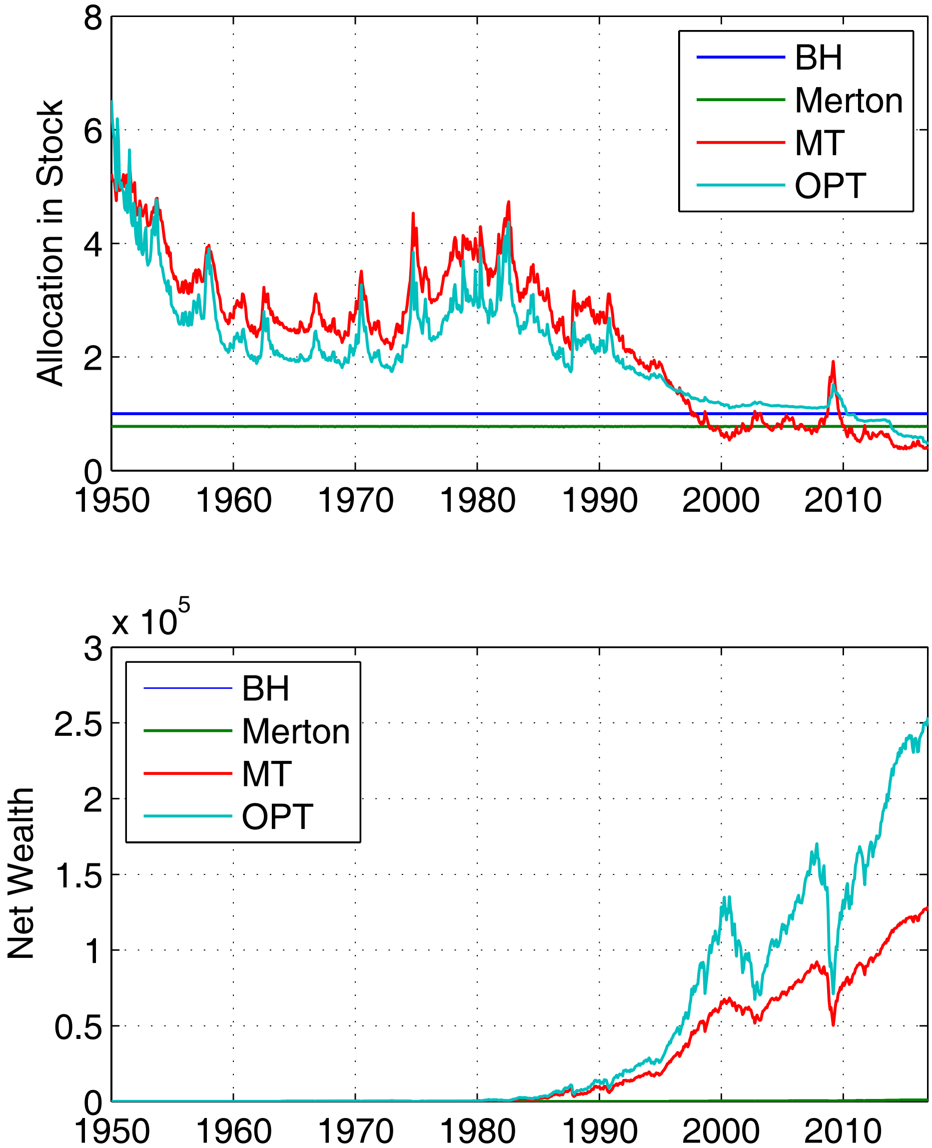

Figure 7: Empirical Simulation: Equity Exposure and Wealth Accumulation
This figure shows the risk exposures to the stock market and the wealth accumulation processes for four different trading strategies: buy-and-hold strategy (BH), Merton's constant ratio strategy (Merton), pure market-timing strategy (MT), and the optimal trading strategy implied by our model (OPT). The simulation uses historical returns of the weighted average U.S. stock market index from January 1950 to December 2016, and assumes constant interest and tax rates over this period. Parameter values: $W_0 = 1$, $T = 67$, $\gamma = 5$, $r = 0.031$, $\mu_0 = 0.116$, $\mu_1 = 4.397$, $\sigma_S = 0.147$, $g_1 = 0.141$, $\bar{z} = 0.026$, $\sigma_z = 0.005$, $\rho = -0.895$, and $\tau = 0.25$.

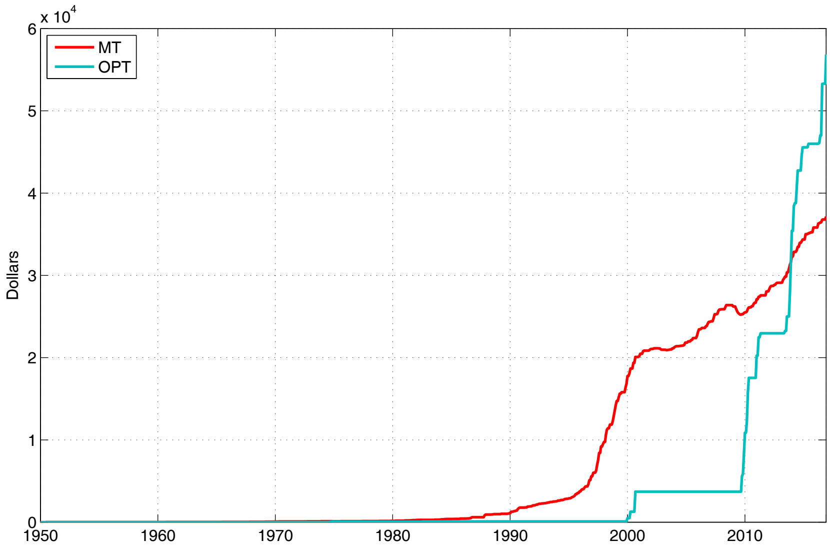

Figure 8: Cumulative Tax Bills in the Empirical Simulation

This figure shows the cumulative tax bills paid for the pure market-timing strategy (MT) and the optimal trading strategy implied by our model (OPT). The simulation uses historical returns of the weighted average U.S. stock market index from January 1950 to December 2016, and assumes constant interest and tax rates over this period. Parameter values: $W_0 = 1$, $T = 67$, $\gamma = 5$, $r = 0.031$, $\mu_0 = 0.116$, $\mu_1 = 4.397$, $\sigma_S = 0.147$, $g_1 = 0.141$, $\bar{z} = 0.026$, $\sigma_z = 0.005$, $\rho = -0.895$, and $\tau = 0.25$.