WIRTSCHAFTSWISSENSCHAFTLICHES ZENTRUM (WWZ)
DER UNIVERSITÄT BASEL

September 2005

Strategic Asset Allocation and the Risk of the Stock Market in the Long Run

WWZ Research Report 2005/ 02

David Rey, Heinz Zimmermann

WWZ Forum, Petersgraben 51, CH - 4003 Basel
Fax +41 61 267 33 33 | e-mail: forum-wwz@unibas.ch | www.wwz.unibas.ch/forum

The Authors:

David Rey
Department of Business and Economics (WWZ) at the University of Basel
Department of Finance
Holbeinstrasse 12
CH-4051 Basel
Email: david.rey@unibas.ch

Heinz Zimmermann
Department of Business and Economics (WWZ) at the University of Basel
Department of Finance
Holbeinstrasse 12
CH-4051 Basel
Email: heinz.zimmermann@unibas.ch

A publication of the University of Basel, Department of Business and Economics (WWZ)

Address for information:
WWZ Forum, Petersgraben 51, CH-4003 Basel Fax +41 61 267 33 33
The WWZ Research Papers document the results of research projects at the WWZ.

ISBN 3-909162-35-5

Mai 2005

Forschungsbericht D-91

Strategic Asset Allocation and the Risk of the Stock Market in the Long Run

A Critical Evaluation of Stock Market Predictability and Estimation Risk for the Swiss and U.S. Stock Market

David Rey, Heinz Zimmermann

Current discussions about public and private pension plans often include a statement that the stock market is less risky in the long run than in the short run. Pension plans with their rather long planning horizon are therefore asked to increase their allocation to the stock market. These statements, however, usually lack of a consistent clarification of the underlying reasons for the asserted positive horizon effect. It is thus important to be precise about any rational reasons for pension plans (and individual long-term investors alike) to allocate a higher fraction of their wealth to the risky stock market. One reason that has become prominent among financial economists only recently, is the mounting empirical evidence of stock market predictability. However, our empirical analysis demonstrates that return predictability and persistent movements in expected excess returns should not be interpreted as a reliable driving force for a positive horizon effect and the corresponding market timing strategies, least of all when estimation risk is taken into account. Although there is some (weak) evidence of stock market predictability, neither the Swiss nor the U.S. stock market seems to be less risky in the long run than in the short run. Motivated by return predictability, thus, long-term investors should not hold more stocks and should not time the market more aggressively than myopic short-term investors.

1. Introduction

Current discussions about public and private pension plans often include a statement that the stock market is less risky in the long run than in the short run. Pension plans with their rather long planning horizon are therefore asked to increase their allocation to the stock market. These statements, however, usually lack of a consistent clarification of the underlying reasons for such a positive horizon effect. Indeed, the question as to whether or not there is any rational reason for long-term investors to allocate a higher fraction of their wealth to the risky stock market created much controversy in the academic literature and the investment community alike (see, e.g., Samuelson (1994), Bodie (1995), Siegel (1998), and particularly Campbell and Viceira (2002)). One source of confusion is the common tendency to measure risk in units of standard deviation rather than variance. For example, the Sharpe ratio of any risky investment grows with the square root of the investment horizon. Referring to long-horizon Sharpe ratios thus somehow reduces risk in a manner analogous to the effect of diversifying a portfolio cross-sectionally. In much the same way, shortfall probabilities imply that it is "almost certain" that the stock market outperforms other asset classes (such as bonds) in the long run. However, it is quite easy to show that these "time diversification" arguments are a fallacy and do not imply too much about the optimal allocation to stocks of long-term investors. Sharpe ratios cannot be compared across different investment horizons, and shortfall probabilities do not take into consideration that any loss – should it occur – is much larger at long investment horizons, and that this possibility is heavily weighted by rational investors.

Any recommendation in favor of a positive horizon effect should rather be derived from an intertemporal expected utility maximization problem, as developed by Samuelson (1969) and Merton (1969, 1971, 1973). On the one hand, if investment opportunities are constant, their asset allocation models imply that investors with power utility who rebalance their portfolio optimally should choose the same portfolio of risky assets, regardless of their planning horizon – a sharp contrast to the common perception of investment managers that the time horizon is a major determinant of asset allocation decisions. On the other hand, however, Merton (1973) shows theoretically that time-varying investment opportunities may indeed introduce intertemporal horizon effects.

Yet, while the economic theory of optimal portfolio choice is by now quite well understood, the renewed interest in portfolio choice problems mainly follows the relatively recent empirical evidence of time-varying return distributions (stock market predictability and conditional heteroskedasticity). Indeed, the mounting empirical evidence of stock market predictability seems to have persuaded the majority of financial economists to abandon the constant expected excess returns paradigm. The time variation and predictability of (stock market) returns, labeled as a "new fact in finance" by Cochrane (1999a), is so widely accepted in the profession that it has generated, with no signs of subsiding, a new wave of conditional asset pricing models and dynamic asset allocation models that attempt to determine Merton's (1973) intertemporal hedging demand. Motivated by the empirical evidence of time-varying investment opportunities and particularly stock market predictability, Brennan, Schwartz, and Lagnado (1997), Campbell and Viceira (1999, 2002), Barberis (2000), and Xia (2001), to name just a few, thus forcefully suggest that long-term investors should allocate substantially more to stocks than short-term investors, i.e., that the horizon effect is positive and economically large.

Our paper, however, challenges this conclusion. Our empirical results rather suggest that stock market predictability and persistent movements in expected excess returns should not be interpreted as a reliable driving force for positive horizon effects, least of all when parameter uncertainty is taken into account. The contribution of this paper is thus primarily an empirical one. We do not intend to enlarge the existing universe of dynamic asset allocation models and econometric approaches (see, e.g., Brandt (2004) or Campbell and Viceira (2002) for up-to-date overviews). Instead, we mainly adopt the discrete-time framework of Barberis (2000), which is, in our view, representative for a large class of portfolio choice problems and econometric approaches.

In particular, there are two assets: a stock market index and a short-term risk-free asset. Investors with power utility over terminal wealth use a homoskedastic vector auto-regression (VAR) for the excess return and the dividend-price ratio to calculate expected excess returns and return volatility.[1] While CRRA preferences are by far the most popular because the value function turns out to be homogeneous in wealth, the homoskedastic VAR is probably not the most realistic data-generating process.[2] Nevertheless, it has been used extensively in the stock

[1] Both intermediate consumption decisions and labor income are thus neglected. Regarding consumption, Wachter (2002) demonstrates that the economic implication of introducing intermediate consumption in a CRRA framework is to shorten the effective planning horizon of the investor. In particular, she shows that the intertemporal hedging demands with intermediate consumption are a weighted sum of the hedging demands of a sequence of terminal wealth problems. Our main conclusions should therefore not be qualitatively affected by additionally taking consumption into account. This is not true with respect to labor income, however. Still, in this case, horizon effects are generally not motivated by stock market predictability (see, e.g., Campbell and Viceira (2002) and Viceira (2001)).

[2] CRRA preferences are not without faults, however. One critique that is particularly relevant in the portfolio choice context is that with power utility the elasticity of intertemporal substitution is directly tied to the level of relative risk aversion (one is the reciprocal of the other), which creates an unnatural link between two very different aspects of the investor's preferences – the willingness to substitute consumption intertemporally versus the willingness to take on risk. Epstein and Zin (1989) and Weil (1989) propose a generalization of CRRA

market predictability and portfolio choice literature (see, e.g., Avramov (2002), Campbell and Viceira (1999, 2002), Kandel and Stambaugh (1996), and Campbell (1990, 1991), to name just a few).[3]

Most importantly, however, Barberis (2000) emphasizes the impact of estimation risk (parameter uncertainty) on the optimal portfolio choice of long-term investors. Observing only a finite sample, investors do not know the true values of the predictive regression parameters. As a result, part of the risk they rationally perceive arises due to estimation risk. It is thus an important aspect of any portfolio choice problem that investors account for the fact that the true extent of stock market predictability is highly uncertain. In particular, Barberis (2000) shows that the perspective of a Bayesian investor (who uses the sample evidence to update prior beliefs about the regression parameters) is particularly useful to deal with estimation risk. Given the data, the uncertainty about the regression parameters is summarized by the posterior distribution of the parameters. Rather than constructing the distribution of future returns conditional on fixed parameter estimates, the Bayesian approach integrates over the uncertainty in the parameters captured by the posterior distribution. The resulting predictive distribution of future returns is then solely based on observed data, and no longer on any fixed parameter values. Barberis' (2000) results reveal that parameter uncertainty induces significant negative horizon effects that eliminate the positive horizon effects due to stock market predictability when estimation risk is ignored. Furthermore, his results also demonstrate that the often recommended allocations to stocks are too high and too sensitive to the prevailing values of the predictive variables parameterizing expected returns.

However, Barberis' (2000) conclusions are solely based on the dividend-price ratio as predictive variable. It is thus not clear how the other predictive variables found important in previous studies of stock market predictability would affect his results. Since existing equilibrium pricing theories are not explicit about which variables should enter the predictive regression, it renders the empirical evidence subject to data mining concerns (see, e.g., Lo and MacKinlay (1990) and Foster, Smith, and Whaley (1997)). For example, Bossaerts and Hillion (1999), Neely and Weller (1999), and particularly Goyal and Welch (2003a,b) point out that even their best prediction models have no out-of-sample forecasting power and fail to generate robust results that outperform simple unconditional benchmark models. Moreover, the multiplicity of potential predictive variables also makes the empirical evidence difficult to

---

preferences based on recursive utility that separates this link between intertemporal substitution and risk aversion.

[3] Of course, the analysis is therefore limited to the risk of a changing equity risk premium, the risk of changing volatility and the risk of changing real interest rates and inflation are not taken into consideration with this specification.

interpret: a predictive variable might be statistically significant based on a particular collection of predictive variables and sample period, but often not based on a competing specification and/or time period. Finally, Ferson, Sarkissian, and Simin (2003, 2004) argue that spurious regression can be a serious concern well outside the classic setting of Yule (1926) and Granger and Newbold (1974). Worse, they show that spurious regression and data mining interact in a very pernicious way. Searching for predictive variables using historical data can increase the likelihood of finding a persistent variable with spurious regression bias. Data mining makes spurious regression more of a problem, and the possibility of spurious regression makes data mining seem more effective in finding predictive variables. Indeed, their simulations imply that virtually all of the suggested predictive variables are consistent with a "spurious mining process". These predictive variables appear to have worked in the past, but will not work in the future and therefore fail to be practically useful. However, Ferson, Sarkissian, and Simin (2003, 2004) also show that a simple form of stochastic detrending lowers the persistence of the predictive variables, and can be used to reduce the risk of finding spurious predictive relations.

Given these specifications, monthly data on equity returns, and, in addition to the dividend-price ratio, a number of other predictive variables (including their stochastically detrended counterparts) for the Swiss and U.S. stock market from 1975 to 2002, we critically review the evidence of stock market predictability and its impact on the optimal allocation to stocks of buy-and-hold as well as optimally rebalancing long-term investors.

First, we show that the predictive power of the predictive variables is generally strong enough to make myopic short-term investors as well as intertemporally optimizing investors time the market. But the extent of these market timing strategies is not very consistent, it rather depends on the choice of the predictive variable, whether the predictive variable is stochastically detrended or not, the time period under consideration, and, of course, whether investors account for estimation risk or not. If they do not, the resulting allocation to stocks is excessively sensitive to the prevailing values of the predictive variables parameterizing expected returns.

Second, optimal allocations to stocks for long-term buy-and-hold investors are generally not very sensitive to the initial value of the predictive variables. There is no reliably positive horizon effect. In the buy-and-hold portfolio choice problem, the planning horizon corresponds to the investment horizon. Therefore, a necessary condition for a positive horizon effect is that expected returns increase (decrease) with the predictive variable while the correlation between innovations in returns and innovations in the predictive variable is negative (positive). In this case, realized returns may show mean reversion and conditional variances of cumulative returns grow less than linearly with the planning horizon. But our estimates reveal that this is solely true if expected excess returns are parameterized by the dividend-price ratio. Because expected excess returns increase with the dividend-price ratio and the correlation between innovations of excess returns and innovations of the dividend-price ratio is strongly negative, stock market risk increases less than linearly across planning horizons and stocks are perceived more attractive in the long run. Investors with a long planning horizon should therefore allocate significantly more to stocks than short-term investors.

After stochastically detrending the dividend-price ratio and/or incorporating estimation risk, however, this positive horizon effect largely disappears. The uncertainty about the predictive power of the dividend-price ratio is large enough to reverse the direction of the results. Long-horizon buy-and-hold investors who ignore estimation risk would overallocate to stocks by a sizeable amount. Thus, the weak statistical significance of stock market predictability makes it important to take parameter uncertainty into account. Indeed, for example, recently, from 1989 to 2002, any evidence of return predictability by the dividend-price ratio seems to have completely disappeared for the Swiss stock market.

Similar results are obtained for the term spread, realized stock market volatility, and the U.S. default risk spread as predictive variables. If these variables (or their stochastically detrended counterparts) parameterize expected excess returns, the implied risk of the stock market does not decrease in the long run. There is no implied mean reversion in realized returns and no positive horizon effect. To understand the absence of a positive horizon effect for these variables, our results show that (i) the evidence of stock market predictability is often very weak (and actually non-existent in most cases), (ii) the correlation between unexpected excess returns and innovations in the predictive variables is often not significantly different from zero, and, (iii), expected excess returns increase (decrease) with the predictive variable while the correlation between innovations in excess returns and innovations in the predictive variable is positive (negative). Therefore, conditional variances of cumulative returns do not grow less than linearly with the planning horizon, and long-term buy-and-hold investors should not hold more equities than short-term investors.

Finally, in the case of the intertemporal portfolio choice problem, any evidence of a positive horizon effect (i.e., a positive intertemporal hedging demand) is again limited to the case where expected returns are parameterized by the dividend-price ratio. The results here must be interpreted differently, however. In the multiperiod optimization framework, the investment horizon is shorter than the planning horizon. Investors thus face not only the current risk inherent in returns, but also the risk about whether the dividend-price ratio will be higher, lower, or the same in the future, and whether, consequently, future investment opportunities will improve, deteriorate or remain the same, respectively. Since our VAR estimates imply that positive (negative) excess returns tend to be associated with a drop (rise) in the dividend-price ratio and an expected utility loss (gain) due to deteriorated (improved) investment opportunities in the future, the investment opportunities risk can be smoothed quite effectively by over-investing in stocks, relative to the myopic allocation. The financial gain (loss) from over-investing partially offsets the expected utility loss (gain) associated with the drop (rise) in the dividend-price ratio.

However, as before, incorporating estimation risk has an adverse effect on the positive intertemporal hedging demand, but the impact is somewhat lower than for long-term buy-and-hold investors. Also, recently from 1989 to 2002, any evidence of stock market predictability by the dividend-price ratio, and therefore of a positive intertemporal hedging demand, seems to have completely disappeared for the Swiss stock market. Similarly, for the remaining and the stochastically detrended predictive variables (including the dividend-price ratio), neither the Swiss nor the U.S. stock market seems to provide a reliable intertemporal hedge to the risk of future changes in investment opportunities.

To sum up, from our empirical point of view, we conclude that stock market predictability and persistent movements in expected excess returns should not be interpreted as a reliable driving force for a positive horizon effect and the corresponding market timing strategies, least of all when estimation risk is taken into account. Although there is some (weak) evidence of stock market predictability, neither the Swiss nor the U.S. stock market seems to be less risky in the long run than in the short run. Long-term investors should therefore not hold more stocks and should not time the market more aggressively than myopic short-term investors.

We proceed as follows. Section 2 reviews the notion of strategic asset allocation and the stock market predictability literature. Section 3 describes the data-generating processes for both a constant and time-varying investment opportunity set. Mean reversion and implied variance ratio statistics in the presence of stock market predictability are discussed in Section 4. Section 5 reviews Barberis' (2000) asset allocation framework for buy-and-hold and intertemporally optimizing investors. Section 6 describes the data and presents the empirical results. Section 7 concludes.

2. Strategic Asset Allocation and Stock Market Predictability

The search for the optimal portfolio of risky assets is a demanding task – for the financial planning industry as well as for financial economists. Optimal portfolio decisions of individual and institutional investors depend on countless details of their circumstances. These circumstances include the financial assets that are available (Treasury bills, bonds, inflation-indexed bonds, stocks, real estate, but also derivatives-based structured products, commodity funds, private equity, hedge funds, and illiquid assets such as family businesses and restricted stock options granted by employers) and their expected returns and risks, the characteristics of the investors' (non-tradable) labor income (human capital) and life-cycle liabilities, as well as their tax code and transaction, participation, and information costs. There are thus many legitimate reasons why different portfolios of risky assets may be appropriate for different investors.

Beyond these circumstances, investors also differ in their degree of risk aversion. For example, conservative investors are typically encouraged to hold more bonds relative to stocks. This financial planning advice, however, does not conform to the well-known mean-variance analysis of Markowitz (1952). Indeed, the Markowitz paradigm (or, equivalently, the two-fund separation theorem of Tobin (1958)) established that investors who care only about the expected return and risk will hold the same portfolio of risky assets, irrespective of their tolerance for risk. Accordingly, Canner, Mankiw, and Weil (1997) call this advice an "asset allocation puzzle".

In addition, although the mean-variance analysis captures the two fundamental aspects of portfolio choice, namely, diversification and the trade-off between expected return and risk, a number of theoretical objections to simple mean-variance optimization have emerged, too. First, the underlying utility function of the mean-variance problem can at best be interpreted as a second-order approximation of expected utility maximization (except for the special case of quadratic utility, which is, however, a problematic preference specification because it is not monotonically increasing in wealth). Second, the mean-variance problem ignores any preferences toward higher-order return moments, such as return skewness and kurtosis. Finally and most importantly, the mean-variance problem is inherently a myopic, single-period portfolio choice problem, whereas most investment strategies likely involve longer planning horizons with intermediate portfolio rebalancing (and consumption decisions). As opposed to making a sequence of myopic single-period portfolio choices, the most straightforward way to address these issues is to formulate the portfolio choice problem explicitly as an intertemporal expected utility maximization problem, called "strategic asset allocation" by Brennan, Schwartz, and Lagnado (1997).[4]

Strategic Asset Allocation

The single-period mean-variance analysis usually (but not necessarily) assumes a short planning horizon, and investment opportunities are well characterized by constant expected returns and risks. In contrast, investors with a longer planning horizon must concern themselves not only with expected returns and risks today, but also with the way in which these may change over time and the way they may optimally respond to such time-varying investment opportunities. Indeed, while the classic work of Samuelson (1969) and Merton (1969, 1971) find that when investment opportunities are constant, investors with power utility who rebalance their portfolio optimally should choose the same portfolio of risky assets, regardless of their planning horizon, already Merton (1973) notes that time-varying investment opportunities can potentially introduce intertemporal horizon effects. In this case, optimal portfolio weights are the sum of two terms, the first being the myopically optimal portfolio weights and the second representing the intertemporal hedging demand (the third fund in Merton's (1973) three-fund separation theorem). Specifically, the myopic portfolio weights depend on the ratio of the first to second moments of the asset returns and on the inverse of the investor's relative risk aversion. The second term depends on the projection of the innovations in the predictive variables onto innovations in the asset returns, on the inverse of the investor's relative risk aversion, and on the sensitivity of the investor's marginal utility to the predictive variables. For each predictive variable suspected relevant, the projection delivers the weights of the corresponding intertemporal hedging portfolios that are maximally correlated with the predictive variables' innovations. Intuitively, investors take positions in each of the maximally correlated portfolios to partially hedge against undesirable innovations in the predictive variables. Analogous to diversifying cross-sectionally, investors thus want to smooth the risk of changes in future investment opportunities intertemporally.

Nevertheless, the buy-and-hold portfolio choice problem is an important special case. First, the portfolio choice problem reduces to a single-period problem when investors have logarithmic preferences. Second, intertemporal hedging demands only arise when investment opportunities vary stochastically through time, constant investment opportunities (independently and identically distributed (i.i.d.) asset returns) do not lead to intertemporal

[4] An alternative definition of strategic (and tactical) asset allocation is suggested in Rey (2004).

hedging demands. Finally, even with time-varying investment opportunities, intertemporal horizon effects only arise when investors can use the available assets to hedge against these changes in future investment opportunities. If the time variation in investment opportunities is completely independent of the asset returns, the optimal portfolio is myopic again.[5]

Yet, while the economic theory of optimal portfolio choice is by now quite well understood, the renewed interest in portfolio choice problems mainly follows the relatively recent empirical evidence of time-varying return distributions (stock market predictability and conditional heteroskedasticity). Indeed, the mounting empirical evidence of stock market predictability seems to have persuaded the majority of financial economists to abandon the constant expected excess returns paradigm. The time variation and predictability of (stock market) returns, labeled as a "new fact in finance" by Cochrane (1999a), is so widely accepted in the profession that it has generated, with no signs of subsiding, a new wave of conditional asset pricing models and dynamic asset allocation models that attempt to determine Merton's (1973) intertemporal hedging demand.

Stock Market Predictability

Campbell and Shiller (1988b) and Fama and French (1988a, 1989), for example, are among the first who document that the dividend yield and particularly the dividend-price ratio on aggregate stock portfolios predict future (long-horizon) (stock market) returns. Other examples of predictive variables include short-term interest rates (e.g., Fama and Schwert (1977) and Campbell (1991)), yield spreads between long-term and short-term interest rates and between low- and high-quality bond yields (e.g., Keim and Stambaugh (1986), Campbell (1987), and Fama and French (1989)), stock market volatility (e.g., French, Schwert, and Stambaugh (1987) and Goyal and Santa-Clara (2003)), Eurodollar-U.S. Treasury (TED) spread (e.g., Ferson and Harvey (1993)), book-to-market ratios (e.g., Kothari and Shanken (1997) and Pontiff and Schall (1998)), dividend-payout and price-earnings ratios (e.g., Lamont (1998) and Campbell and Shiller (1988a)), and more complex measures based on analysts' forecasts (e.g., Lee, Myers, and Swaminathan (1999)). Recently, Baker and Wurgler (2000) have shown that the share of equity in new finance is a negative predictor of future

[5] The myopic portfolio choice problem is an important special case for practitioners and financial economists alike. There are, to our knowledge, few financial institutions that implement multiperiod investment strategies involving intertemporal hedging demands. A common justification from practitioners is that the expected utility loss from errors that could creep into the solution of a complicated dynamic optimization problem outweighs the expected utility gain from investing optimally as opposed to myopically (see, e.g., Brandt (2004)). Our empirical findings following below strongly support this view.

equity returns. Lettau and Ludvigson (2001) find evidence of predictability using a consumption-wealth ratio, the level of consumption relative to income and wealth.

At the same time, however, there is less consensus on what drives this stock market predictability. Bekaert (2001) differentiates between three possibilities: it may indeed reflect time-varying investment opportunities (which he calls the "risk view"), it may reflect irrational behavior on the part of market participants (the "behavioral view") or it may simply not present in the data – a statistical fluke due to poor statistical inference (the "statistical view").

Of course, whether stock market predictability is consistent with market efficiency can only be interpreted in conjunction with an intertemporal equilibrium model of the economy. Inevitably, thus, all theoretical attempts at interpretation of predictability will be model-dependent, and hence inconclusive. Nevertheless, while much work has focused on the second possibility, recent advances in rational asset pricing theory seem to demonstrate that a certain degree of time-varying expected returns is necessary to reward investors for bearing certain dynamic risks associated with the business cycle. Loosely, it is claimed that the equity premium rises during an economic slow-down and falls during periods of economic growth, so that expected returns and business conditions move in opposite directions (see, e.g., Fama and French (1989), Chen (1991), Fama (1991), and Ferson and Harvey (1991)). Consequently, stock market predictability on its own would not imply stock market inefficiency. On the other hand, the risk view has sometimes been discredited because the empirical evidence suggests substantial time variation in risk premiums that can not (yet) be delivered by standard (still extremely stylized) models of risk. Moreover, other aspects of the empirical research on predicting returns remain controversial. Specifically, it seems that the statistical view gains increasing credibility considering the long list of authors criticizing the statistical methodologies in the stock market predictability literature. Indeed, a number of recent contributions (see, e.g., Rey (2003a,b) for an overview) suggest that standard statistical inference procedures overreject the null hypothesis of no predictability, so that any evidence of predictability may have more to do with poorly behaved test statistics than with stock market predictability. Properly adjusting for small sample biases, near unit roots (particularly so for the price-scaled predictive variables such as the dividend-price ratio), and other statistical issues associated primarily with long-horizon regressions weakens and often reverses many of the standard conclusions. It thus seems that the evidence of stock market predictability is much less transparent than the early work may have suggested.

Despite these critical objections to the empirical evidence of stock market predictability, Brennan, Schwartz, and Lagnado (1997), Campbell and Viceira (1999, 2002), Barberis (2000), and Xia (2001), to name just a few, express the view that long-term investors should allocate substantially more to stocks than short-term investors, i.e., that the buy-and-hold horizon effect as well as the intertemporal hedging demand are positive and economically large.

Our paper, however, challenges this conclusion. In the spirit of Bekaert's (2001) statistical view, our empirical results rather suggest that stock market predictability and persistent movements in expected excess returns should not be interpreted as a reliable driving force for positive horizon effects, least of all when parameter uncertainty is taken into account.

3. Constant vs. Time-Varying Investment Opportunities

Investment opportunities are either constant or time-varying. The obvious case of constant investment opportunities is the temporally independently and identically distributed (i.i.d.) normal model,

$$e_t = \mu + \xi_t,$$

where the continuously compounded excess return, $e_t = r_t - r_{f,t}$, is defined as the log nominal return, r_t, less the prevailing log nominal short-term risk-free rate, $r_{f,t}$. Thus, e_t follows a random walk with mean μ and normally distributed increments,

$$\xi_t \sim \text{i.i.d. } N\left(0, \sigma_\xi^2\right).$$

In the case of a time-varying return distribution, we follow Barberis (2000) and implement stock market predictability with the following restricted and homoskedastic vector auto-regression (VAR) for the log excess return and, in general, a number of predictive variables, $\mathbf{x}_t$,

$$\mathbf{z}_t = \mathbf{a} + \mathbf{B}\mathbf{x}_{t-1} + \boldsymbol{\xi}_t,$$

with

$$\mathbf{z}_t' = [e_t \ \mathbf{x}_t'], \ \mathbf{x}_t = [x_{1,t} \cdots x_{n,t}]', \text{ and } \boldsymbol{\xi}_t \sim \text{i.i.d. } N(\mathbf{0}, \boldsymbol{\Sigma}).$$

In the empirical analysis following below, we consider the special case of n = 1, i.e., $\mathbf{z}_T$ contains only two components: the log excess return and a single predictive variable, x_t. We can thus write the predictive regression model as

$$\begin{aligned} e_t &= \alpha + \beta x_{t-1} + \xi_t \\ x_t &= \gamma + \delta x_{t-1} + \eta_t, \end{aligned}$$

where

$$\begin{pmatrix} \xi_t \\ \eta_t \end{pmatrix} \sim N\left(\mathbf{0}, \begin{pmatrix} \sigma_\xi^2 & \sigma_{\xi\eta} \\ \sigma_{\xi\eta} & \sigma_\eta^2 \end{pmatrix}\right).$$

and $\sigma_{\xi\eta}$ denotes the covariance between the innovations ξ_t and η_t.

4. Mean Reversion: Variance Ratio Statistics

The risk of the stock market is commonly measured by the variance of continuously compounded one-period excess returns, $Var(e_t) = \sigma^2_\xi$, and, in the long run, by the variance of cumulative (multi-period) log excess returns, $Var(e_{t\text{->}t+q\text{-}1}) = Var(e_t + e_{t+1} + \ldots + e_{t+q\text{-}1})$. At first glance, thus, whether the stock market is less risky in the long run than in the short run crucially depends on the relation between $Var(e_t)$ and $Var(e_{t\text{->}t+q\text{-}1})$. Depending on the underlying data-generating process, the q-period variance ratio statistic, $VR(q)$,

$$VR(q) \equiv \frac{Var(e_{t\to t+q-1})}{qVar(e_t)} = 1 + 2\sum_{j=1}^{q-1}\left(1 - \frac{j}{q}\right) Corr(e_t, e_{t+j}),$$

is either below, equal to, or above one, indicating that the risk of the stock market will either grow slower than linearly, linearly, or faster than linearly with the time horizon, respectively. Hence, in the case of the random walk i.i.d. normal model, estimates for $VR(q)$ should be statistically indistinguishable from one: log excess returns are not autocorrelated and the variance is a linear function of the time horizon.[6] In contrast, in the case of stock market predictability, estimates for $VR(q)$ may deviate from one. Indeed, the stock market predictability literature often claims that stock market returns show mean reversion, i.e., $VR(q) < 1$, indicating that the risk grows less than linearly with the time horizon when investment opportunities are time-varying.

Variance ratio statistics can be estimated in two ways. On the one hand, estimates for $VR(q)$ may be obtained directly from the raw return data, i.e., without first specifying a data-generating process. Because $VR(q)$ is a particular linear combination of the first $q - 1$ autocorrelation coefficients of $\{e_t\}$, $Corr(e_t, e_{t+j})$, estimates for $VR(q)$ are easily obtained from estimates of the respective autocorrelation coefficients.

---

[6] Of course, the volatility/standard deviation of log excess returns grows less than linearly with the time horizon, namely with the square root of q. However, interpreting this "time diversification" argument as evidence that the stock market is less risky in the long run than in the short run is a fallacy. The same is true for other measures of risk, such as long-horizon Sharpe ratios and shortfall probabilities (see, e.g., Campbell and Viceira (2002) and Zimmermann (2005)).

On the other hand, estimates for $VR(q)$ may be obtained from estimates of the VAR model. In this case, both univariate (unconditional) and multivariate (conditional) variance ratio statistics can be calculated. Given the predictive regression model in equation (.), the calculation of unconditional and conditional variances of one-period and multi-period log excess returns, and hence of unconditional and conditional variance ratio statistics, is straightforward. Moreover, it is easy to show that the unconditional variance ratio statistic can be expressed as

$$VR(q) = VR_{t-1}(q) \frac{1 - R^2(e_t)}{1 - R^2(e_{t\to t+q-1})},$$

where $VR_{t-1}(q)$ denotes the conditional q-period variance ratio statistic, and $R^2(e_t)$ and $R^2(e_{t->t+q-1})$ denote the coefficients of determination of the predictive regression in equation (.) and a long-horizon predictive regression of the form

$$e_{t\to t+q-1} = \alpha(q) + \beta(q) x_{t-1} + \xi_{t\to t+q-1},$$

respectively. Consequently, as long as $R^2(e_{t->t+q-1}) > R^2(e_t)$, the multivariate variance ratio statistic is below the corresponding univariate variance ratio statistic. Therefore, (strong) mean reversion is sufficient, but not necessary for the existence of a positive horizon effect: univariate variance ratio statistics understate the risk-reduction that is relevant for long-term investors.

Given the predictive regression model in equation (.), estimates for the implied multivariate $R^2(e_{t->t+q-1})$ may be obtained as

$$R^2(e_{t\to t+q-1}) = 1 - \frac{1 - Var_{t-1}(e_{t\to t+q-1})}{1 - Var(e_{t\to t+q-1})},$$

where $Var_{t-1}(e_{t->t+q-1})$ denotes the conditional variance of cumulative log excess returns. The corresponding implied univariate coefficient of determination corresponds to the regression

$$e_{t\to t+q-1} = \alpha(q) + \beta(q) e_{t-q\to t-1} + \xi_{t\to t+q-1},$$

and is simply the square of the regression coefficient $\beta(q)$. Fama and French (1988b) show that the regression coefficient is related to the univariate q-period variance ratio statistic by

$$\beta(q) = \frac{VR(2q)}{VR(q)} - 1.$$

Note, however, variance ratio statistics below one do not necessarily imply a positive horizon effect. Besides the dynamics of the variance ratio statistics, the dynamics of conditional expected cumulative log excess returns, and, depending on the underlying utility function, even higher moments play a crucial role in the determination of the optimal allocation to stocks. Furthermore, when risk grows less than linearly with the time horizon, investment opportunities are time-varying and it is optimal to formulate the portfolio choice problem as an intertemporal expected utility maximization. In this case, variance ratio statistics do not capture the relevant trade-off between expected return and risk of the corresponding intertemporal strategy. In general, thus, variance ratio statistics below one are not sufficient for long-term investors to allocate a higher fraction of wealth to the risky stock market than short-term investors.

5. The Portfolio Choice Problem

As described above, there are two different ways of formulating the portfolio choice problem (see Table 1 for an overview).

[Insert Table 1 about here]

The first possibility is a buy-and-hold strategy. In this case, investors with a planning horizon of *T* months choose an allocation at the beginning of the first month, and do neither adjust nor rebalance their portfolio again until the *T* months are over. The second possibility of formulating the portfolio choice problem is a dynamic rebalancing strategy. Following this strategy, investors split their planning horizon of *T* months into several investment horizons (rebalancing intervals). They then choose an allocation at the beginning of the first investment period, knowing that at the start of every new investment period, they will optimally adjust their portfolio using the new information available at that time. In contrast to the buy-and-hold strategy, thus, the planning horizon is longer than the investment horizon for intertemporally optimizing investors. In this case, it is also important to distinguish between "dynamic rebalancing" and "static rebalancing". Following a myopic rebalancing strategy, investors choose an allocation at the beginning of the first investment period, knowing that they will reset (i.e., effectively rebalance) their portfolio to that same allocation at the start of every new investment period. While this is myopic in the sense that investors do not use any of the new information they have once an investment period has passed, it is the optimal intertemporal strategy if investment opportunities are constant.[7] In contrast, dynamic rebalancing means that investors adjust their portfolio using the new information at each time, likely resulting in optimal portfolio weights varying over time. Finally, in both cases of

[7] Static rebalancing thus corresponds to the cases (4a) and (4b) in Figure 1, respectively. Given the empirical results for the cases (1a) and (1b), however, cases (4a) and (4b) will not add any new insights. In what follows, we therefore skip the discussion of static rebalancing.

formulating the portfolio choice problem, investors may ignore estimation risk or take it into account.

The following sections briefly review Barberis' (2000) framework for investigating how stock market predictability and parameter uncertainty affect the optimal portfolio choice of buy-and-hold and intertemporally optimizing investors.

5.1. The Buy-and-Hold Portfolio Choice Problem

Suppose that a buy-and-hold investor is at time *T*, has a planning horizon of *T* months, and can choose from the following two assets: a short-term risk-free interest rate and a stock market index. For simplicity, the continuously compounded real monthly risk-free rate is assumed to be constant, r_f.

If initial wealth is normalized to $W_T = 1$, and $\omega = \omega_{T->T+T}$ is the allocation to stocks, then end-of-horizon wealth is trivially given by

$$W_{T+\hat{T}} = (1-\omega)\exp\left(r_f\hat{T}\right) + \omega\exp\left(r_f\hat{T} + e_{T+1} + ... + e_{T+\hat{T}}\right).$$

As in Barberis (2000), we further assume that the preferences over terminal wealth are described by the constant relative risk aversion (CRRA) power utility function

$$v\left(W_{T+\hat{T}}\right) = \frac{W_{T+\hat{T}}^{1-\zeta}}{1-\zeta},$$

where ζ denotes the coefficient of relative risk aversion. Given the cumulative log excess return over *T* periods by

$$e_{T\to T+\hat{T}} \equiv e_{T+1} + e_{T+2} + ... + e_{T+\hat{T}},$$

a buy-and-hold investor maximizes expected utility according to

$$\max_{\omega} E_T\left(\frac{\left\{(1-\omega)\exp\left(r_f\hat{T}\right) + \omega\exp\left(r_f\hat{T} + e_{T\to T+\hat{T}}\right)\right\}^{1-\zeta}}{1-\zeta}\right).$$

Depending on whether investors account for estimation risk or not, the distribution they should use in calculating the above conditional expectation is very different. Ignoring estimation risk, the conditional distribution of cumulative excess returns is given by

$$p\left(e_{T\to T+\hat{T}} \mid \hat{\boldsymbol{\theta}}, \mathbf{z}\right),$$

where **z** is the data observed by the investors up until the start of their planning horizon and $\boldsymbol{\theta}$ denotes the set of regression parameter estimates. Investors then maximize expected utility according to

$$\max_{\omega} \int \upsilon\left(W_{T+\hat{T}}\right) p\left(e_{T\to T+\hat{T}} \mid \hat{\boldsymbol{\theta}}, \mathbf{z}\right) de_{T\to T+\hat{T}}.$$

While this approach proceeds from the assumption that the investors know the parametric form of the return distribution, it ignores the fact that the true values of the regression parameters in $\boldsymbol{\theta}$ are not known precisely. To take parameter uncertainty into account, Barberis (2000) proposes a Bayesian approach of a posterior distribution, which summarizes the uncertainty about the regression parameters. This posterior distribution can then be used to integrate out the unknown regression parameters from the return distribution to obtain the investor's subjective (since it involves subjective priors) return distribution. In this case, the predictive distribution is conditional only on the sample observed, and no longer on any fixed parameter estimate,

$$p\left(e_{T\to T+\hat{T}} \mid \mathbf{z}\right) = \int p\left(e_{T\to T+\hat{T}} \mid \boldsymbol{\theta}, \mathbf{z}\right) p\left(\boldsymbol{\theta} \mid \mathbf{z}\right) d\boldsymbol{\theta},$$

and the resulting portfolio choice problem for the buy-and-hold investor to solve is

$$\max_{\omega} \int \upsilon\left(W_{T+\hat{T}}\right) p\left(e_{T\to T+\hat{T}} \mid \mathbf{z}\right) de_{T\to T+\hat{T}}.$$

In both cases, we follow Barberis (2000) and evaluate these integrals numerically by simulation.

When estimation risk is ignored, the conditional distribution of cumulative excess returns is normal. Therefore, the integral in equation (.) is approximated by generating a large sample of independent draws from the corresponding Normal distribution, and averaging $\upsilon(W_{T+T})$ over all these draws. Alternatively, Campbell and Viceira (2002) show that the optimal allocation to stocks can be approximated by

$$\omega \approx \frac{1}{\zeta} \frac{E_T\left(e_{T\to T+\hat{T}}\right) + \frac{1}{2} Var_T\left(e_{T\to T+\hat{T}}\right)}{Var_T\left(e_{T\to T+\hat{T}}\right)}.$$

Hence, the optimal allocation to stocks depends on the ratio between expected cumulative log excess returns with the addition of one-half the conditional variance (to convert from log returns to simple returns that are ultimately of concern to investors) and the conditional variance.

When estimation risk is taken into account, we first sample from the posterior distribution and then from the conditional

$$p\left(e_{T\to T+\hat{T}} \mid \boldsymbol{\theta}, \mathbf{z}\right).$$

Thus, rather than solving the optimization problem for a single choice of parameter values, the investor effectively solves an average problem over all possible set of parameter values.

The choice of either an informative or uninformative prior is crucial in Bayesian statistics. Most applications in finance employ uninformative priors, with the reasoning that empirical results with uninformative priors are most comparable to results obtained through classical statistics and therefore are easier to relate to the literature (see, e.g., Brandt (2004)). To construct the posterior distribution, we follow Barberis (2000) and use the following two standard diffuse priors,

$$p\left(\mu, \sigma_\xi^2\right) \propto \frac{1}{\sigma_\xi^2}$$

in the case of i.i.d. excess returns, and

$$p\left(\mathbf{C}, \mathbf{\Sigma}\right) \propto \left[\mathbf{\Sigma}\right]^{-3/2}$$

in the case of stock market predictability. $\mathbf{C}$ is a matrix with top row $\mathbf{a}'$ and the vector $\mathbf{B}'$ below that (see also Hamilton (1994) and Stambaugh (1999)).[8]

To get an accurate representation of the posterior distribution $p(\mu,\ \sigma^2_\xi;\ \mathbf{e})$ in the case of constant investment opportunities, we first sample from the marginal $p(\sigma^2_\xi;\ \mathbf{e})$, an inverse Gamma distribution, and then, given the σ^2 draw, from the conditional $p(\mu,\ \sigma^2_\xi;\ \mathbf{e})$, a Normal distribution. To obtain a large sample from the predictive distribution, we sample one point from the Normal distribution with mean $T\mu$ and variance $T\sigma^2_\xi$ for each draw of μ and σ^2_ξ from the posterior distribution. In the case of stock market predictability, we sample from the posterior distribution by first drawing from the marginal $p(\mathbf{\Sigma}^{-1};\ \mathbf{z})$, a Wishart distribution, and then, given the $\mathbf{\Sigma}^{-1}$ draw, from the conditional $p(\mathbf{C},\ \mathbf{\Sigma};\ \mathbf{z})$, a multivariate Normal distribution. Repeating this many times gives an accurate representation of the posterior distribution. Second, for each draw of $\mathbf{a}$, $\mathbf{B}$ and $\mathbf{\Sigma}$ from the posterior distribution, we sample from the Normal distribution with mean vector

$$\begin{aligned}\mu_{T\rightarrow T+\hat{T}} &= \hat{T}\mathbf{a} + \left(\hat{T}-1\right)\mathbf{B}_0\mathbf{a} + \left(\hat{T}-2\right)\mathbf{B}_0^2\mathbf{a} + \ldots + \mathbf{B}_0^{\hat{T}-1}\mathbf{a} \\ &\quad + \left(\mathbf{B}_0 + \mathbf{B}_0^2 + \ldots + \mathbf{B}_0^{\hat{T}}\right)\mathbf{z}_T,\end{aligned}$$

variance matrix

[8] As indicated in Barberis (2000, footnote 10), another reasonable approach would be to use a more informative prior that puts zero weight on negative values of μ, reflecting the observation in Merton (1980) that the expected market risk premium should be positive; see also Boudoukh, Richardson, and Smith (1993), Brandt (2004), and Campbell and Thompson (2005).

$$
\begin{aligned}
&\mathbf{\Sigma}_{T\to T+\hat{T}} \\
=\ &\mathbf{\Sigma} \\
&+\left(\mathbf{I}+\mathbf{B}_0\right)\mathbf{\Sigma}\left(\mathbf{I}+\mathbf{B}_0\right)' \\
&+\left(\mathbf{I}+\mathbf{B}_0+\mathbf{B}_0^2\right)\mathbf{\Sigma}\left(\mathbf{I}+\mathbf{B}_0+\mathbf{B}_0^2\right)' \\
&\vdots \\
&+\left(\mathbf{I}+\mathbf{B}_0+\mathbf{B}_0^2+...+\mathbf{B}_0^{\hat{T}-1}\right)\mathbf{\Sigma}\left(\mathbf{I}+\mathbf{B}_0+\mathbf{B}_0^2+...+\mathbf{B}_0^{\hat{T}-1}\right)',
\end{aligned}
$$

with

$$
\mathbf{B}_0=[\mathbf{0}\ \mathbf{B}].
$$

This gives a large sample of the predictive distribution, which we can use to compute the optimal allocation, ω, when taking both return predictability and estimation risk into account. Note that the initial value of the predictive variable, x_T, enter the conditional mean vector through $\mathbf{z}_T$.

The numerical procedure is repeated for planning horizons ranging from one year to ten years in one-year increments, for several values of risk aversion coefficients, and, in the case of return predictability, for several initial values of the predictive variable.

5.2. The Intertemporal Portfolio Choice Problem

An investor who optimally adjusts his portfolio at regular intervals is confronted with a dynamic programming problem. We follow Barberis (2000) and solve this problem by employing the standard technique of discretizing the state space and using backward induction.

Again, suppose that at time T, an investor has a planning horizon of T months. His planning horizon may then be divided into K (rebalancing) intervals of equal length, $[t_0 = T, t_1]$, $[t_1, t_2]$,…, $[t_{K-1}, t_K = T+T]$. The investor is thus allowed to adjust his portfolio K times over the planning horizon, at points $(t_0, t_1,…, t_{K-1})$, and the control variables at his disposal are $(\omega_0, \omega_1,…, \omega_{K-1})$. Specifically, we consider an investor who is allowed to adjust his portfolio annually using the new information at the end of each year.

Barberis (2000) shows that the Bellman equation of optimality is given by

$$
J\left(W_k,x_k,t_k\right)=\max_{\omega_k} E_{t_k}\left(J\left(W_{k+1},x_{k+1},t_{k+1}\right)\right),
$$

and that derived utility for ζ <> 1 may be written as

$$
J\left(W_k,x_k,t_k\right)=\frac{W_k^{1-\zeta}}{1-\zeta}Q\left(x_k,t_k\right),
$$

where, for the ease of exposition, we write W_k and x_k in place of W_{tk} and x_{tk}, respectively. The Bellman equation can thus be rewritten as

$$\begin{aligned} & Q(x_k, t_k) \\ = \; & \max_{\omega_k} E_{t_k}\left\{\left[(1-\omega_k)\exp\left(r_f\frac{\hat{T}}{K}\right)+\omega_k\exp\left(r_f\frac{\hat{T}}{K}+e_{t_k\to t_k+\hat{T}/K}\right)\right]^{1-\zeta}\right\} \\ & \times Q(x_{k+1}, t_{k+1}), \end{aligned}$$

with

$$e_{t_k\to t_k+\hat{T}/K} \equiv e_{t_k+1}+e_{t_k+2}+\ldots+e_{t_k+\hat{T}/K}$$

and $k = 0,\ldots, K-1$.

The usual technique for solving a Bellman equation is to discretize the state space and then use backward induction. Slightly different to Barberis (2000), we take the interval ranging from the maximum and minimum value of the predictive variable over the respective sample period, and discretize this range with 25 equally spaced grid points, x^j_T for $j = 1,\ldots,25$.[9] Suppose that $Q(x_{k+1}, t_{k+1})$ is known for all $x_{k+1} = x^j_{k+1}$. Clearly, this is true in the last period as $Q(x_K, t_K) = 1$. Then we can use equation (.) to obtain $Q(x^j_{K-1}, t_{K-1})$ and, in general, $Q(x^j_k, t_k)$. Backward induction through all K rebalancing points eventually gives $Q(x^j_0, t_0)$, and hence the optimal allocations ω_0.

When estimation risk is ignored, the conditional expectation in equation (.) is taken over the Normal distribution $p(e_{tk\text{->}etk+T/K}, x_{k+1}; \boldsymbol{\theta}, x_k)$, conditional on fixed parameter values. Taking parameter uncertainty into account, however, is much more involved in an intertemporal setting. On the one hand, a first effect is analogous to the one investors face in the buy-and-hold portfolio choice problem. As before, when investors calculate the value function in equation (.) the conditional expectation should be taken over a predictive distribution that incorporates parameter uncertainty. On the other hand, parameter uncertainty may change over time as more data are received. With fresh data, investors may update their posterior distribution of the parameters. If investors anticipate this learning, it may affect their portfolio holdings.

Williams (1977) and Gennotte (1986) analyze the effect of a learning-based hedging demand theoretically. In an intertemporal context, however, additionally incorporating learning is a formidable problem. In this case, the investment opportunity set would no longer be characterized by the predictive variable alone, but also by variables summarizing investors' beliefs about the regression parameters $\boldsymbol{\theta}$. These additional variables would dramatically

---

[9] We refer to Barberis (2000) for a more detailed description of the numerical procedure and its accuracy.

increase the size of the state space, making the dynamic programming problem difficult to solve. To our knowledge, there is no suggestion as to how learning can be incorporated in a discrete-time setting. The notable exception is Xia (2001), but she examines the effects of uncertainty about return predictability in a continuous-time setting.

As in Barberis (2000), we thus simplify the problem and assume that although investors acknowledge that they are uncertain about the regression parameters, they ignore the impact on today's optimal allocation of the fact that their beliefs about those parameters may change over time. As in the buy-and-hold case, these beliefs are summarized by the posterior distribution calculated conditional only on data up until *T*. Consequently, the investment opportunity set is still described by the prevailing predictive variable alone and investors still use equation (.) to calculate the value function. When investors account for parameter uncertainty, the conditional expectation in equation (.) is taken over the predictive distribution $p(e_{tk\text{->}etk+T/K}, x_{k+1}; x_k)$, rather than over $p(e_{tk\text{->}etk+T/K}, x_{k+1}; \boldsymbol{\theta}, x_k)$. Investors thus sample from the predictive distribution by taking a large sample from the posterior distribution $p(\boldsymbol{\theta}; \mathbf{z})$ and then for each set of parameter values drawn, make a draw from $p(e_{tk\text{->}etk+T/K}, x_{k+1}; \boldsymbol{\theta}, x_k)$, a Normal distribution.[10]

6. Empirical Results

Using data from the Swiss and U.S. stock market, this section describes the data, presents empirical evidence of mean reversion, and demonstrates the impact of stock market predictability and estimation risk on the optimal allocation to stocks of myopic and intertemporally optimizing investors.

6.1. The Data

The investment universe consists of monthly observations on log excess stock market returns (including dividends).[11] For the Swiss stock market, the log excess return is the log nominal return less the prevailing one-month Swiss interbank rate, for the U.S. stock market, less the prevailing Treasury bill return.

---

[10] Of course, this approach is possibly too simple, investors should at least recognize that the precision of the parameter estimates will improve over time. However, as suggested by Barberis (2000, footnote 20) and our empirical results following below, we might justify this by thinking of it in the context of a model with time-varying regression parameters. Indeed, in this reasonable case, it would no longer be true that the posterior distribution becomes tighter as more data are received; in fact, it may become even more dispersed.

[11] The data for the stock indices is collected from Datastream Thomson Financial (mnemonics: TOTMKSW(RI,DY), TOTMKUS(RI,DY)).

In deciding which predictive variables to include, attention was given to those variables found important in previous studies of stock market predictability. In particular, we consider the following four predictive variables: (i) dividend-price ratio, log (DPR), (ii) term spread (TERM), (iii) realized stock market volatility, log (VOLA), and (iv) U.S. default risk spread (DEF).

The dividend-price ratio is measured as the sum of dividends paid on the index over the previous year, divided by the current level of the index. The term spread is the difference between the (log) nominal yield on long-term government bonds and the (log) nominal three-month Swiss interbank rate/Treasury bill return. In the same way as Goyal and Santa-Clara (2003), we compute the monthly realized variance of the nominal stock market returns using within-month daily return data for each month as

$$Var_t = \sum_{d=1}^{D_t} r_d^2 + 2\sum_{d=2}^{D_t} r_d r_{d-1},$$

where D_t is the number of days in month t and r_d is the log nominal return on day d. The second term on the right-hand side adjusts for the autocorrelation in daily returns using the approach proposed by French, Schwert, and Stambaugh (1987). Finally, the U.S. default risk spread is formed as the difference in annualized (log) yields of Moody's Baa and Aaa rated bonds.

Motivated by the recent contributions of Ferson, Sarkissian, and Simin (2003, 2004), we also consider the stochastically detrended predictive variables,

$$x_{t-1}^{s.d.} = x_{t-1} - \frac{1}{12}\sum_{\tau=1}^{12} x_{t-1-\tau}.$$

We thus subtract a backward one-year moving average of past values from the prevailing value of the predictive variable. Accordingly, the detrended time series is stationary if changes in the predictive variable are stationary. While this stochastic detrending method has already been used by Campbell (1991) and Hodrick (1992), Ferson, Sarkissian, and Simin (2003, 2004) show that this is the most practically useful insurance against spurious regression bias (and therefore data mining). Since most of the predictive variables are either manifestly non-stationary (realized stock market volatility is the exception), or, if not, their behavior is close enough to unit-root non-stationarity for small-sample statistics to be affected, it is interesting to compare the characteristics of these two data subsets.

The data covers the time period from January 1975 to December 2002 (336 monthly observations) and includes two subsamples of equal length. The first subsample uses data from January 1975 to December 1988, covering the first half of the total time period, the second subsample is based on data from January 1989 to December 2002, covering the second half of the full sample.

Table 2 presents summary statistics of the log excess returns. Compared to the Swiss stock market, the U.S. stock market exhibits a slightly better trade-off between the average excess return and volatility over the full sample period; the reverse is true for the period from 1975 to 1988. With respect to the two subsamples, the average excess return/volatility trade-off is fairly constant for the U.S. stock market, but quite worse for the Swiss stock market over the more recent time period.

[Insert Table 2 about here]

Figure 1 shows the time series of the predictive variables from December 1974 to November 2002, and the corresponding autocorrelation coefficients of both the original and stochastically detrended predictive variables. In general, the predictive variables vary considerably over time. The dividend-price ratio seems to be the only variable that exhibits a (negative) time trend. With the exception of realized stock market volatility, the autocorrelation coefficients are typically large and tend to fall only for a large number of lags, indicating that the behavior of the predictive variables is close enough to unit-root non-stationarity for small-sample statistics to be affected. In contrast, the autocorrelation coefficients of the stochastically detrended predictive variables decrease much faster.[12]

[Insert Figure 1 about here]

Table 3 shows selected parameter estimates for the VAR model in equation (.). The table exhibits the slope coefficients and the adjusted coefficients of determination of the predictive regression, and the correlation coefficients between unexpected excess returns and innovations in the predictive variables.

[Insert Table 3 about here]

[12] Summary statistics and Augmented Dickey-Fuller (ADF) tests of a unit root for the predictive variables, using the procedure discussed in Doldado, Jenkinson, and Sosvilla-Rivero (1990) and Enders (1995), can be obtained from the authors. The consideration of the two subperiods from 1975 to 1988 and 1989 to 2002 do not change the conclusions.

In general, irrespective of whether the predictive variables are stochastically detrended or not, the evidence of stock market predictability is weak, and neither consistent over time nor across the Swiss and U.S. stock market. For example, while the coefficient on the dividend-price ratio of the Swiss stock market is statistically significant over the first subperiod, the same is true for the U.S. stock market over the more recent subperiod. When the dividend-price ratio is stochastically detrended, however, the slope coefficients are not significantly different from zero in all cases.

The risk view of stock market predictability claims that the equity premium rises during an economic slow-down and falls during periods of economic growth, so that expected returns and business conditions move in opposite directions (see, e.g., Fama and French (1989), Chen (1991), Fama (1991), and Ferson and Harvey (1991)). The slope coefficients on the term spread and the U.S. default risk spread are predominantly positive. Therefore, likely reflecting recession states, high term spreads and high U.S. default risk spreads predict high future excess returns (but notice the counterexamples over some of the subsamples). However, when expected excess returns are parameterized either by realized stock market volatility or the dividend-price ratio, the argumentation is more involved. The slope coefficients on realized stock market volatility are predominantly negative. These results are in line with Campbell (1987) and Glosten, Jagannathan, and Runkle (1993), who also find a negative relation between realized stock market volatility and future excess returns, and provide some perspective on the so-called flight-to-quality phenomenon, which refers to investors moving capital from stock markets to government bond markets when the stock markets are more volatile than usual. In contrast, however, French, Schwert, and Stambaugh (1987), Campbell and Hentschel (1992), and, in a recent contribution, Goyal and Santa-Clara (2003) find a significant positive relation.

Finally, the hypothesis that the dividend-price ratio predicts future returns has a long tradition among practitioners and academics alike (see, e.g., Graham and Dodd (1934) and Rozeff (1984)). The intuition of the "efficient markets" version of the hypothesis is that stock prices are low relative to dividends when discount rates and expected returns are high, and vice versa, so that the dividend-price ratio varies with expected returns. Interestingly, thus, in contrast to the bulk of the other predictive variables, the risk-based business-cycle argumentation is not explicitly needed to explain the predictive power of the dividend-price ratio. It can be shown that, without relying on any rational asset pricing model, the dividend-price ratio can only move at all if it predicts expected future returns, if it forecasts expected future dividend growth, or if there is a bubble – if the inverse of the dividend-price ratio is non-stationary and is expected to grow explosively (see, e.g., Campbell and Shiller (1987, 1988a,b), Campbell (1991), Cochrane (1992, 1997, 2001), and Campbell, Lo, and MacKinlay (1997)). When prices are high relative to dividends (or earnings, cash flow, book value or some other divisor), one of three things must be true. First, investors expect dividends to rise in the future. Second, investors expect returns to be low in the future. Future cash flows are discounted at a lower than usual rate, leading to higher prices. Third, investors expect prices to rise forever, giving an adequate return even if there are no dividends. This statement is not a theory, it is an identity. If the dividend-price ratio is low, either dividends must rise, prices must decline, or the inverse of the dividend-price ratio must grow explosively. The open question is, which option holds for the stock market. Confirming his earlier findings and those of Campbell (1991), Cochrane (2001) establishes that, historically, virtually all variation in dividend-price ratios has reflected time-varying expected returns. Similarly, Campbell and Shiller (1998) show that after a deviation from its average value, the dividend-price ratio (and the price-earnings ratio) does a poor job as a forecaster of future dividend growth to the date when the ratio is borne back to its average value again. It is instead the stock price that moves to restore the ratio to its mean value, an offsetting adjustment to dividends (and earnings) seems to be absent. The standard conclusion in the literature thus seems to be threefold: (i) the bulk of the variance of dividend-price ratios must be accounted for by changing forecasts of discount rates, (ii) expected dividend growth is approximately constant, and (iii) high prices reflect low expected returns. Indeed, Table 3 indicates that the slope coefficients on the dividend-price ratio are positive (with the exception of the Swiss stock market over the more recent subperiod).

After all, adjusted R^2s are generally low, about half of them are even negative. When expected excess returns are parameterized either by the dividend-price ratio or the realized stock market volatility (irrespective of whether stochastically detrended or not), the correlation coefficients between unexpected excess returns and innovations in these predictive variables are significantly negative in most cases. The reverse is true for the term spread, but only over the first time period. The correlation coefficients resulting for the U.S. default spread are not significantly different from zero.

6.2. Mean Reversion: Variance Ratio Statistics

Figure 2 shows estimates for the q-period variance ratio statistics, $VR(q)$, that are directly obtained from the raw return data. The dotted lines represent estimates calculated according to equation (.). In order to improve the finite-sample properties, the solid lines use overlapping

q-period log excess returns and correct the bias in the variance estimators as suggested in Lo and MacKinlay (1988, 1989) and Campbell, Lo, and MacKinlay (1997). q varies from one to 60 months.

[Insert Figure 2 about here]

Over the full sample period, the variance ratio statistics are above one for the Swiss stock market, irrespective of the estimation method. For the U.S. stock market, however, the variance ratio statistics indicate mean reversion, at least if the estimates are based on overlapping returns and corrected for the bias in the variance estimators. But the estimates are not significantly different from one. Estimates for the variance ratio statistics are remarkably different in the two subperiods. Over the first subperiod, the variance ratio statistics indicate mean reversion, especially so if they are calculated according to equation (.). In contrast, over the more recent subperiod, the variance ratio statistics indicate strong mean aversion (the estimates for the Swiss stock market are even, at least partially, significantly above one).

Overall, thus, neither the Swiss nor the U.S. stock market exhibits mean reversion. Estimates for the variance ratio statistics are generally not significantly different from one, and the slight tendency of mean reversion over the first subperiod completely vanishes over the more recent time period. These results are thus hard to reconcile with time-varying investment opportunities.

Figures 3, 4, and 5 show estimates for the implied variance ratio statistics and the implied coefficients of determination, obtained from the predictive regression model in equation (.). Figures 3 and 4 are based on the dividend-price ratio from the Swiss and U.S. stock market, respectively. Figure 5 shows the implied variance ratio statistics for the other predictive variables from both stock markets.

[Insert Figure 3 about here]

The solid lines in Figure 3 correspond to the multivariate statistics, the dotted lines to the respective univariate statistics. The black and gray lines represent the cases where expected excess returns are parameterized by the dividend-price ratio and the stochastically detrended dividend-price ratio, respectively.

The general pattern is quite different in the two subperiods. From 1975 to 1988, the implied univariate variance ratio statistics decline steadily and approach a limit below 0.3. Thus, the corresponding implied autocorrelation coefficients are negative, indicating strong mean reversion. As discussed above, the multivariate variance ratio statistics are even lower. Recently, however, Figure 3b shows that the variance ratio statistics are approximately one across all time horizons, indicating no stock market predictability by the dividend-price ratio and thus constant i.i.d. investment opportunities (see also Table 3). In the case of the stochastically detrended dividend-price ratio, the implied variance ratio statistics even suggest mean aversion instead of mean reversion. Figure 3 also shows the implied R^2s for time horizons up to 10 years. In the first subperiod, the multivariate R^2 statistics rise steeply from their initial values at the one-month time horizon to a peak at about 36% at a time horizon of about 4 to 5 years. This is a result of the fact that the dividend-price ratio is highly persistent. On the other hand, the univariate Fama-French R^2s peak at about 16%, indicating that the conditional R^2 statistics are about twice as high as the unconditional Fama-French R^2s. As emphasized in Campbell (1991), this is an indication of the benefits obtainable from a multivariate rather than a univariate approach to stock market returns. In the recent subsample, however, the dividend-price ratio does not predict excess returns, and the implied R^2 statistics do not rise with the forecast horizon. Given the low persistence of the expected excess return when parameterized by the stochastically detrended dividend-price ratio, the small implied R^2 statistics in both subperiods are just consistent.

[Insert Figure 4 about here]

Figure 4 extends the results to the U.S. stock market. In contrast to the Swiss stock market, the general pattern is similar in the two subperiods. The dividend-price ratio of the U.S. stock market is a more robust predictor of future excess returns than the dividend-price ratio of the Swiss stock market. We thus see strong mean reversion and increasing multivariate R^2 statistics in both subperiods. While the R^2s remain rather small, the variance ratio statistics even indicate some moderate mean reversion when the dividend-price ratio is stochastically detrended.

[Insert Figure 5 about here]

Figure 5 shows the implied variance ratio statistics when expected excess returns are parameterized by the term spread, realized stock market volatility, and the U.S. default risk spread. In comparison to the dividend-price ratio, the general pattern is quite different. Instead of mean reversion, excess returns rather exhibit mean aversion, i.e., univariate and multivariate variance ratio statistics often exceed one, indicating implied autocorrelation coefficients that are predominantly positive. This is true for both subperiods.

Overall, thus, the evidence of mean reversion in excess returns is rather low based on variance ratio statistics. The exception is when expected excess returns are parameterized by the dividend-price ratio, especially for the U.S. stock market. For the remaining predictive variables and the stochastically detrended variables (including the dividend-price ratio), there is at best some evidence of mean aversion, but hardly any reliable evidence of mean reversion.

The following section shows what these results mean for the optimal allocation to stocks.

6.3. Buy-and-Hold Investment Strategies

This section demonstrates the impact of stock market predictability and estimation risk on the optimal portfolio choice of buy-and-hold investors. Optimal allocations to stocks are obtained using the four different choices of the distribution of future excess returns introduced above. These distributions differ in whether they take stock market predictability and estimation risk into account. As in Barberis (2000), we do these calculations for planning horizons T ranging from one year to ten years in one-year increments, and, in the case of return predictability, for several initial values of the predictive variables. We start with the optimal allocations to stocks of investors who follow a buy-and-hold strategy, use an i.i.d. constant investment opportunities model for the log excess return, and have power utility over terminal wealth. Only then we move to the more general case and additionally allow for time-varying investment opportunities. Throughout, the continuously compounded real monthly Swiss interbank rate, r_f, is fixed at the arithmetic average over the sample period under consideration.

Constant Investment Opportunities

Figure 6 shows the optimal percentage 100ω percent allocated to the stock market index, plotted against the planning horizon in years. The solid lines correspond to the case where investors ignore estimation risk, the dotted lines where they account for it. Parameter estimates are obtained using data from 1975 to 1988 and 1989 to 2002, respectively. Black lines represent the Swiss stock market, gray lines the U.S. stock market.

If parameter uncertainty is ignored, buy-and-hold investors maximize expected utility using the return distribution conditional on the parameter values fixed at their estimated values. In the case of i.i.d. excess returns, we thus need estimates of the realized mean excess return, μ, and the volatility of the realized risk premium, σ^2_ξ. These estimates are shown in Table 2.

[Insert Figure 6 about here]

The solid lines in Figure 6 are horizontal in all cases. Thus, investors who ignore the uncertainty about the mean and volatility of the risk premium allocate the same amount to stocks, regardless of their planning horizon. However, when estimation risk is explicitly incorporated into the investor's decision-making framework, Figure 6 demonstrates that the allocation to stocks falls as the planning horizon increases. Estimation risk thus introduces a negative horizon effect when investment opportunities are constant. The magnitude of the negative horizon effect induced by parameter uncertainty is quite substantial. At the ten-year horizon, with $\zeta = 5$ and from 1975 to 1988, for example, the optimal allocation to stocks is $\omega = 71\%$ (61%) for Swiss (U.S.) investors who ignore estimation risk, and only $\omega = 45\%$ (38%) for investors who account for it.

These results can be explained as follows. When estimation risk is ignored, both the mean and the variance grow linearly with the planning horizon. However, when parameter uncertainty is taken into account, the distribution of long-horizon excess returns faces an additional source of uncertainty. Estimation risk thus makes the variance of the distribution of cumulative excess returns increase faster than linearly with the planning horizon. Indeed, in the presence of estimation risk and from the viewpoint of the investors, returns are no longer i.i.d., but rather positively serially correlated. Positive serial correlation makes stocks look riskier in the long run, and investors therefore reduce the amount they allocate to equities.

Finally, Figure 6 also shows that the impact of different values of relative risk aversion should not be neglected. If risk aversion is $\zeta = 10$ instead of $\zeta = 5$, for example, the optimal allocation to stocks approximately reduces by half, and the effect of estimation risk is more pronounced (expressed as a percentage).

Time-Varying Investment Opportunities

Now that the impact of estimation risk alone has been illustrated, stock market predictability can be introduced as well. We begin with the dividend-price ratio and discuss these results in detail. Only then we continue with the analysis of the other predictive variables.

Figures 7 shows optimal allocations to stocks for the dividend-price ratio as predictive variable and Swiss stock market data from 1975 to 1988 and 1989 to 2002, respectively.[13] Recall that the distribution of future excess returns depends on the value of the dividend-price ratio at the beginning of the planning horizon. The initial value of the dividend-price ratio enters equation (.) through $\mathbf{z}_T$. Accordingly, the five lines within each graph of Figure 7 correspond to different initial values of the dividend-price ratio over the respective time period.[14]

[Insert Figure 7 about here]

When estimation risk is ignored, buy-and-hold investors maximize expected utility using the return distribution conditional on the parameter values fixed at their estimated values. Given the above VAR framework, they thus need estimates for **a**, **B** and **Σ**. Estimates of these are (partly) shown in Table 2. The table indicates that, over the first subperiod, the slope coefficient on the dividend-price ratio is significantly positive at the 5% significance level, and the adjusted R^2 is 2.14%. The correlation coefficient between unexpected excess returns and innovations in the dividend-price ratio is strongly negative, estimated at -0.9574.

Equation (.) shows that the (approximate) optimal allocation to stocks depends on the dynamics of both the conditional mean and the conditional variance. With respect to the implied dynamics of the conditional variance, Figure 3 shows that the conditional variance ratio statistics are below one. Consequently, when investors take the predictive power of the dividend-price ratio into account, conditional variances grow slower than linearly with the planning horizon, lowering the perceived risk of stocks in the long run. As a result, the optimal allocation to stocks rises with the planning horizon for all initial values of the dividend-price ratio.

The intuition behind this effect is the following. Since $\sigma_{\xi\eta} < 0$ and $\beta > 0$, an unexpected fall of the dividend-price ratio is likely to be accompanied by a contemporaneous positive shock to realized returns. Given that the dividend-price ratio is now lower, however, returns are

---

[13] The analysis of the full sample period from 1975 to 2002 does not change the conclusions.

[14] We take the interval ranging from the maximum and minimum value of the dividend-price ratio (and, in general, the predictive variable under consideration) over the respective sample period, and discretize this range with 25 equally spaced grid points, x^j_T for $j = 1,\ldots,25$. In particular, the five values we use are x^6_T, x^{10}_T, x^{13}_T, x^{16}_T, and x^{20}_T.

expected to be lower in the future, too. This pattern generates a component of negative serial correlation in realized returns. With a growing planning horizon, this perceived negative serial correlation slows the evolution of the conditional variance of cumulative returns. Indeed, the idea that time variation in expected returns may induce mean reversion in realized returns has a strong economic intuition. If expected returns suddenly increase, it is reasonable to assume that realized returns suffer a contemporaneous negative shock: after all, the discount rate for discounting future cash flows has suddenly increased (see, e.g., Campbell, Lo, and MacKinlay (1997)). Negative shocks to current realized returns, which are followed by the higher returns predicted for the future, are the source of this perceived mean reversion.

Because the dividend-price ratio also affects the conditional mean of the distribution of future excess returns, conditioning on the dividend-price ratio has a second implication for the optimal allocation to stocks. Consistent with the fact that the equity risk premium increases with the dividend-price ratio ($\beta > 0$), Figure 7 indicates that the optimal allocation to stocks increases with the dividend-price ratio, too. The intuition is simple. When the risk premium increases, stocks become more attractive (higher expected excess return for the same level of risk), and, consequently, investors allocates more wealth to stocks. The extend to which investors try to time the market decreases with risk aversion, however.

This pattern significantly changes when investors take estimation risk into account. The graphs on the right-hand side of Figure 7 show that for low initial values of the dividend-price ratio, the optimal allocation to stocks rises with the planning horizon. For high initial values of the dividend-price ratio, however, the allocation to equities falls in the long run. Obviously, thus, when investors take parameter uncertainty into account, the allocation lines converge, mirroring the dynamics of the conditional mean. If investors ignore estimation risk, return predictability makes stocks look less risky in the long run. On the other hand, incorporating estimation risk makes them look more risky. These two effects go in the opposite directions. Portfolio decisions are therefore not necessarily monotonic as a function of the planning horizon.[15] Overall, the resulting optimal allocation to stocks in the long run is less sensitive to the initial value of the dividend-price ratio than the allocation to stocks in the short run, and

[15] Barberis (2000), partly based on Kandel and Stambaugh (1996), highlight another surprising fact. For a given planning horizon and risk aversion level, the optimal allocation to stocks is not necessarily increasing in the initial value of the dividend-price ratio. If the initial value of the dividend-price ratio is $x^{20}{}_T$ rather than $x^{16}{}_T$ (note that $x^{20}{}_T > x^{16}{}_T$), for example, the predictive distribution has a higher posterior mean. This should lead to a higher allocation to equities. Because, in addition, the variance of the predictive distribution is not sensitive to the initial value of the dividend-price ratio, this cannot explain the non-monotonicity result in Figure 7. In fact it is the third moment of the predictive distribution, skewness, that is important here. While for low initial values of the dividend-price ratio, incorporating estimation risk generates positive skewness in the predictive distribution, it generates negative skewness for high initial values of the dividend-price ratio. In the case of power utility, this negative skewness makes stock less attractive, and makes optimal stock holdings non-monotonic in the initial value of the dividend-price ratio.

much less sensitive than the allocation to stocks of investors with a long planning horizon who ignore estimation risk. When investors take estimation risk into account, they acknowledge that they are not only uncertain about the average equity premium (as in the i.i.d. case), but also about the true predictive power of the dividend-price ratio and whether the dividend-price ratio really does induce conditional negative serial correlations in realized returns. Consequently, with respect to their asset allocation, investors are more cautious about stocks and allocate less to them.[16]

The true predictive power of the dividend-price ratio may change over time. Thus, investors may prefer to estimate the relationship over the more recent subsample from 1989 to 2002. Indeed, Table 2 indicates that any evidence of stock market predictability by the dividend-price ratio seems to have disappeared over the more recent subsample. The coefficient on the dividend-price ratio is far from being statistically significant at any reasonable significance level, and, even worse, the sign is wrong ($\beta < 0$).

[Insert Figure 8 about here]

Figure 8 repeats the calculation of Figure 7 for the case where investors use the more recent subsample in making their portfolio decisions. Not surprisingly, the resulting allocations to stocks show neither a positive horizon effect nor are they sensitive to the initial value of the dividend-price ratio. Conversely, if estimation risk is incorporated, the recommended allocation to equities decreases with the planning horizon – in the same way as discussed in the case of i.i.d. excess returns.

Figure 9 shows optimal allocations to equities over the whole state space, i.e., for all initial values of the dividend-price ratio from x^{1}_{T} to x^{25}_{T}. When investors account for estimation risk, the non-monotonicity and lower level of the allocations to stocks is only obvious. For the period from 1989 to 2002, there is no positive horizon effect at all – the predictive power of the dividend-price ratio has completely disappeared.

[Insert Figure 9 about here]

---

[16] As discussed in Barberis (2000, footnote 16), when investors take estimation risk into account, they acknowledge both that the predictive power of the dividend-price ratio may be weaker than the parameter estimates suggest or that it may be stronger. In the first case they would allocate less to stocks at long horizons, in the second case, however, they would increase their equity holdings. In sum, thus, because investors are risk averse and hence dislike the mean-preserving spread that accounting for estimation risk adds to the distribution of future excess returns, they invest less at long horizons.

Figure 10 confirms the findings in Table 2 that the predictive power of the dividend-price ratio is more robust for the U.S. stock market than for the Swiss stock market. U.S. investors who ignore estimation risk increase their allocation to stocks across their planning and investment horizon. However, as in the Swiss case, if parameter uncertainty is taken into account, the negative impact of stock market predictability on the growth of long-term variances (and hence the positive horizon effect) is substantially reduced. Long-term investors should not invest more in stocks than short-term investors, the allocation lines merely mirror the dynamics of the conditional mean implied by the VAR.

[Insert Figure 10 about here]

Table 2 indicates that the predictive power of the stochastically detrended dividend-price ratio is marginal compared to the more persistent original time series – both for the Swiss and the U.S. stock market. In addition, Figures 3 and 4 indicate that the implied variance ratio statistics show much less evidence of mean reversion. Altogether, thus, the findings illustrated in Figure 11 suggest at best some short-term market timing due to the dynamics of the stochastically detrended dividend-price ratio, but no positive horizon effect anymore. When expected future excess returns are parameterized by the much less persistent stochastically detrended dividend-price ratio, the optimal allocation to stocks does not increase with the length of the planning horizon, irrespective of whether estimation risk is taken into account or not.

[Insert Figure 11 about here]

It is now important to discuss whether the above results for the dividend-price ratio generalize to the other predictive variables. In what follows, we repeat the same calculations for those variables as well and demonstrate the resulting optimal allocations to equities. To save space, we restrict the analysis to the full sample from 1975 to 2002. Figures 12 and 13 show the results for the Swiss and the U.S. stock market, respectively.

[Insert Figure 12 about here]

When time-varying investment opportunities are parameterized by either the term spread, realized stock market volatility, or the U.S. default risk spread, Figures 12 and 13 show that there is no positive horizon effect. Optimal stock holdings of long-term investors are generally not higher than those of short-term investors. In contrast, the additional source of uncertainty due to estimation risk rather make stocks look riskier in the long run, and investors reduce both the extend to which they try to time the market and the amount they allocate to equities in the long run.

[Insert Figure 13 about here]

To understand the absence of a positive horizon effect, note that (i) the evidence of stock market predictability is often very weak (and actually non-existent in most cases), (ii) the correlation coefficient between unexpected excess returns and innovations in the predictive variables is often not significantly different from zero, and, (iii), expected excess returns increase (decrease) with the predictive variable while the correlation between innovations in excess returns and innovations in the predictive variable is positive (negative), thereby not leading to a diminishing perceived variance of excess returns across the planning horizon and a higher allocation to stocks.

Overall, thus, a positive horizon effect is only evident if investors believe that expected excess returns are parameterized by the dividend-price ratio. Even in this case, however, the evidence is rather weak. Parameter instability (for the Swiss stock market, for example, the positive horizon effect is limited to the time period from 1975 to 1988), estimation risk, and stochastically detrending the dividend-price ratio (which reduces the evidence of return predictability and makes the variable less persistent) may still induce some short-term market timing strategies, but no longer a positive horizon effect due to decreasing conditional variances of cumulative excess returns. The same is true for the other predictive variables. The dynamic interaction of realized excess returns and expected excess returns parameterized by either the term spread, realized stock market volatility, or the U.S. default risk spread suggest that long-term allocations to stocks behave as if investment opportunities were constant.

6.4. Intertemporal Investment Strategies

Using the same data from the Swiss and U.S. stock market as above, this section examines the size of the intertemporal hedging demand.

As in Barberis (2000), the planning horizons range from one year to ten years in one-year increments, and investors are allowed to adjust their portfolio annually using the new information at the end of each year. Again, we begin with the dividend-price ratio as the single predictive variable and discuss these results in detail. The analysis of the other predictive follows, but does not reveal any new evidence.

[Insert Figure 14 about here]

As in Kim and Omberg (1996) and Campbell and Viceira (2002), Figure 14 shows optimal allocations to stocks for the Swiss stock market and for all initial values of the dividend-price ratio from x^1_T to x^{25}_T. The graphs on the left-hand (right-hand) side of Figure 14 ignore (account for) parameter uncertainty, the top row graphs use data from 1975 to 1988, the lower two use data from 1989 to 2002, respectively.

Using data from 1975 to 1988, Figure 14 shows that the optimal allocation to stocks rises with the planning horizon. For example, when the dividend-price ratio is 2.60%, the optimal allocation to stocks invested over the following year is 23%, 60%, and 75% for a planning horizon of one, five, and ten years, respectively.[17] Although these results appear to be the same as in the buy-and-hold case (which is only true for a planning horizon of one year, however), they must be explained differently. In the multiperiod framework, any observed increase in the allocation to equities across planning horizons is due to the intertemporal hedging demand. According to equation (.), expected excess returns are governed by the dividend-price ratio. As the dividend-price ratio varies over time, investment opportunities faced by the investors change accordingly. Merton (1973) shows that investors who are more risk averse than log-utility maximizers optimally hedge these changes by investing in a way that gives them higher wealth precisely when investment opportunities are unattractive, i.e., when expected excess returns are low. Since shocks to expected excess returns are negatively correlated with shocks to realized returns, holding more in stocks seems to be an ideal way of hedging against movements in expected excess returns. The financial gain (loss) from over-investing partially offsets the expected utility loss (gain) associated with the drop (rise) in the dividend-price ratio.

However, as above, incorporating estimation risk has an adverse effect on the positive intertemporal hedging demand, but the impact is less distinct than for the long-term buy-and-hold portfolio choice problem. Continuing the above example, when the dividend-price ratio

[17] The intertemporal hedging demand seems to converge for long planning horizons. The difference between the resulting hedging demand for a five-year (dotted lines) and a one-year planning horizon is larger than the difference between a ten-year and a five-year horizon hedging demand. In addition, investors who follow a dynamic rebalancing strategy do hardly time the market more aggressively than myopic buy-and-hold investors. The allocation lines in Figure 14 are largely parallel to one another.

is 2.60%, the optimal allocation to stocks invested over the following year is 21%, 44%, and 57% for a planning horizon of one, five, and ten years, respectively. In general, if the uncertainty in the predictive power of the dividend-price ratio is taken into account, the graphs on the right-hand side of Figure 14 demonstrate that the intertemporal hedging demand is somewhat smaller and that the allocation lines are flatter across the state space. If investors acknowledge estimation risk, they become more skeptical whether investment opportunities do change over time, and consequently reduce the magnitude of the intertemporal hedging demand. In addition, the presence of parameter uncertainty decrease the sensitivity of the optimal allocation to stocks to the initial value of the dividend-price ratio, especially for low and high values of the dividend-price ratio. Changes in portfolio compositions thus occur more gradually over time. Therefore, the often cited results of Brennan, Schwartz, and Lagnado (1997) and Campbell and Viceira (1999), for example, may need to be interpreted with caution. Because their dynamic asset allocation strategies ignore estimation risk, the recommended allocations to stocks are probably too high and too sensitive to their variables parameterizing expected excess returns.

Recently, from 1989 to 2002, any evidence of stock market predictability by the dividend-price ratio seems to have disappeared. Correspondingly, the lower graphs of Figure 14 show that the resulting allocations to stocks neither show a positive intertemporal hedging demand nor are they sensitive to the initial value of the dividend-price ratio.

[Insert Figure 15 about here]

Figure 15 extends the above analysis to the U.S. stock market. Because the dividend-price ratio is a more robust predictor for the U.S. stock market than for the Swiss stock market, intertemporal hedging demands arise over both time periods. Both the size and the pattern of the intertemporal hedging demand is very similar to the Swiss stock market over the first sample period.

Finally, the analysis of intertemporal hedging demands due to the stochastically detrended dividend-price ratios and the other predictive variables introduced above does not reveal any evidence of a positive horizon effect. As can be suspected from the earlier results for the buy-and-hold portfolio choice problem, there is no intertemporal hedging demand in these cases. The optimal allocations to stocks for the first year are the same across all planning horizons. Therefore, the optimal allocations to stocks are already plotted in Figures 11, 12, and 13, respectively (for the planning horizon of one year).

In sum, mirroring the earlier findings for the buy-and-hold portfolio choice problem, intertemporal hedging demands are evident solely if investors believe that expected excess returns are parameterized by the dividend-price ratio. If the time variation in investment opportunities is parameterized by the stochastically detrended dividend-price ratio or by either the term spread, realized stock market volatility, or the U.S. default risk spread (stochastically detrended or not), neither the Swiss nor the U.S. stock market seems to provide a promising intertemporal hedge to future changes in investment opportunities. The resulting long-term allocations to stocks behave as if investment opportunities were constant. Motivated by return predictability, thus, long-term investors should not hold more stocks and should not time the market more aggressively than myopic short-term investors.

7. Conclusion

Recommendations that long-term investors should tilt their portfolio toward stocks are commonplace, but often lack of a consistent clarification of the underlying reasons. Not surprisingly, thus, the question as to whether or not there is any rational reason for long-term investors to allocate a higher fraction of their wealth to the risky stock market created much controversy in the academic literature and the investment community alike. One promising reason that has become prominent among financial economists only recently, is the mounting empirical evidence of stock market predictability. Indeed, the time variation and predictability of stock market returns is meanwhile so widely accepted in the profession that it has renewed the interest in portfolio choice problems in order to estimate Merton’s (1973) intertemporal hedging demand. Most of these dynamic asset allocation models indicate that long-term investors should allocate substantially more to stocks than short-term investors, i.e., that the horizon effect is positive and economically large.

However, certain aspects of the empirical research on stock market predictability remain controversial and should be regarded with caution. Notably, a number of recent contributions suggest that standard statistical inference procedures overreject the null hypothesis of no predictability, so that any evidence of return predictability may have more to do with poorly behaved test statistics than with stock market predictability. In addition, studies such as Bossaerts and Hillion (1999), Goyal and Welch (2003a,b), and Schwert (2003) conclude that the out-of-sample predictive power of the dividend-price ratio and other predictive variables is abysmal and, in the words of Schwert (2003), disastrous, suggesting that the evidence of stock market predictability and its implications for the optimal allocation to the stock market is much less transparent than common beliefs may have suggested.

Our results support these critical considerations. From our empirical point of view, we conclude that stock market predictability and persistent movements in expected excess returns should not be interpreted as a reliable driving force for a positive horizon effect and the corresponding market timing strategies, least of all when estimation risk is taken into account. Although there is some (weak) evidence of stock market predictability, it is not very consistent across the Swiss and the U.S. stock market, depends on the particular specification of the information set which investors use to predict equity returns, and the time period under consideration (indicating parameter instability).[18] Notably, despite the fact that the dividend-price ratio is a more robust predictor for the U.S. stock market, our conclusions for the Swiss stock market apply to the U.S. stock market as well. Thus, neither the Swiss nor the U.S. stock market seems to be less risky in the long run than in the short run. Long-term investors should therefore not hold more stocks and should not time the market more aggressively than myopic short-term investors.

Overall, thus, the standard lesson that stock market predictability and persistent movements in expected excess returns give rise to dramatic portfolio implications for long-term investors may be somewhat premature, at least incomplete. Our results suggest that the real existence of horizon effects due to stock market predictability is a very open question. In the same way, it seems very questionable whether time-varying investment opportunities can be sufficiently predicted by simple regression techniques and whether the respective results should be considered as a serious input for dynamic asset allocation strategies.

This conclusion is further supported by the following discussion borrowed from Kandel and Stambaugh (1996) and Campbell and Viceira (2002). They note that based on a partial equilibrium framework (which is standard in the dynamic asset allocation literature), all investors are forced to buy and sell assets at the very same time. With a constant supply of stocks, however, this cannot be consistent with a general equilibrium model. One possible resolution of this difficulty is that the representative investor has different preferences from those commonly assumed (e.g., power utility), perhaps the habit-formation preferences of Campbell and Cochrane (1999). Models with time-varying risk aversion may solve the complementary problem of finding preferences that make a representative agent content to buy and hold the market in the face of a time-varying equity premium, but under this

[18] Additionally incorporating model risk would only corroborate our conclusions (see, e.g., Avramov (2002) and Rey (2005)).

interpretation, any results on market timing should be used only by investors with constant risk aversion, who cannot be typical of the market as a whole.

This point is similarly stressed in Cochrane (1999b). He emphasizes that the average investor must hold the market. Market timing can only work if it involves buying stocks when anybody else wants them and selling them when everybody else wants them. Portfolio advice to follow these strategies must fall on deaf ears for the average investor, and a large class of investors must want to head in exactly the other direction. So, at a basic level, return predictability would (and should) not affect the portfolio decisions of average investors! Put it differently, all arguments that return predictability is real are inconsistent with market timing strategies. If stock market predictability is real, i.e., an equilibrium time-varying reward for holding risk (Bekaert's (2001) risk view), then the average investor knows about it but does not invest because the extra risk exactly counteracts the extra average return. If more than a minuscule fraction of investors are not already at their best allocations, then the market has not reached equilibrium and the premiums will change. If the risk is irrational, then by the time the average investor knows about it, it is gone. But expected excess returns corresponding to an irrational risk premium would have the strongest portfolio implications – everyone should do it – but the shortest lifetime. If the average return comes from a behavioral aversion to risk (the behavioral view), it is just as inconsistent with market timing portfolio advice as if it were real. Investors cannot all be less behavioral than average, just as investors cannot all be less exposed to a risk than average. Finally, if stock market predictability is actually a statistical fluke due to poor statistical inference (the statistical view), any advice to time the market is questionable anyway.

The world of investment opportunities is probably much simpler than indicated by the stock market predictability and dynamic asset allocation literature. For investment purposes, the natural alternative to return predictability (and thus time-varying investment opportunities) is that investment opportunities are constant. After all, it is already difficult enough to estimate and justify the unconditional equity premium (just recall the equity premium puzzle of Mehra and Prescott (1985), for example) and the degree of the investors' risk tolerance. Admittedly, compared with common sense and much industry practice, this seems to be radical advice. Still, this does not mean that horizon effects are completely absent. Strategically important variables such as human capital (labor income), endogenous labor supply, and, in general, life-cycle considerations (asset- and liability management) may give rise to intertemporal hedging demands and horizon effects as well, even if investment opportunities themselves are constant (see, e.g., Samuelson (1991, 1994), Bodie, Merton, and Samuelson (1992), Viceira

(2001), Campbell and Viceira (2002), and Drobetz and Rey (2004)). In any case, a comprehensive consideration of other strategic variables seems to be much more important to the average investor than quibbling over sophisticated myopic or intertemporal market timing strategies based on the rather weak if not non-existent evidence of stock market predictability.

References

Avramov, D. (2002). "Stock Return Predictability and Model Uncertainty", Journal of Financial Economics 64, 423-458.

Baker, M. and J. Wurgler (2000). "The Equity Share in New Issues and Aggregate Stock Returns", Journal of Finance 55, 2219-2257.

Barberis, N. (2000). "Investing for the Long Run when Returns are Predictable", Journal of Finance 55, 225-264.

Bekaert, G. (2001). "Editor's Foreword to the Special Issue: 'On the Predictability of Asset Returns'", Journal of Empirical Finance 8, 451-457.

Bodie, Z. (1995). "On the Risk of Stocks in the Long Run", Financial Analysts Journal, May/June, 18-22.

Bodie, Z., R.C. Merton, and W. Samuelson (1992). "Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model", Journal of Economic Dynamics and Control 16, 427-449.

Bossaerts, P. and P. Hillion (1999). "Implementing Statistical Criteria to Select Return Forecasting Models: What Do We Learn?", Review of Financial Studies 12, 405-428.

Boudoukh, J., M. Richardson, and T. Smith (1993). "Is the Ex-Ante Risk Premium Always Positive?", Journal of Financial Economics 34, 387-408.

Brandt, M.W. (2004). "Portfolio Choice Problems", to appear in Y. Ait-Sahalia and L. P. Hansen, eds., "Handbook of Financial Econometrics", Elsevier Science.

Brennan, M.J., E.S. Schwartz, and R. Lagnado (1997). "Strategic Asset Allocation", Journal of Economic Dynamics and Control 21, 1377-1403.

Campbell, J.Y. (1987). "Stock Returns and the Term Structure", Journal of Financial Economics 18, 373-399.

Campbell, J.Y. (1990). "Measuring the Persistence of Expected Returns", American Economic Review 80, 43-47.

Campbell, J.Y. (1991). "A Variance Decomposition for Stock Returns", Economic Journal 101, 157-179.

Campbell, J.Y. and J.H. Cochrane (1999). "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior", Journal of Political Economy 107, 205-251.

Campbell, J.Y. and L. Hentschel (1992). "No News is Good News: An Asymmetric Model of Changing Volatility in Stock Returns", Journal of Financial Economics 31, 281-318.

Campbell, J.Y. and R. J. Shiller (1988a). "Stock Prices, Earnings, and Expected Dividends", Journal of Finance 43, 661-676.

Campbell, J.Y. and R. J. Shiller (1988b). "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 (1998). "Valuation Ratios and the Long-Run Stock Market Outlook", Journal of Portfolio Management 24, 11-26 (updated 2001; NBER, Working Paper 8221).

Campbell, J.Y. and S. B. Thompson (2005). "Predicting the Equity Premium Out of Sample: Can Anything Beta the Historical Average?", Working Paper (Harvard University), April 2005.

Campbell, J.Y. and L.M. Viceira (1999). "Consumption and Portfolio Decisions when Expected Returns are Time-Varying", Quarterly Journal of Economics 144, 433-495.

Campbell, J.Y. and L.M. Viceira (2002). "Strategic Asset Allocation: Portfolio Choice for Long-Term Investors", Oxford University Press.

Campbell, J.Y., A. Lo, and A. MacKinlay (1997). "The Econometrics of Financial Markets", Princeton University Press.

Canner, N., N.G. Mankiw, and D.N. Weil (1997). "An Asset Allocation Puzzle", American Economic Review 87, 181-191.

Chen, N. (1991). "Financial Investment Opportunities and the Macroeconomy", Journal of Finance 46, 529-554.

Cochrane, J.H. (1992). "Explaining the Variance of Price-Dividend Ratios", Review of Financial Studies 5, 243-280.

Cochrane, J.H. (1997). "Where is the Market Going? Uncertain Facts and Novel Theories", Economic Perspectives 21, 3-37.

Cochrane, J.H. (1999a). "New Facts in Finance", Economic Perspectives 23, 36-58.

Cochrane, J.H. (1999b). "Portfolio Advice for a Multifactor World", Economic Perspectives 23, 59-78.

Cochrane, J.H. (2001). "Asset Pricing", Princeton University Press.

Doldado, J., T. Jenkinson, and S. Sosvilla-Rivero (1990). "Cointegration and Unit Roots", Journal of Economic Surveys 4, 249-273.

Drobetz, W. and D. Rey (2004). "Neue Ansätze der Risikoverteilung – Mosaiksteine einer altergerechten Vermögensaufteilung", Neue Zürcher Zeitung, August 30, p. 22.

Enders, W. (1995). "Applied Econometric Time Series", Wiley.

Epstein, L. and S. Zin (1991). "Substitution, Risk Aversion, and Temporal Behavior of Consumption and Asset Returns II: An Empirical Analysis", Journal of Political Economy 99, 263-286.

Fama, E.F. (1991). "Efficient Capital Markets: II", Journal of Finance 46, 1575-1617.

Fama, E.F. and K.R. French (1988a). "Dividend Yields and Expected Stock Returns", Journal of Financial Economics 22, 3-25.

Fama, E.F. and K.R. French (1988b). "Permanent and Temporary Components of Stock Prices", Journal of Political Economy 96, 246-273.

Fama, E.F. and K.R. French (1989). "Business Conditions and Expected Returns on Stocks and Bonds", Journal of Financial Economics 25, 23-49.

Fama, E.F. and G.W. Schwert (1977). "Asset Returns and Inflation", Journal of Financial Economics 5, 115-146.

Ferson, W.E. and C.R. Harvey (1991). "The Variation of Economic Risk Premiums", Journal of Political Economy 99, 385-415.

Ferson, W.E. and C.R. Harvey (1993). "The Risk and Predictability of International Equity Returns", Review of Financial Studies 6, 527-566.

Ferson, W.E., S. Sarkissian, and T. Simin (2003). "Spurious Regression in Financial Economics?", Journal of Finance 58, 1393-1413.

Ferson, W.E., S. Sarkissian, and T. Simin (2004). "Is Stock Return Predictability Spurious?", Journal of Investment Management 3, 1-10.

Foster, F.D., T. Smith, and R.E. Whaley (1997). "Assessing Goodness-of-Fit of Asset Pricing Models: The Distribution of the Maximal R^2", Journal of Finance 52, 591-607.

French, K.R., G.W. Schwert, and R.F. Stambaugh (1987). "Expected Stock Returns and Volatility", Journal of Financial Economics 19, 3-29.

Gennotte, G. (1986). "Optimal Portfolio Choice Under Incomplete Information", Journal of Finance 41, 733-746.

Glosten, L.R., R. Jagannathan, and D.E. 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.

Goyal, A. and P. Santa-Clara (2003). "Idiosyncratic Risk Matters!", Journal of Finance 58, 975-1007.

Goyal, A. and I. Welch (2003a). "Predicting the Equity Premium with Dividend Ratios", Management Science 49, 639-654.

Goyal, A. and I. Welch (2003b). "A Note on 'Predicting Returns with Financial Ratios'", Working Paper (Yale International Center for Finance), December 2003.

Graham, B. and D. L. Dodd (1934). "Security Analysis", McGraw-Hill.

Granger, C.W.J. and P. Newbold (1974). "Spurious Regression in Economics", Journal of Econometrics 4, 111-120.

Hamilton, J.D. (1994). "Time Series Analysis", Princeton University Press.

Hodrick, R.J. (1992). "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement", Review of Financial Studies 5, 357-386.

Jarque, C.M. and A.K. Bera (1980). "Efficient Tests for Normality, Homoskedasticity, and Serial Independence of Regression Residuals", Economic Letters 6, 255-259.

Kandel, S. and R.F. 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, 357-390.

Kim, T.S. and E. Omberg (1996). "Dynamic Nonmyopic Portfolio Behavior", Review of Financial Studies 9, 141-161.

Kothari, S.P. and J. Shanken (1997). "Book-to-Market, Dividend Yield, and Expected Market Returns: A Time-Series Analysis", Journal of Financial Economics 44, 169-203.

Lamont, O. (1998). "Earnings and Expected Returns", Journal of Finance 53, 1563-1587.

Lee, C., J. Myers, and B. Swaminathan (1999). "What is the Intrinsic Value of the Dow?", Journal of Finance 54, 1693-1742.

Lettau, M. and S. Ludvigson (2001). "Consumption, Aggregate Wealth, and Expected Stock Returns", Journal of Finance 56, 815-849.

Lo, A.W. and A.C. MacKinlay (1988). "Stock Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test", Review of Financial Studies 1, 41-66.

Lo, A.W. and A.C. MacKinlay (1989). "The Size and Power of the Variance Ratio Test in Finite Samples: A Monte Carlo Investigation", Journal of Econometrics 40, 203-238.

Lo, A.W. and A.C. MacKinlay (1990). "Data-Snooping Biases in Tests of Financial Asset Pricing Models", Review of Financial Studies 3, 431-467.

Markowitz, H. (1952). "Portfolio Selection", Journal of Finance 7, 77-91.

Mehra, R. and E.C. Prescott (1985). "The Equity Premium: A Puzzle", Journal of Monetary Economics 15, 145-161.

Merton, R.C. (1969). "Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case", Review of Economics and Statistics 51, 247-257.

Merton, R.C. (1971). "Optimum Consumption and Portfolio Rules in a Continuous-Time Model", Journal of Economic Theory 3, 373-413.

Merton, R.C. (1973). "An Intertemporal Capital Asset Pricing Model", Econometrica 41, 867-887.

Merton, R.C. (1980). "On Estimating the Expected Return on the Market: An Exploratory Investigation", Journal of Financial Economics 8, 323-361.

Neely, C.J. and P. Weller (1999). "Predictability in International Asset Returns: A Reexamination", Working Paper (Federal Reserve Bank of St. Louis), February 1999.

Pontiff, J. and L. Schall (1998). "Book-to-Market as a Predictor of Market Returns", Journal of Financial Economics 49, 141-160.

Rey, D. (2003a). "Stock Market Predictability: Is it There? A Critical Review", Working Paper (Department of Finance, WWZ University of Basel), December 2003.

Rey, D. (2003b). "Current Research Topics – Stock Market Predictability: Is it There?", Financial Markets and Portfolio Management 17, 379-387.

Rey, D. (2004). "Tactical Asset Allocation: An Alternative Definition", Working Paper (Department of Finance, WWZ University of Basel), March 2004.

Rey, D. (2005). "Market Timing and Model Uncertainty: An Exploratory Study for the Swiss Stock Market", Financial Markets and Portfolio Management, forthcoming.

Rozeff, M.S. (1984). "Dividend Yields Are Equity Risk Premiums", Journal of Portfolio Management 11, 68-75.

Samuelson, P.A. (1969). "Lifetime Portfolio Selection by Dynamic Stochastic Programming", Review of Economics and Statistics 51, 239-246.

Samuelson, P.A. (1991). "Long-Run Risk Tolerance when Equity Returns are Mean Regressing: Pseudoparadoxes and Vindication of Businessman's Risk", in Brainard, W.C., W.D. Nordhaus, and H.W. Watts, eds., "Money, Macroeconomics, and Economic Policy", MIT Press.

Samuelson, P.A. (1994). "The Long-Term Case for Equities", Journal of Portfolio Management 21, 15-24.

Schwert, G.W. (1989). "Why does Stock Market Volatility Change over Time?", Journal of Finance 44, 1115-1153.

Schwert, G.W. (2003). "Anomalies and Market Efficiency", in Constantinides, G.M., M. Harris, and R. Stulz, eds., "Handbook of the Economics of Finance", Elsevier Science.

Siegel, J.J. (1998). "Stocks for the Long Run", McGraw-Hill.

Stambaugh, R.S. (1999). "Predictive Regressions", Journal of Financial Economics 54, 315-421.

Tobin, J. (1958). "Liquidity Preference as Behavior Towards Risk", Review of Economic Studies 25, 65-85.

Viceira, L.M. (2001). "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income", Journal of Finance 56, 433-470.

Wachter, J.A. (2002). "Portfolio and Consumption Decisions Under Mean-Reverting Returns: An Exact Solution for Complete Markets", Journal of Financial and Quantitative Analysis 37, 63-91.

Weil, P. (1989). "The Equity Premium Puzzle and the Risk-Free Rate Puzzle", Journal of Monetary Economics 24, 401-421.

Williams, J. (1977). "Capital Asset Prices with Heterogeneous Beliefs", Journal of Financial Economics 5, 219-239.

Xia, Y. (2001). "Learning About Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation", Journal of Finance 56, 205-246.

Yule, G.U. (1926). "Why Do We Sometimes Get Nonsense Correlations Between Time Series? A Study in Sampling and the Nature of Time Series", Journal of the Royal Statistical Society 89, 1-64.

Zellner, A. (1971). "An Introduction to Bayesian Inference in Econometrics", Wiley.

Zimmermann, H. (2005). "Aktienrisiko und Anlagehorizont – Eine alte und wenig schlüssige Kontroverse", Neue Zürcher Zeitung, April 1, p. 31.

	Swiss Market			U.S.Market		
	1975-2002	1975-1988	1989-2002	1975-2002	1975-1988	1989-2002
Mean	6.02%	6.65%	5.40%	6.28%	6.20%	6.37%
Median	0.86%	0.51%	1.43%	0.76%	0.49%	1.05%
Max.	16.81%	16.81%	11.65%	12.10%	12.10%	10.72%
Min.	-27.57%	-27.57%	-19.92%	-23.78%	-23.78%	-16.38%
Volatility	16.13%	14.75%	17.45%	15.54%	15.68%	15.43%
Skewness	-1.1301	-1.4156	-0.9306	-0.7172	-0.7765	-0.6547
Kurtosis	8.2512	14.4346	4.7606	5.5830	7.2782	3.7710
Jarque-Bera	449.76 ***	943.14 ***	44.14 ***	119.57 ***	139.58 ***	15.48 ***
No. of Obs.	336	168	168	336	168	168

Table 2:
Summary Statistics for the Swiss and U.S. Stock Market.

The table presents summary statistics of the continuously compounded monthly excess returns. The summary statistics include mean (annualized), median, maximum and minimum values, volatility (annualized), skewness, kurtosis, the Jarque and Bera (1980) test of normality, and the number of monthly observations. Estimates are given for three different time periods: 1975 to 2002 (full sample), 1975 to 1988, and 1989 to 2002. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% significance level, respectively.

		Original Pred. Variables			Stochastically Detrended Pred. Variables		
		β	Adj. RSq.	Corr(ξ,η)	β	Adj. RSq.	Corr(ξ,η)
Swiss Market							
DPR	1975-2002	0.0020	-0.12%	-0.8664 ***	0.0002	-0.30%	-0.8499 ***
	1975-1988	0.0070 **	2.14%	-0.9574 ***	0.0030	-0.10%	-0.9439 ***
	1989-2002	-0.0003	-0.60%	-0.8168 ***	-0.0019	-0.45%	-0.7944 ***
TERM	1975-2002	0.0038	0.38%	0.0974 *	0.0021	-0.10%	0.0727
	1975-1988	0.0032	-0.03%	0.2037 ***	0.0038	0.18%	0.1737 **
	1989-2002	0.0045	0.20%	-0.0745	-0.0011	-0.56%	-0.0963
VOLA	1975-2002	-0.0027	0.03%	-0.2514 ***	-0.0026	0.02%	-0.2575 ***
	1975-1988	-0.0022	-0.34%	-0.2033 ***	-0.0049	0.71%	-0.2366 ***
	1989-2002	-0.0030	-0.24%	-0.2962 ***	-0.0004	-0.59%	-0.2685 ***
DEF	1975-2002	0.0002	-0.30%	0.0570	0.0033	0.21%	0.0684
	1975-1988	0.0033	0.00%	0.1326 *	0.0055 *	1.06%	0.1261
	1989-2002	-0.0071 *	1.37%	-0.0590	-0.0007	-0.58%	-0.0134
U.S. Market							
DPR	1975-2002	0.0036	0.34%	-0.9792 ***	0.0024	-0.02%	-0.9682 ***
	1975-1988	0.0044	0.37%	-0.9889 ***	0.0039	0.13%	-0.9762 ***
	1989-2002	0.0070 **	1.91%	-0.9694 ***	0.0007	-0.57%	-0.9611 ***
TERM	1975-2002	0.0002	-0.30%	0.0788	0.0002	-0.30%	0.0755
	1975-1988	0.0024	-0.31%	0.1530 **	0.0040	0.18%	0.1426 *
	1989-2002	-0.0028	-0.20%	-0.0582	-0.0061 *	1.31%	-0.0398
VOLA	1975-2002	-0.0004	-0.29%	-0.1566 ***	0.0001	-0.30%	-0.1555 ***
	1975-1988	0.0038	0.11%	-0.1515 **	0.0003	-0.60%	-0.1990 ***
	1989-2002	-0.0036	0.07%	-0.1581 **	-0.0002	-0.60%	-0.1203
DEF	1975-2002	0.0027	0.06%	0.0490	0.0053 **	1.08%	0.0421
	1975-1988	0.0072 **	1.96%	0.1142	0.0082 **	2.71%	0.0817
	1989-2002	-0.0030	-0.14%	-0.0575	-0.0001	-0.60%	-0.0322

Table 3:
Parameter Estimates for the VAR Model.

The results in this table are based on the model given in equation (.). The table exhibits the slope coefficients, β, obtained by regressing continuously compounded monthly excess returns on a constant and one of the following predictive variables: dividend-price ratio (DPR), term spread (TERM), realized stock market volatility (VOLA), and U.S. default risk spread (DEF). Estimates for the corresponding stochastically detrended variables are shown on the right hand side. All predictive variables are standardized with mean zero and variance one. Estimates are given for three different time periods: 1975 to 2002 (full sample), 1975 to 1988, and 1989 to 2002. Adj. RSq. denotes the adjusted coefficient of determination, and Corr(ξ,η) the correlation coefficient between innovations in the continuously compounded excess return and the respective predictive variable. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% significance level, respectively.

		Stochastic Properties of (Continuously Compounded Excess) Returns	
		Constant IOS (i.i.d.)	Time-Varying IOS (Return Predictability)
Portfolio Choice Problem	Buy-and-Hold Optimization	1a. Ignoring Estimation Risk 1b. Accounting for Estimation Risk	2a. Ignoring Estimation Risk 2b. Accounting for Estimation Risk
	Intertemporal Optimization	4a. Ignoring Estimation Risk 4b. Accounting for Estimation Risk	3a. Ignoring Estimation Risk 3b. Accounting for Estimation Risk

Table 1:
Portfolio Choice Problems and Stochastic Properties of Stock Market Returns: An Overview.

There are two possibilities of formulating the portfolio choice problem: first, a buy-and-hold strategy, where the investment horizon equals the planning horizon, second, an intertemporal strategy, where the investment horizon is shorter than the planning horizon. In the static buy-and-hold optimization problem, investors do not rebalance portfolio weights over time. The dynamic optimization problem instead leads to investors rebalancing portfolio weights regularly over the planning horizon. In addition, investment opportunities (the investment opportunity set, IOS) can be either constant or time-varying. A constant IOS corresponds to independently and identically distributed (i.i.d.) stock market returns, time-varying investment opportunities to the case of stock market predictability. Finally, investors may ignore estimation risk or take it into account.

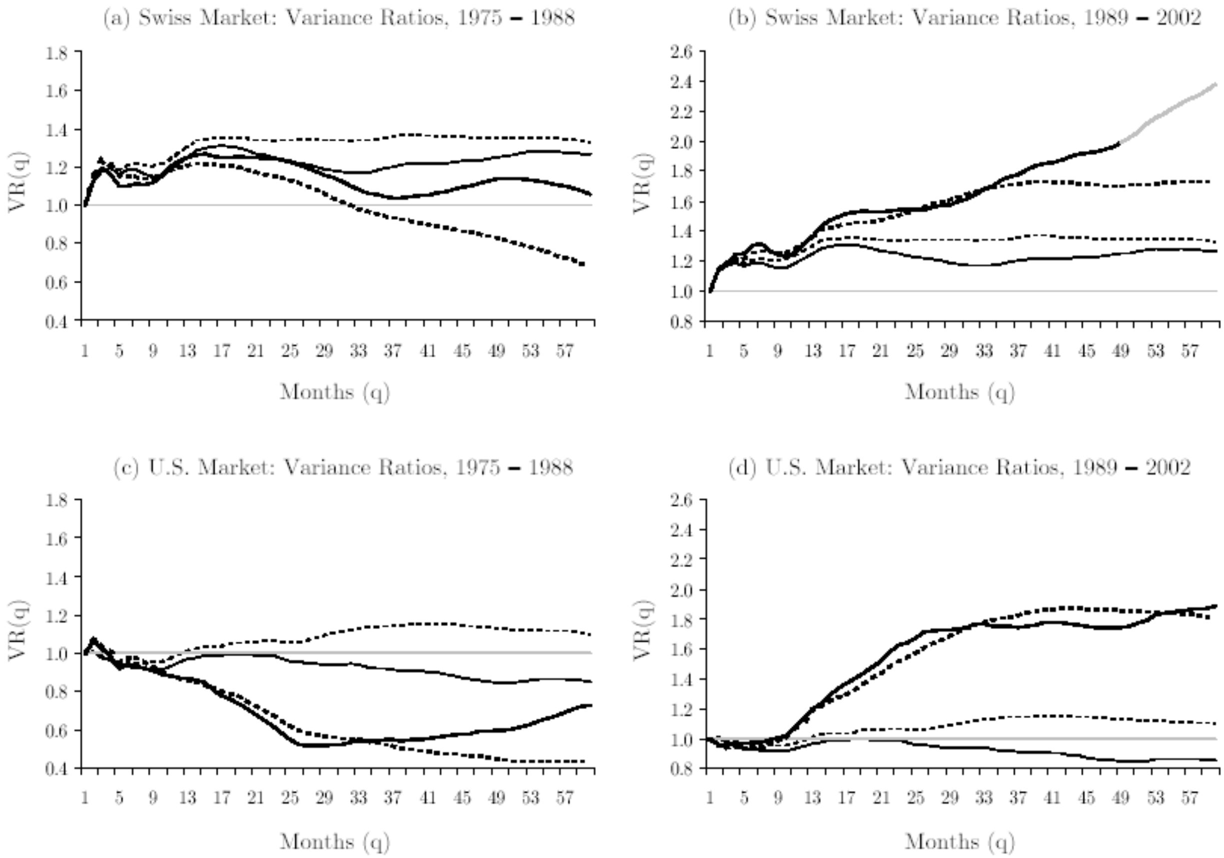

Figure 2:
Mean Reversion and Variance Ratio Statistics for the Swiss and the U.S. Stock Market.

The graphs show variance ratio statistics, $VR(q)$, for time horizons from q = 1 to 60 months. A variance ratio below one indicates mean reversion, a ratio above one mean aversion. The graphs on the left are based on continuously compounded monthly excess returns from 1975 to 1988. The graphs on the right on monthly data from 1989 to 2002. The fine lines correspond to variance ratio statistics calculated based on the full sample of monthly data (1975 to 2002). To improve the finite-sample properties, the solid lines use overlapping q-period returns and correct the bias in the variance estimators as suggested in Lo and MacKinlay (1988, 1989) and Campbell, Lo, and MacKinlay (1997, Ch. 2). In contrast, the dotted lines are calculated according to equation (.). With the exception of the Swiss stock market from 1989 to 2002 and for $q > 48$ months (gray), the variance ratio statistics are statistically not different from one at the 10% significance level.

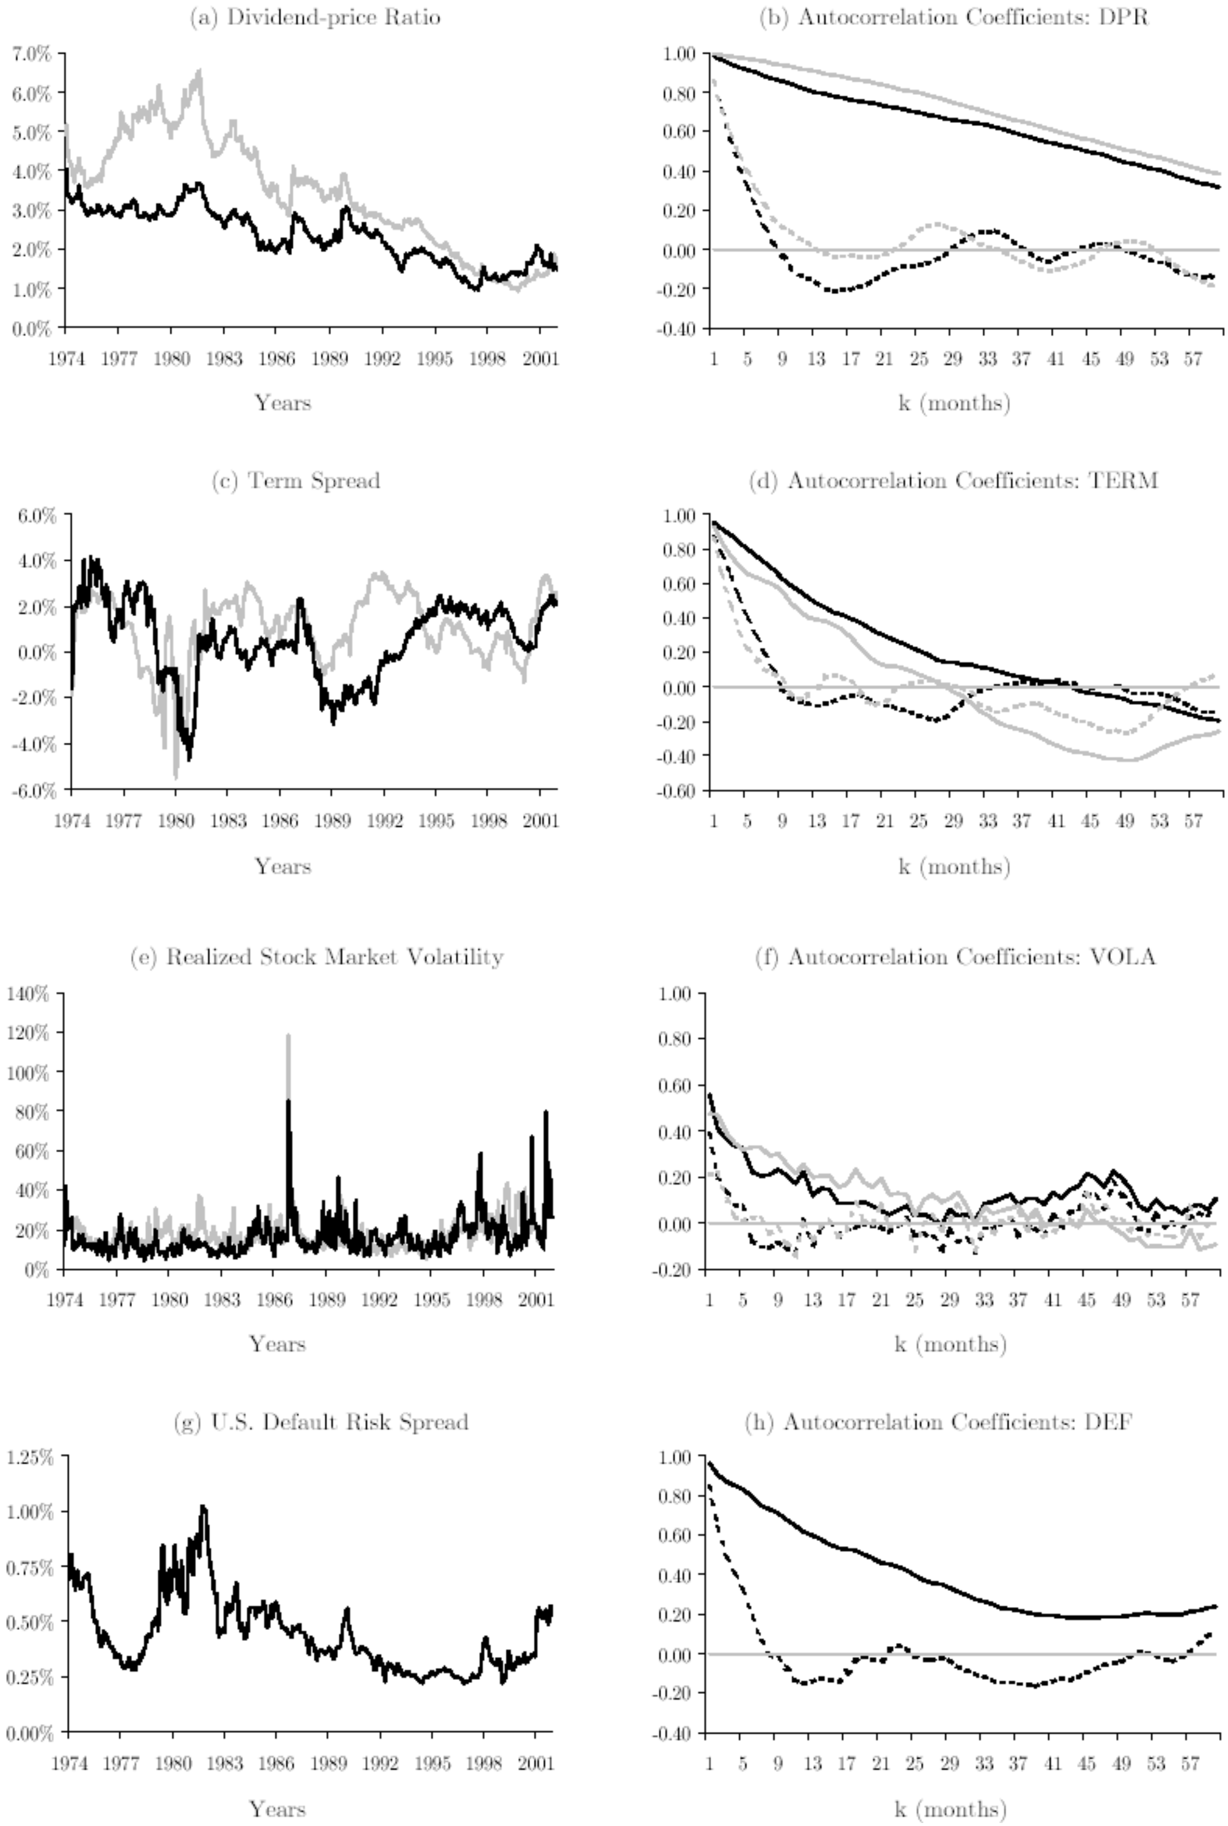
(a) Dividend-price Ratio
Years
(b) Autocorrelation Coefficients: DPR
k (months)
(c) Term Spread
Years
(d) Autocorrelation Coefficients: TERM
k (months)
(e) Realized Stock Market Volatility
Years
(f) Autocorrelation Coefficients: VOLA
k (months)
(g) U.S. Default Risk Spread
Years
(h) Autocorrelation Coefficients: DEF
k (months)

Figure 1:
The Predictive Variables for the Swiss and the U.S. Stock Market: DPR, TERM, VOLA, and DEF.

The graphs on the left show the (original) time series of the predictive variables from December 1974 to November 2002 (note, however, that in the case of the dividend-price ratio and realized stock market volatility, all calculations are based on their natural logarithms). The graphs on the right present autocorrelation coefficients estimated over the full sample of monthly data. The solid lines represent the autocorrelation coefficients of the original time series, the dotted lines of the stochastically detrended predictive variables. Black lines correspond to the Swiss stock market, gray lines to the U.S. stock market.

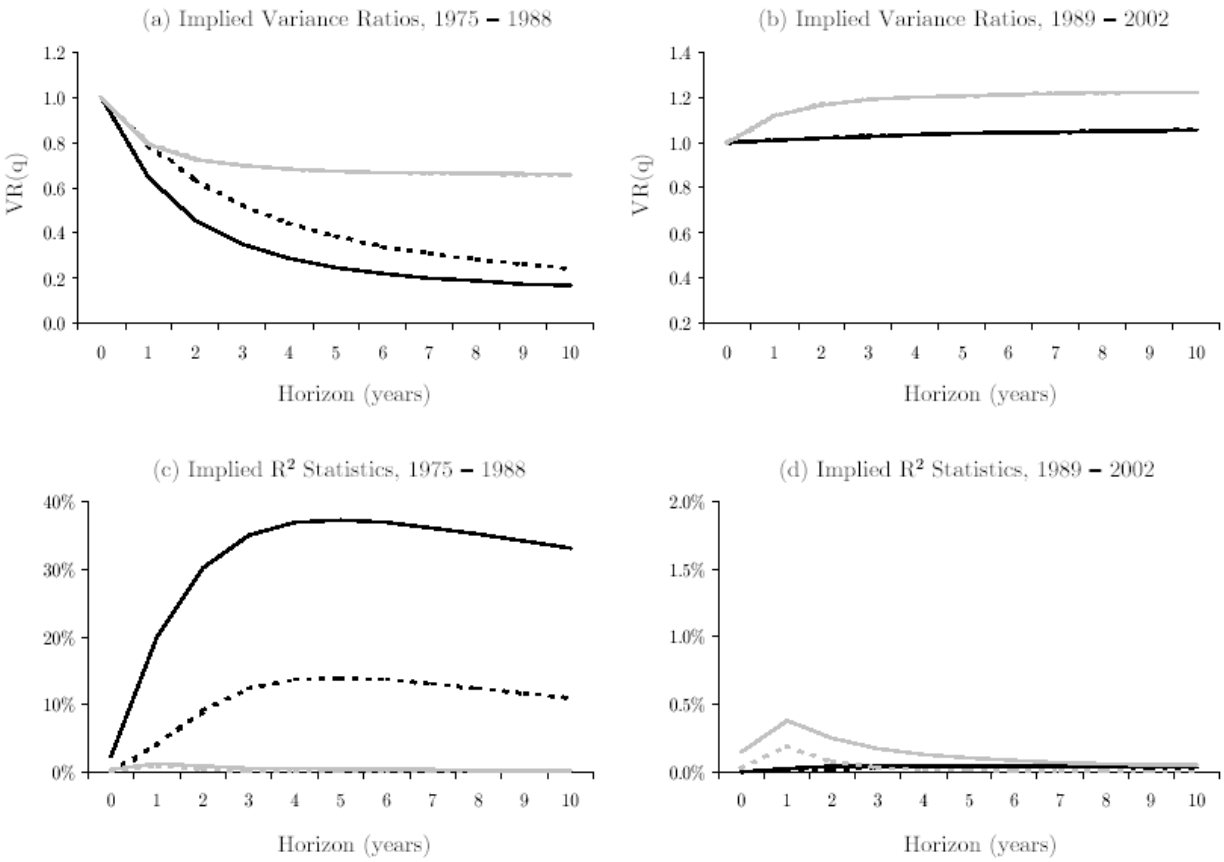

Figure 3:
Dividend-price Ratio and Swiss Stock Market. Implied Variance Ratio and R^2 Statistics.

The two graphs on the top show implied variance ratio statistics, $VR(q)$, the graphs below implied coefficients of determination, R^2. The solid lines correspond to conditional statistics, the dotted lines to the respective unconditional statistics. The black/gray lines represent the cases where continuously compounded expected excess returns are parameterized by the dividend-price ratio/the stochastically detrended dividend-price ratio. Parameter estimates are obtained using monthly data from 1975 to 1988 (graphs on the left), and 1989 to 2002 (graphs on the right).

.

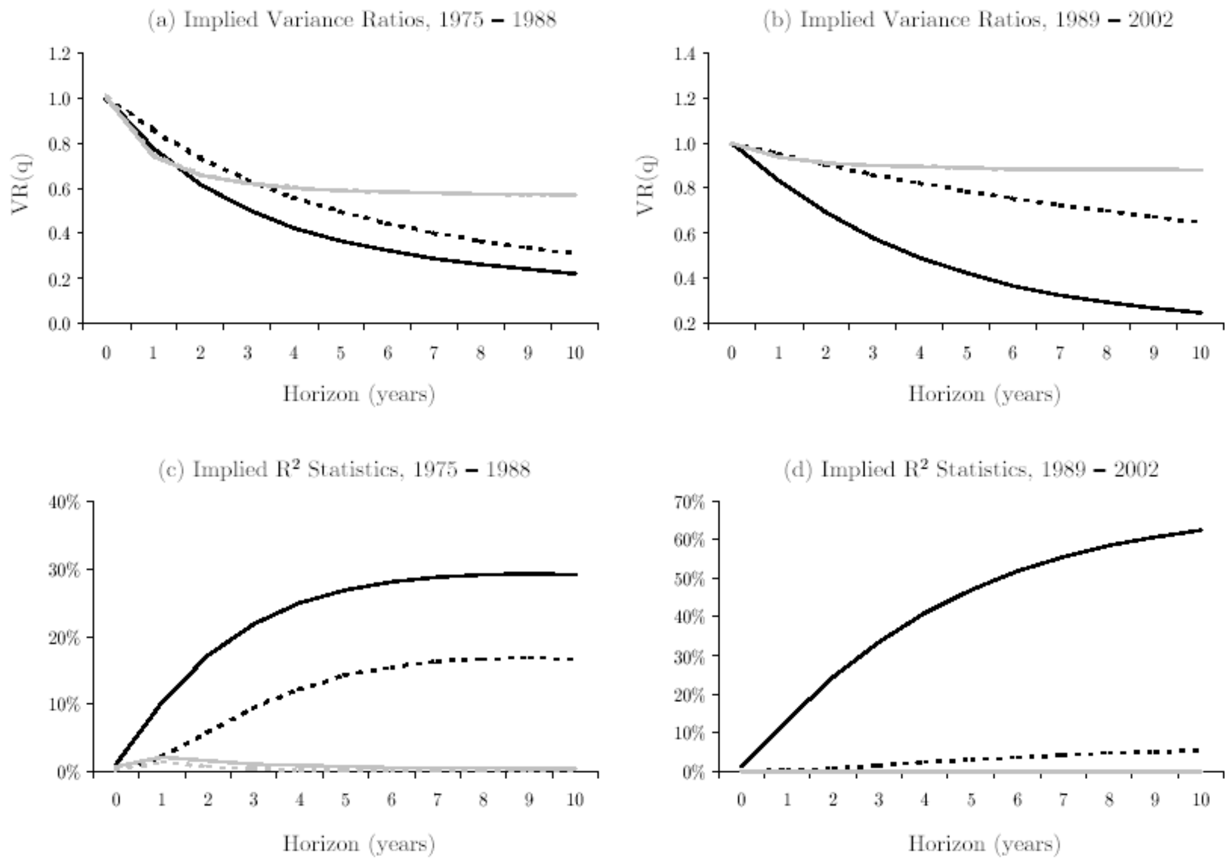

Figure 4:
Dividend-price Ratio and U.S. Stock Market. Implied Variance Ratio and R^2 Statistics.

The two graphs on the top show implied variance ratio statistics, $VR(q)$, the graphs below implied coefficients of determination, R^2. The solid lines correspond to conditional statistics, the dotted lines to the respective unconditional statistics. The black/gray lines represent the cases where continuously compounded expected excess returns are parameterized by the dividend-price ratio/the stochastically detrended dividend-price ratio. Parameter estimates are obtained using monthly data from 1975 to 1988 (graphs on the left), and 1989 to 2002 (graphs on the right).

.

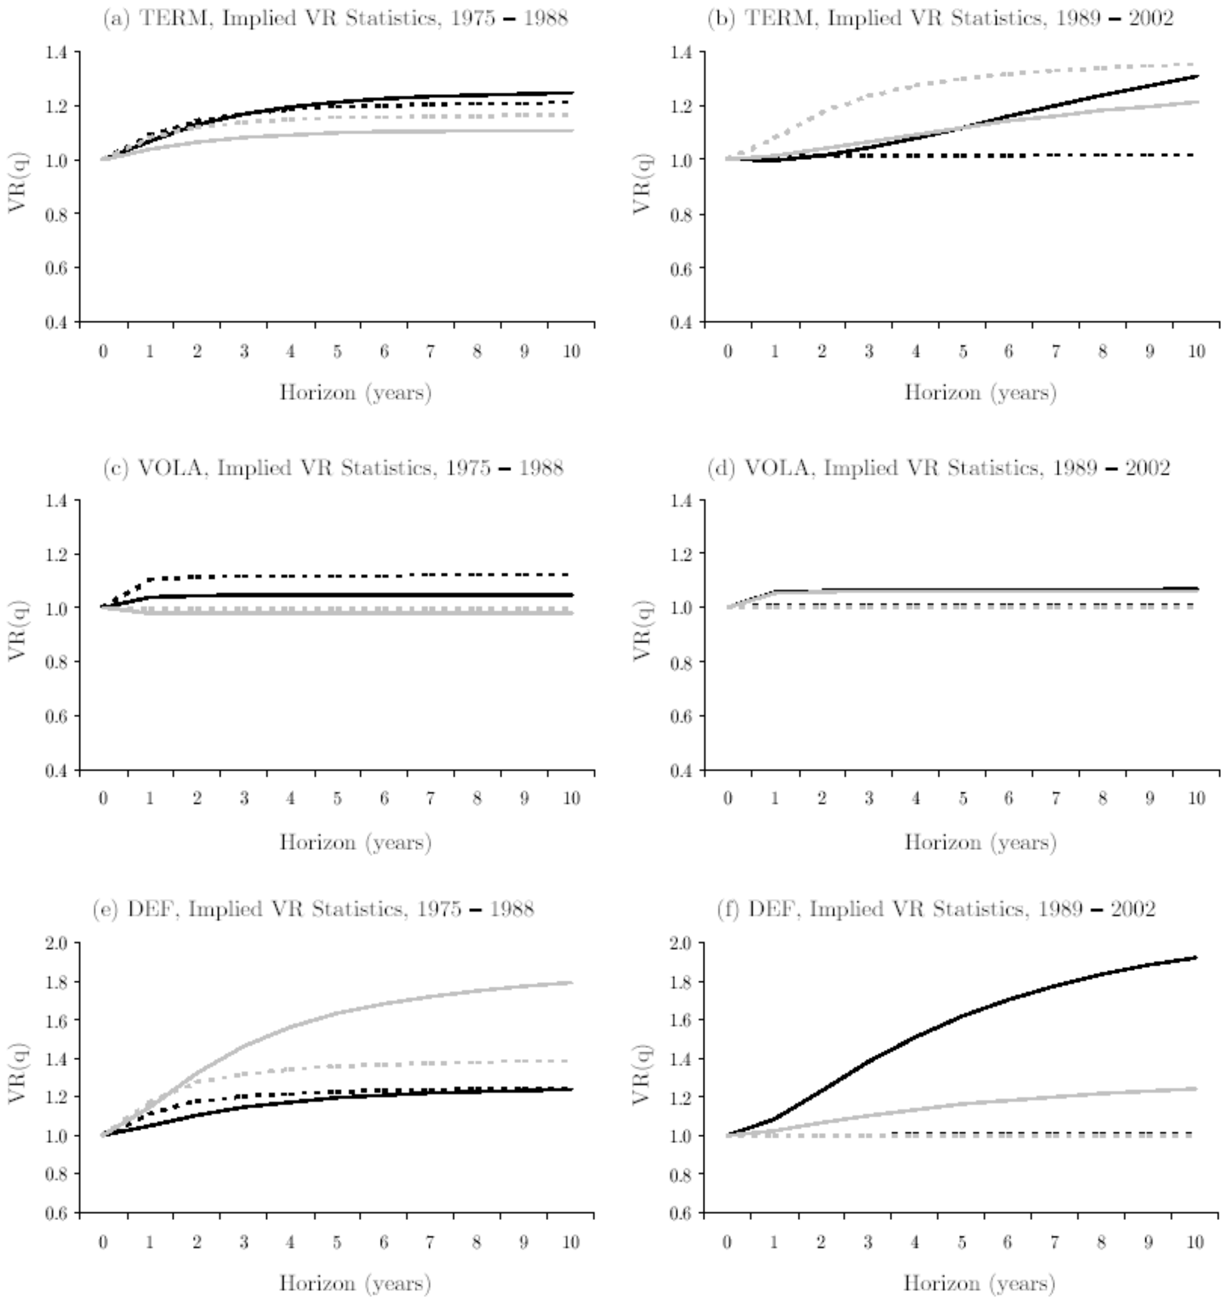

Figure 5:
Swiss and U.S. Stock Market. Implied Variance Ratio Statistics for TERM, VOLA, and DEF.

The graphs show implied conditional variance ratio statistics, $VR(q)$, for the following predictive variables: term spread (TERM), realized stock market volatility (VOLA), and U.S. default risk spread (DEF). The solid lines correspond to the original time series, the dotted lines to the stochastically detrended predictive variables. Black lines represent the Swiss stock market, gray lines the U.S. stock market. Parameter estimates are obtained using monthly data from 1975 to 1988 (graphs on the left), and 1989 to 2002 (graphs on the right).

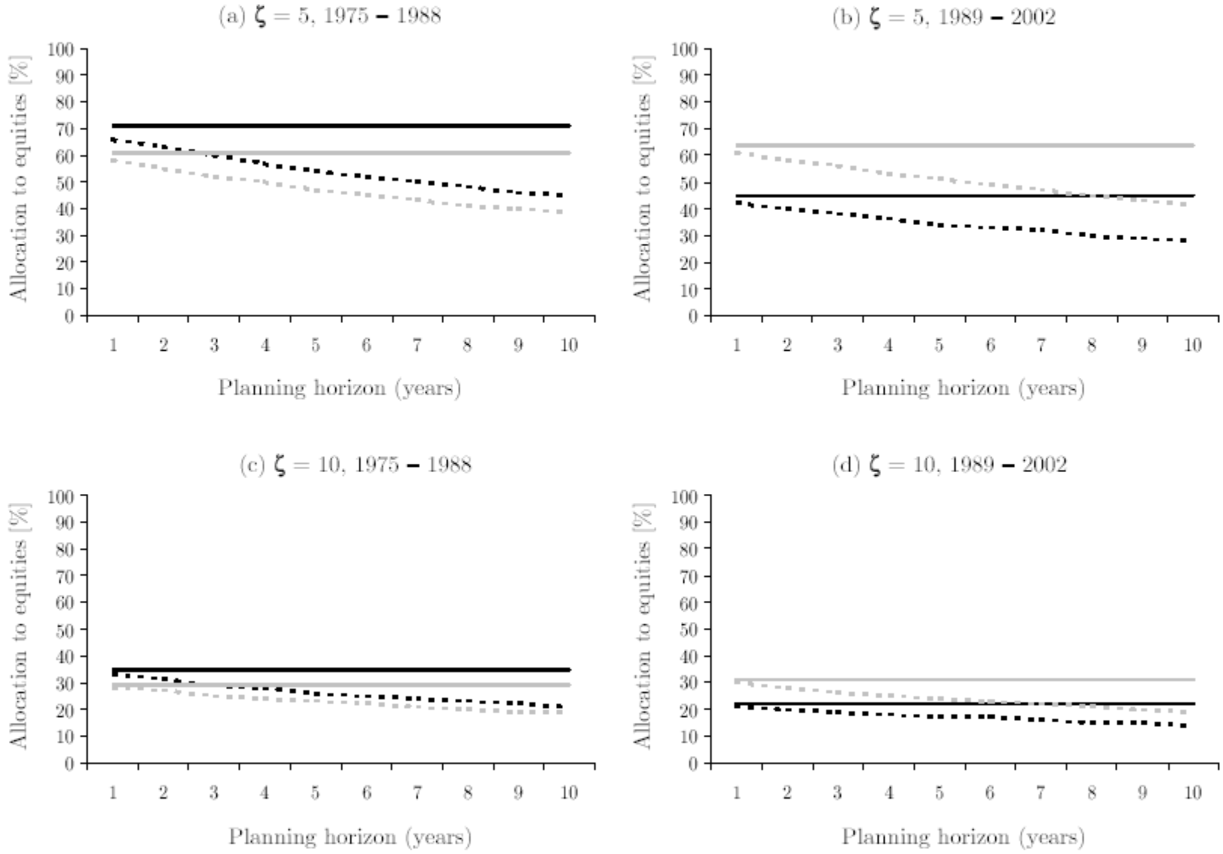

Figure 6:
Swiss and U.S. Stock Market. Optimal Allocation to Stocks Plotted Against the Planning Horizon: Buy-and-Hold Optimization and Constant Investment Opportunity Set (i.i.d.).

The graphs show the optimal allocation to stocks (in percentages) for an investor who follows a buy-a-hold strategy, uses an i.i.d. model for continuously compounded excess returns, and has power utility over terminal wealth. The solid lines correspond to the case where the investor ignores estimation risk (1a), the dotted lines where he accounts for it (1b). Parameter estimates are obtained using data from 1975 to 1988 (graphs on the left), and 1989 to 2002 (graphs on the right). Black lines represent the Swiss stock market, gray lines the U.S. stock market. ζ denotes the coefficient of relative risk aversion. The time horizon is measured on the horizontal axis (in years).

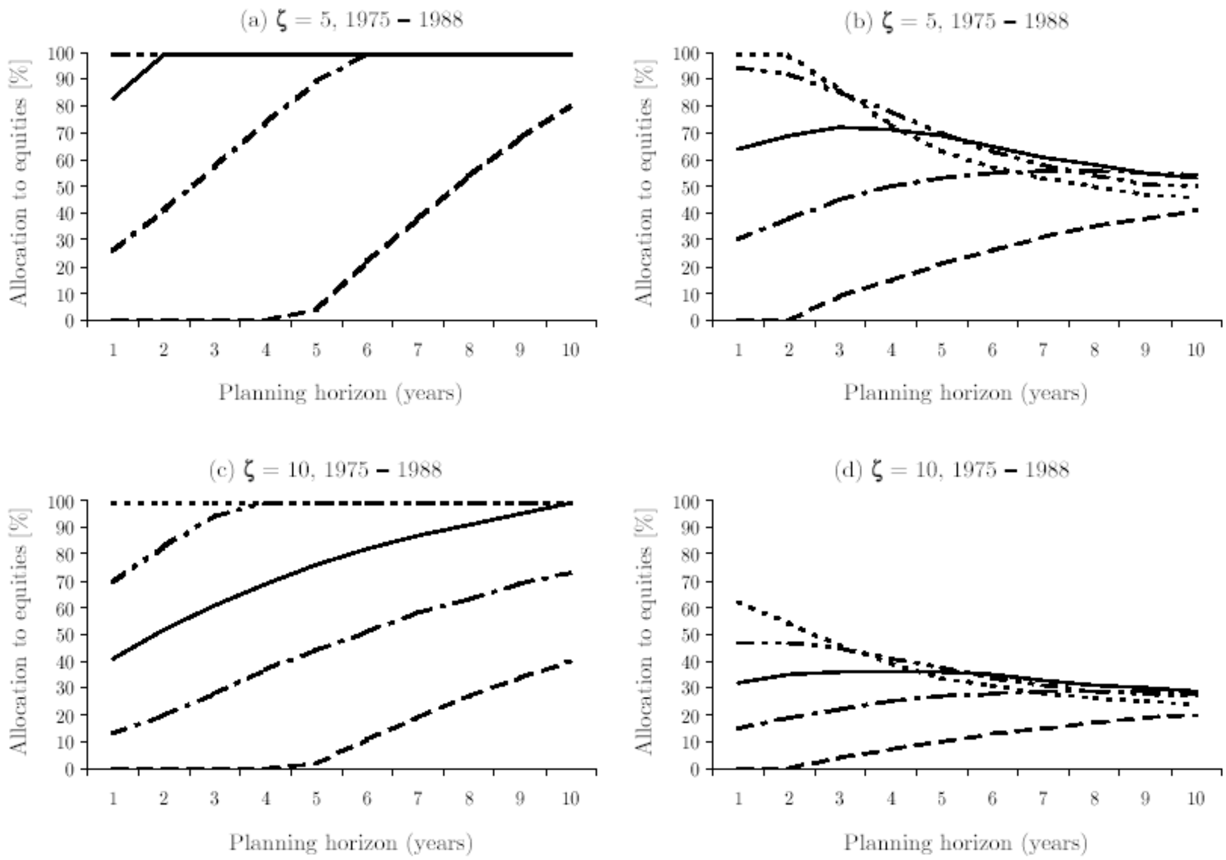

Figure 7:
Dividend-price Ratio. Swiss Stock Market. Optimal Allocation to Stocks Plotted Against the Planning Horizon: Buy-and-Hold Optimization and Time-Varying Investment Opportunity Set (Stock Market Predictability).

The graphs show the optimal allocation to stocks (in percentages) for an investor who follows a buy-a-hold strategy, uses a VAR model which allows for predictability in continuously compounded excess returns, and has power utility over terminal wealth. The graphs on the left ignore parameter uncertainty (2a), those on the right account for it (2b). The model is estimated over the 1975 to 1988 sample period using Swiss stock market data. The five lines within each graph correspond to different initial values of the (log) dividend-price ratio: 2.23% (dashed), 2.52% (dash dot), 2.77% (solid), 3.05% (dash dot dot), and 3.45% (dotted). ξ denotes the coefficient of relative risk aversion. The time horizon is measured on the horizontal axis (in years).

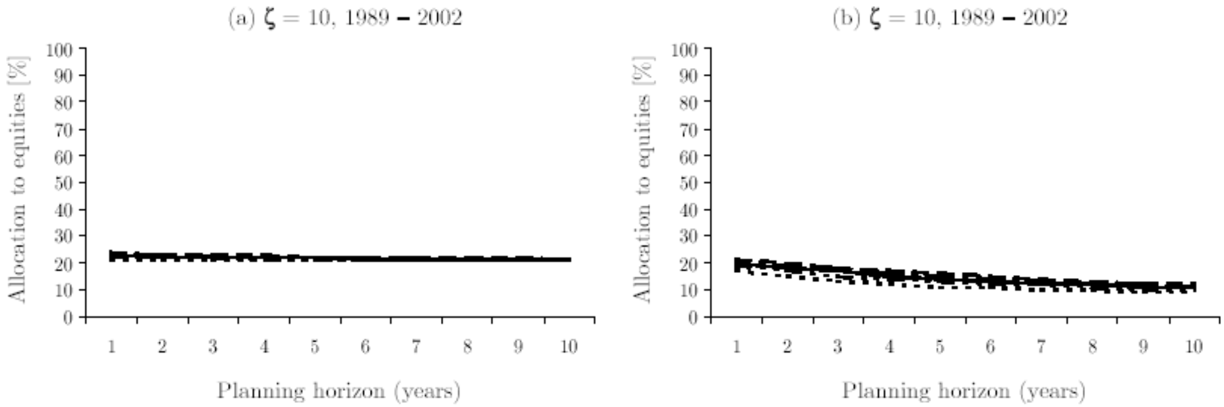

Figure 8:
Dividend-price Ratio. Swiss Stock Market. Optimal Allocation to Stocks Plotted Against the Planning Horizon: Buy-and-Hold Optimization and Time-Varying Investment Opportunity Set (Stock Market Predictability).

The graphs show the optimal allocation to stocks (in percentages) for an investor who follows a buy-a-hold strategy, uses a VAR model which allows for predictability in continuously compounded excess returns, and has power utility over terminal wealth. The graph on the left ignores parameter uncertainty (2a), the one on the right accounts for it (2b). The model is estimated over the 1989 to 2002 sample period using Swiss stock market data. The five lines within each graph correspond to different initial values of the (log) dividend-price ratio: 1.22% (dashed), 1.48% (dash dot), 1.71% (solid), 1.98% (dash dot dot), and 2.41% (dotted). ζ denotes the coefficient of relative risk aversion. The time horizon is measured on the horizontal axis (in years).

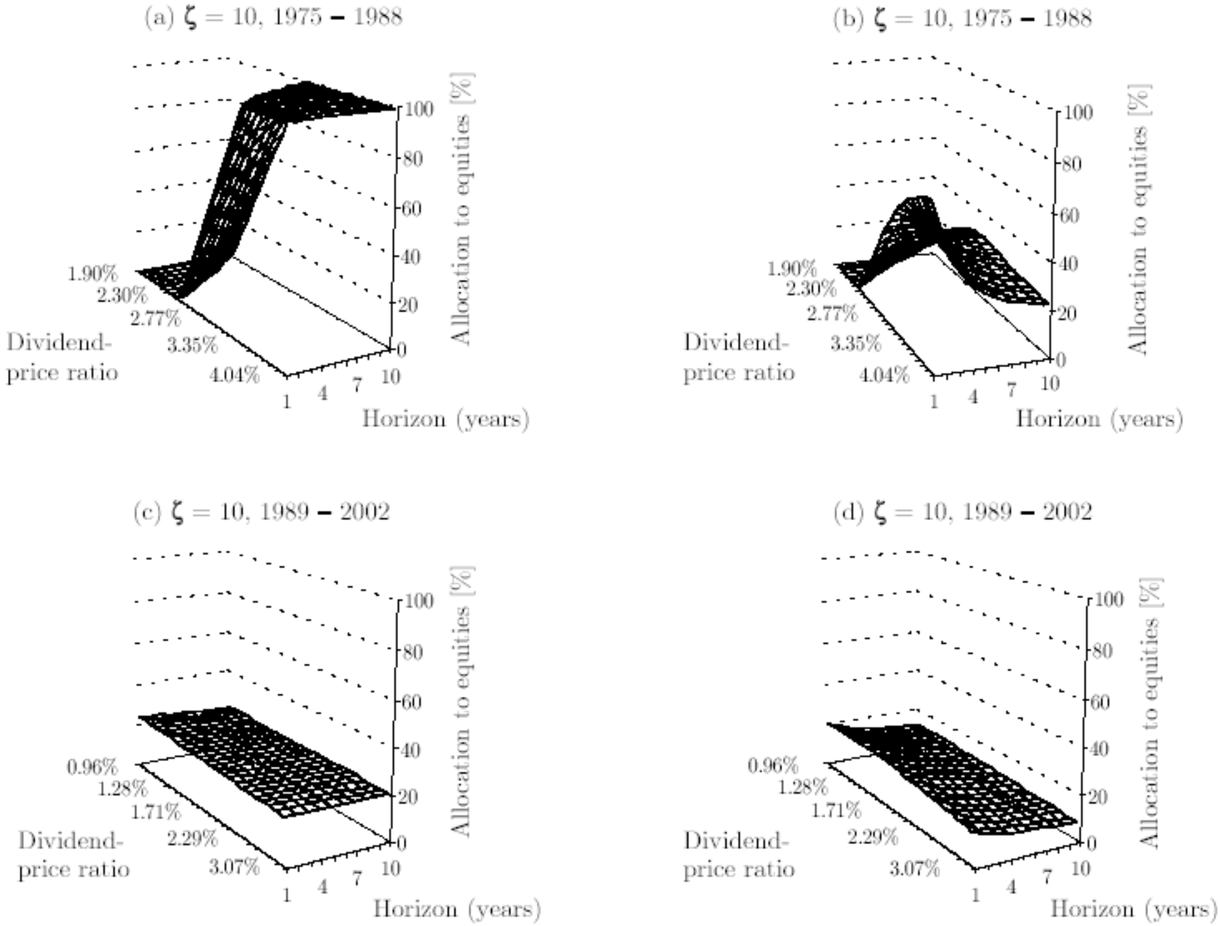

Figure 9:
Dividend-price Ratio. Swiss Stock Market. Optimal Allocation to Stocks Plotted Against the Planning Horizon: Buy-and-Hold Optimization and Time-Varying Investment Opportunity Set (Stock Market Predictability).

The graphs show the optimal allocation to stocks (in percentages) for an investor who follows a buy-a-hold strategy, uses a VAR model which allows for predictability in continuously compounded excess returns, and has power utility over terminal wealth. The graphs on the left ignore parameter uncertainty (2a), those on the right account for it (2b). The top row graphs use Swiss stock market data from 1975 to 1988, the lower two use data from 1989 to 2002. The predictive variable is the (log) dividend-price ratio, plotted for the interval ranging from its maximum and minimum value over the respective sample period, discretized with 25 equally spaced grid points. ζ denotes the coefficient of relative risk aversion. The time horizon is measured on the horizontal axis (in years).

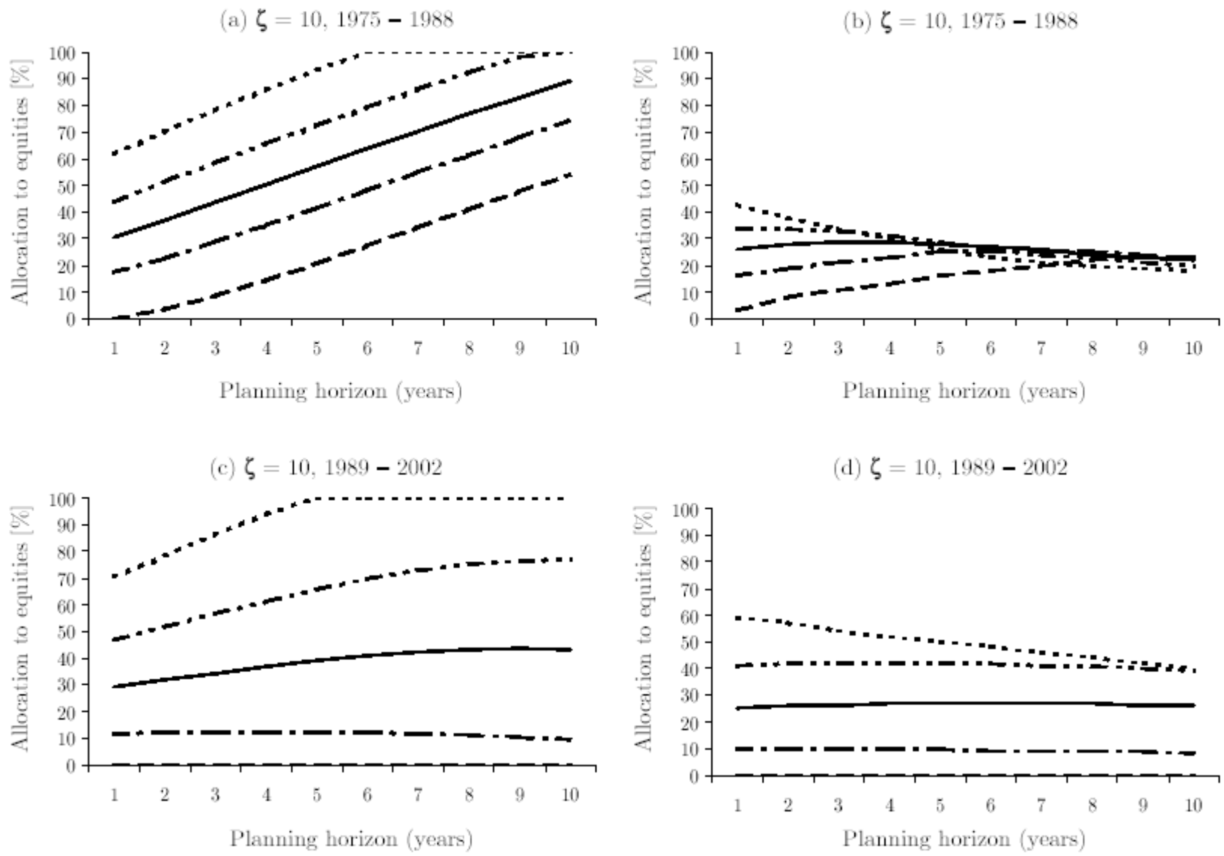

Figure 10:
Dividend-price Ratio. U.S. Stock Market. Optimal Allocation to Stocks Plotted Against the Planning Horizon: Buy-and-Hold Optimization and Time-Varying Investment Opportunity Set (Stock Market Predictability).

The graphs show the optimal allocation to stocks (in percentages) for an investor who follows a buy-a-hold strategy, uses a VAR model which allows for predictability in continuously compounded excess returns, and has power utility over terminal wealth. The graphs on the left ignore parameter uncertainty (2a), those on the right account for it (2b). The top row graphs use U.S. stock market data from 1975 to 1988, the lower two use data from 1989 to 2002. The five lines within each graph correspond to different initial values of the (log) dividend-price ratio: 3.38%/1.28% (dashed), 3.89%/1.61% (dash dot), 4.32%/1.93% (solid), 4.79%/2.30% (dash dot dot), and 5.51%/2.91% (dotted) for the first/second time period. ζ denotes the coefficient of relative risk aversion. The time horizon is measured on the horizontal axis (in years).

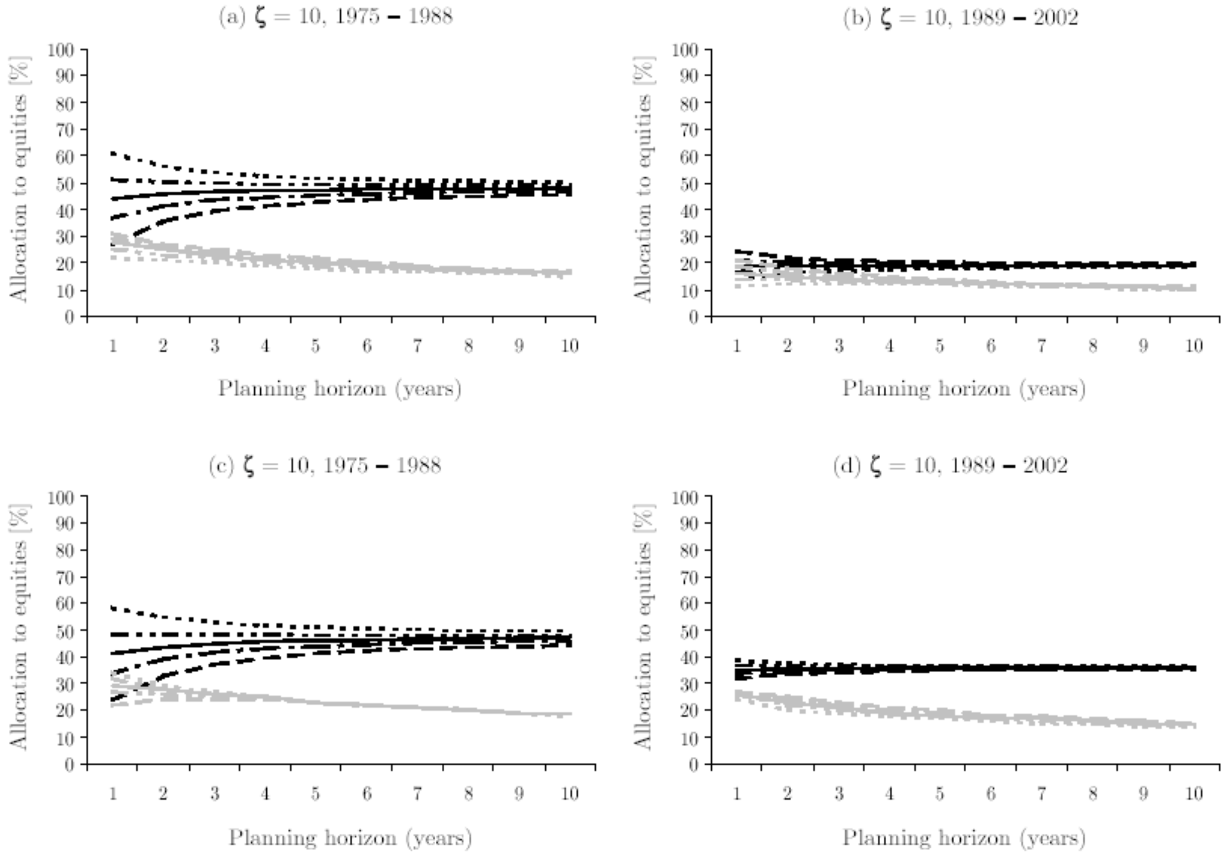

Figure 11:
Stochastically Detrended Dividend-price Ratio. Swiss and U.S. Stock Market. Optimal Allocation to Stocks Plotted Against the Planning Horizon: Buy-and-Hold Optimization and Time-Varying Investment Opportunity Set (Stock Market Predictability).

The graphs show the optimal allocation to stocks (in percentages) for an investor who follows a buy-a-hold strategy, uses a VAR model which allows for predictability in continuously compounded excess returns, and has power utility over terminal wealth. The top row graphs use Swiss stock market data from 1975 to 1988 (left) and 1989 to 2002 (right), the lower two use U.S. stock market data over the same time periods. Black lines ignore parameter uncertainty (2a), gray lines account for it (2b). In all cases, the five lines correspond to different initial values of the stochastically detrended dividend-price ratio, ascending from the lowest (dashed) to the highest (dotted) initial value. ζ denotes the coefficient of relative risk aversion. The time horizon is measured on the horizontal axis (in years).

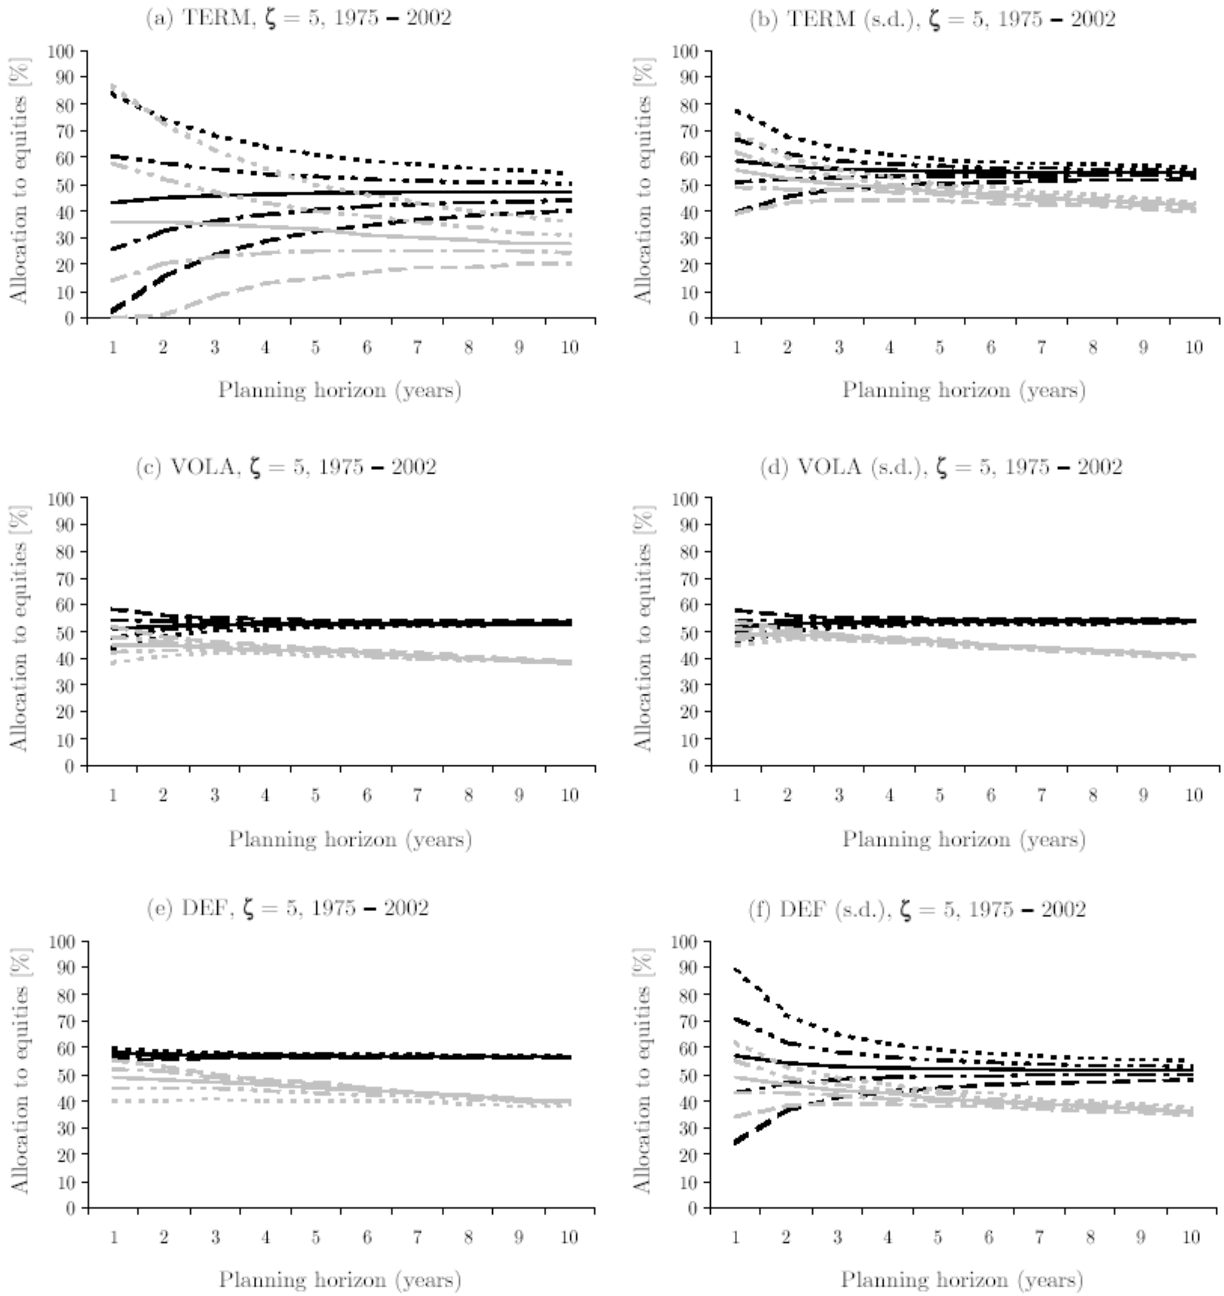

Figure 12:
Swiss Stock Market. TERM, VOLA, and DEF. Optimal Allocation to Stocks Plotted Against the Planning Horizon: Buy-and-Hold Optimization and Time-Varying Investment Opportunity Set (Stock Market Predictability).

The graphs show the optimal allocation to stocks (in percentages) for an investor who follows a buy-a-hold strategy, uses a VAR model which allows for predictability in continuously compounded excess returns, and has power utility over terminal wealth. The predictive variables are: term spread (TERM), realized stock market volatility (VOLA), and U.S. default risk spread (DEF). The graphs on the left correspond to the original time series, those on the right to the stochastically detrended predictive variables. Black lines ignore parameter uncertainty (2a), gray lines account for it (2b). Parameter estimates are obtained over the full period from 1975 to 2002 using Swiss stock market data. In all cases, the five lines correspond to different initial values of the predictive variables, ascending from the lowest (dashed) to the highest (dotted) initial value. ζ denotes the coefficient of relative risk aversion. The time horizon is measured on the horizontal axis (in years).

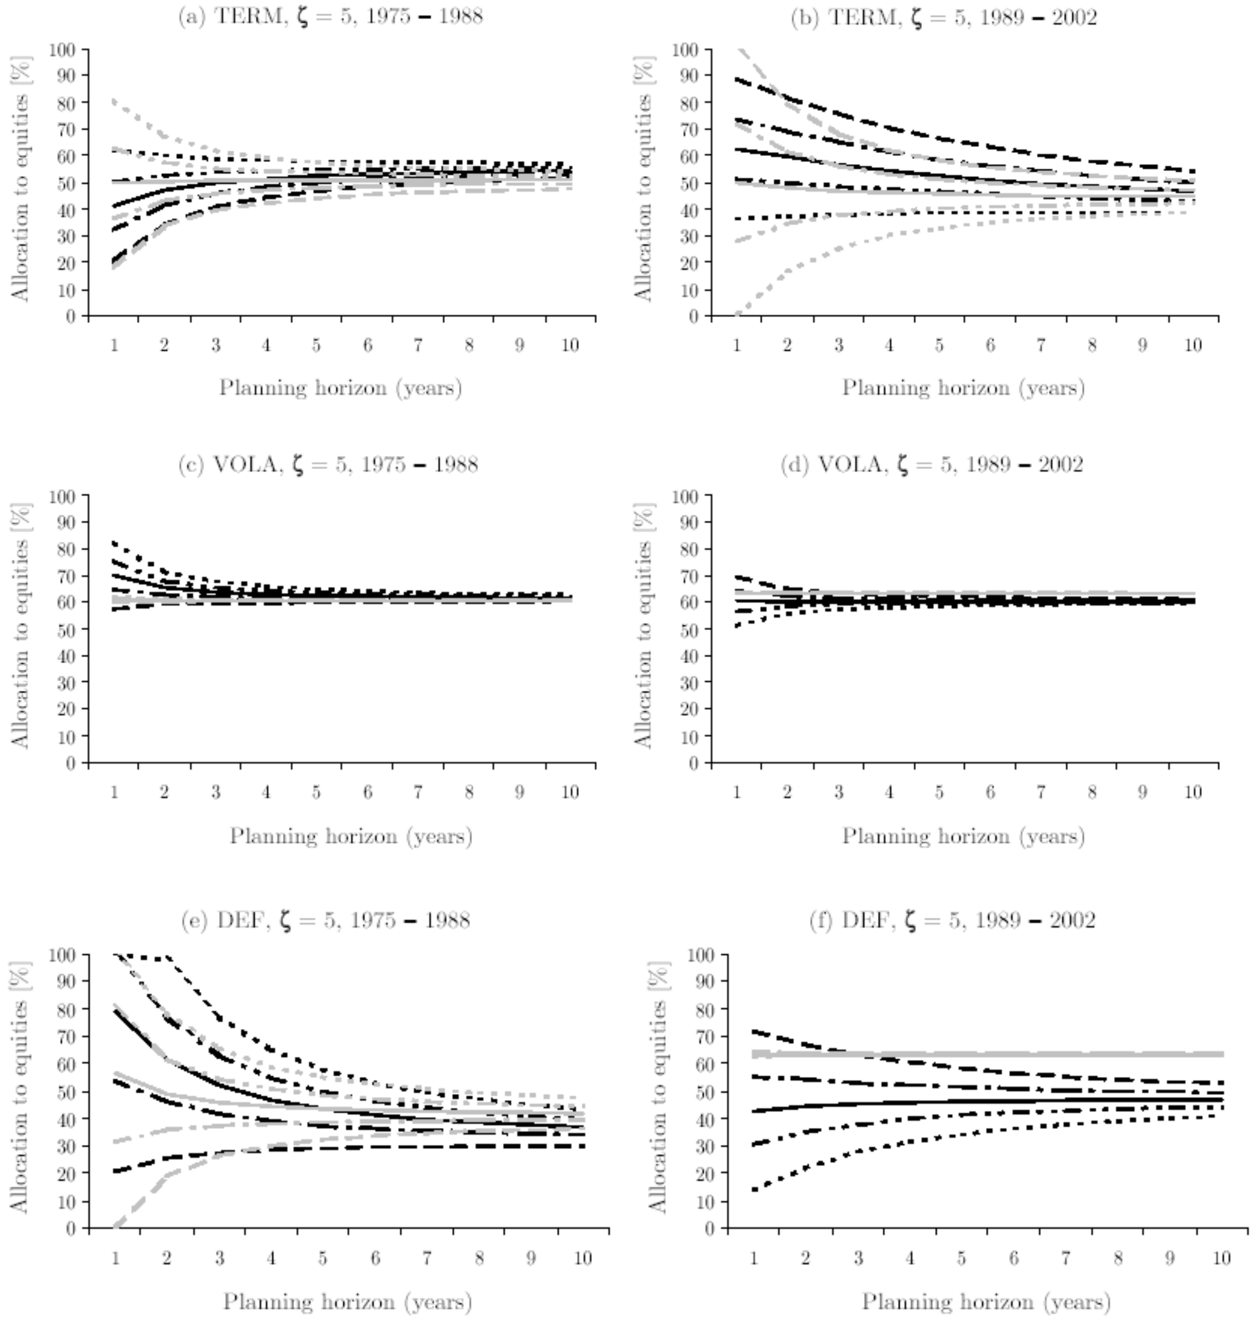

Figure 13:
U.S. Stock Market. TERM, VOLA, and DEF. Optimal Allocation to Stocks Plotted Against the Planning Horizon: Buy-and-Hold Optimization and Time-Varying Investment Opportunity Set (Stock Market Predictability).

The graphs show the optimal allocation to stocks (in percentages) for an investor who follows a buy-a-hold strategy, uses a VAR model which allows for predictability in continuously compounded excess returns, and has power utility over terminal wealth. The predictive variables are: term spread (TERM), realized stock market volatility (VOLA), and U.S. default risk spread (DEF). The graphs on the left use U.S. stock market data from 1975 to 1988, those on the right are based on data from 1989 to 2002. Black lines correspond to the original time series, gray lines correspond to the stochastically detrended predictive variables. Parameter uncertainty is ignored throughout (2a). In all cases, the five lines correspond to different initial values of the predictive variables, ascending from the lowest (dashed) to the highest (dotted) initial value. ξ denotes the coefficient of relative risk aversion. The time horizon is measured on the horizontal axis (in years).

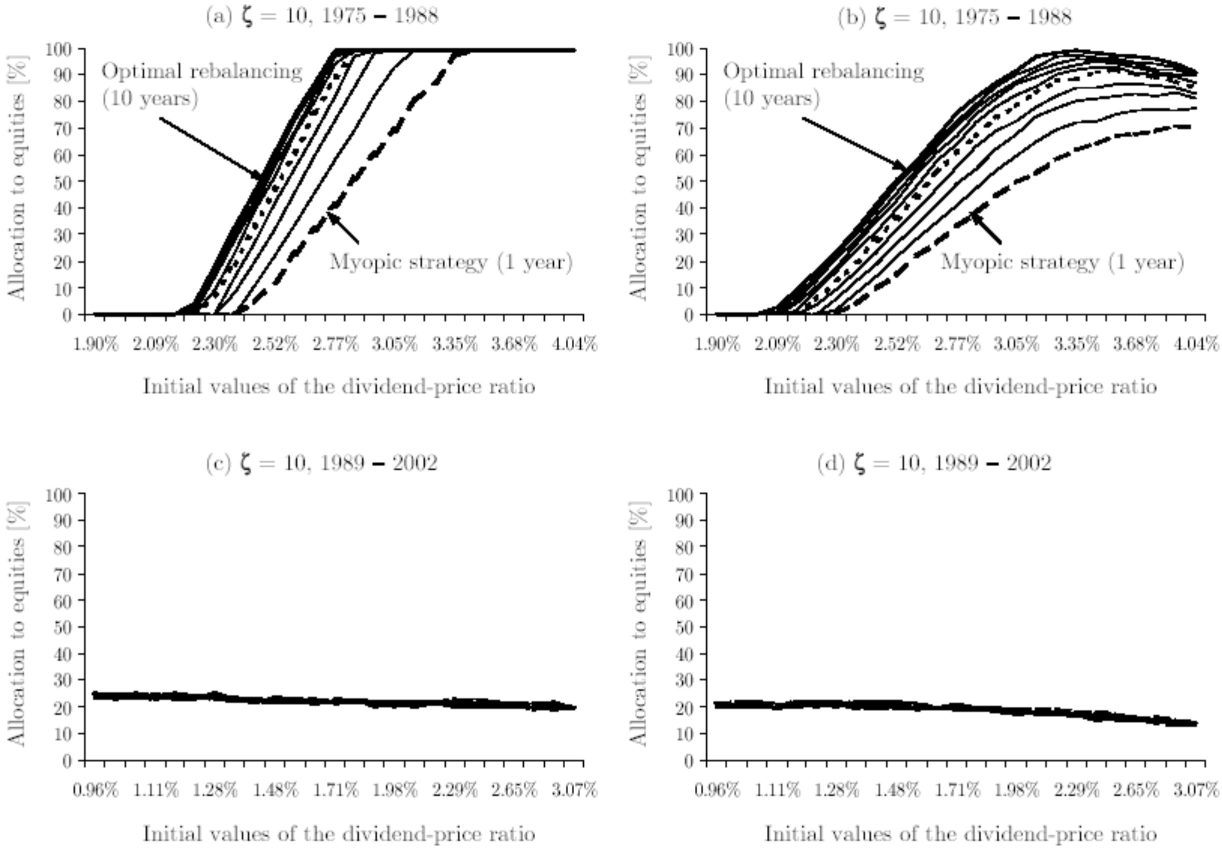

Figure 14:
Dividend-price Ratio. Swiss Stock Market. Optimal Allocation to Stocks Plotted Against the Planning Horizon: Intertemporal Optimization and Time-Varying Investment Opportunity Set (Stock Market Predictability).

The graphs show the optimal allocation to stocks (in percentages) for an investor who rebalances optimally once a year, uses a VAR model which allows for predictability in continuously compounded excess returns, and has power utility over terminal wealth. The graphs on the left ignore parameter uncertainty (3a), those on the right account for it (3b). The top row graphs use Swiss stock market data from 1975 to 1988, the lower two use data from 1989 to 2002. Initial values of the (log) dividend-price ratio are depicted for the interval ranging from the maximum and minimum value over the respective sample period, discretized with 25 equally spaced grid points. The ten lines within each graph correspond to different planning horizons, ranging from one year (dashed) to 10 years (solid). ζ denotes the coefficient of relative risk aversion.

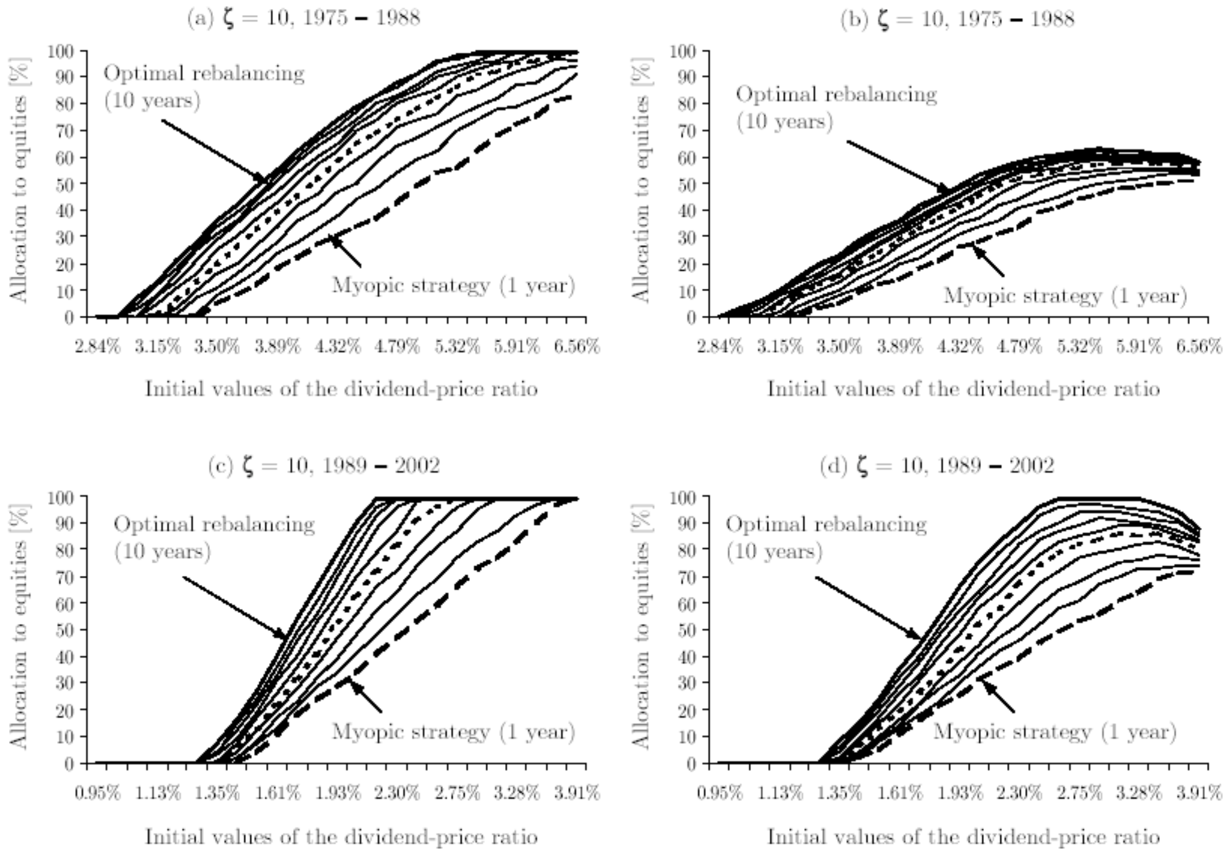

Figure 15:
Dividend-price Ratio. U.S. Stock Market. Optimal Allocation to Stocks Plotted Against the Planning Horizon: Intertemporal Optimization and Time-Varying Investment Opportunity Set (Stock Market Predictability).

The graphs show the optimal allocation to stocks (in percentages) for an investor who rebalances optimally once a year, uses a VAR model which allows for predictability in continuously compounded excess returns, and has power utility over terminal wealth. The graphs on the left ignore parameter uncertainty (3a), those on the right account for it (3b). The top row graphs use U.S. stock market data from 1975 to 1988, the lower two use data from 1989 to 2002. Initial values of the (log) dividend-price ratio are depicted for the interval ranging from the maximum and minimum value over the respective sample period, discretized with 25 equally spaced grid points. The ten lines within each graph correspond to different planning horizons, ranging from one year (dashed) to 10 years (solid). ζ denotes the coefficient of relative risk aversion.