Market Timing Does Work: Evidence from the NYSE[1]

Devraj Basu

Alexander Stremme

Warwick Business School, University of Warwick

November 2005

address for correspondence:

Alexander Stremme

Warwick Business School

University of Warwick

Coventry, CV4 7AL

United Kingdom

e-mail: `alex.stremme@wbs.ac.uk`

[1]We would like to thank John Cochrane for helpful comments on an earlier version of this paper. All remaining errors are ours.

Devraj Basu
Alexander Stremme

November 2005

Abstract

In this paper we use a new approach to construct unconditionally efficient market-timing strategies that optimally utilize the predictive information inherent in commonly used macroeconomic and term-structure variables. We also construct a statistical test to evaluate the performance of our strategies. We find that in-sample, our strategies almost double the unconditional Sharpe ratio of our benchmark index. We also compare the performance of our unconditionally efficient strategies with that of more traditional conditionally efficient portfolios. We find that our strategies not only show clearly superior performance, but their portfolio allocation weights respond much more conservatively to changes in the predictive information, resulting in significantly lower transaction costs. The out-of-sample performance of our strategies is broadly consistent with the in-sample estimates. For example, our market-timing strategy achieves an out-of-sample alpha of 13.5% against the benchmark, suggesting that a portfolio manager who followed this strategy could have made significant gains.

JEL Classification: C31, G11, G12

1 Introduction

The purpose of this note is to investigate if optimal market timing can generate excess performance and hence produce significant economic gains. To this end, we construct unconditionally efficient, dynamically managed trading strategies via the optimal use of conditioning information in portfolio formation as developed in Ferson and Siegel (2001), and Abhyankar, Basu, and Stremme (2005b). We evaluate the in-sample and out-of-sample performance of these strategies using a statistical test developed in Abhyankar, Basu, and Stremme (2005a), based on an idea first considered in Cochrane (1999). This statistical test initially allows us to determine whether *any* market-timing strategy will work. Moreover, implicit in its construction is the best-performing such strategy. This strategy optimally utilizes the predictive information inherent in commonly used indicators in order to efficiently time the market.

We consider benchmark-timing with the CRSP value-weighted index as the risky asset, and the 1-month Treasury bill rate as the risk-free asset. We employ a number of commonly used macro-economic and term-structure variables as conditioning instruments. The former include inflation, growth in aggregate consumption and unemployment. The latter include the short rate, term spread and convexity, as well as the credit yield spread. Although in a predictive regression, less than 5% of the time-variation of returns is explained by the predictive instruments, our framework allows us to use even these small levels of predictability to obtain large economic gains. To this end, we construct efficiently managed market-timing strategies, and assess their ex-post performance using a variety of standard performance measures. In these strategies the investment-mix between the risky and risk-free assets is optimally managed as a function of the predictive instruments. This is in contrast to the simple switching strategies prevalent in much of the literature, which simply move in and out of the market.

We find that in-sample Sharpe ratios more than double when all our predictive variables are optimally utilized. The p-value associated with our test statistic is indistinguishable from zero, indicating that the increase in Sharpe ratio is highly significant. The term-structure and macro-economic variables used separately also lead to significant economic gains. Our optimal market-timing strategies have considerably higher means than the corresponding fixed-weight strategies. Constraining the portfolio weights to exclude short selling does not lead to much reduction in performance. This is due in part to the 'conservative response' of the optimal portfolio weights, first noted in Ferson and Siegel (2001). A more detailed analysis of the strategies reveals that the behavior of the optimal portfolio weights is largely driven by time-variations in the conditional Sharpe ratio. In particular, we show that the gain from predictability is higher the higher the volatility of the conditional Sharpe ratio, a point also noted in Cochrane (1999).

Our strategies are designed to be *unconditionally* efficient and hence optimal with respect to *ex-post* performance criteria, in contrast to traditional conditionally efficient portfolios that are optimal with respect to the one-period-ahead conditional distribution of returns. While unconditionally efficient strategies are necessarily also conditionally efficient, the converse is generally not true. Our strategies not only show significantly better performance, but also display much less variability in portfolio weights, both in and out-of-sample. The latter is important in particular when considering transaction costs. Moreover, the conditionally efficient strategies often require extreme long and short positions, which makes them much more sensitive to the introduction of short-sale restrictions.

The out-of-sample performance of our strategies is broadly consistent with the in-sample results. In particular, while the performance of a traditional efficient static portfolio deteriorates considerably out-of-sample (barely matching the benchmark performance), our efficiently managed active strategies maintain their performance (achieving Sharpe ratios in excess of 0.6 and alphas relative to the benchmark of more than 13%).

The remainder of this paper is organized as follows. In Section 2, we briefly review the theoretical background, discuss the construction of our efficient strategies and define our statistical test. In Section 3, we outline our empirical methodology and present the results of our analysis. Section 4 concludes. For brevity, most mathematical proofs are omitted in this note. Details are available from the authors upon request.

2 Efficient Market Timing Strategies

The flow of information in the economy is described by a discrete-time filtration $\{\mathcal{F}_t\}_t$, defined on some probability space $(\Omega, \mathcal{F}, P)$. There are two traded assets, a market index portfolio M, and a risk-free asset (e.g. a Treasury bill). Denote by r_t^M the gross return from time $t-1$ to time t on the index portfolio (i.e. the future value at time t of \$1 invested at time $t-1$). Similarly, denote by r_{t-1}^f the (gross) return on the risk-free asset over the same period. The difference in time indexing indicates that, while the return r_{t-1}^f on the risk-free asset is known at the beginning of the period (i.e. at time $t-1$), the return r_t^M on the index portfolio is uncertain *ex-ante* and only realized at the end of the period (i.e. at time t). Note however that we do *not* assume r_{t-1}^f to be unconditionally constant. In other words, while the return on the risk-free asset is known with certainty at the beginning of any one investment period, it may (and in general will) vary over time from one period to the next. In this sense, r_{t-1}^f is only *conditionally* risk-free but not unconditionally.

Conditioning Information

Denote by $\mathcal{G}_{t-1} \subseteq \mathcal{F}_{t-1}$ the information set on which investors base their asset allocation decisions at time $t-1$. In our empirical applications, $\mathcal{G}_{t-1}$ will be given by a set of lagged *instruments* y_{t-1}, variables observable at time $t-1$ that contain information about the conditional distribution of risky asset returns. Finally, we denote by $E_{t-1}(\,\cdot\,)$ the conditional expectation with respect to $\mathcal{G}_{t-1}$. A dynamically managed market timing strategy therefore is a sequence of 'weights' $\{\theta_t\}_t$, where θ_{t-1} denotes the fraction of the investor's wealth that is invested in the risky asset (i.e. the index portfolio) at time $t-1$. The return on such a strategy over the period from time $t-1$ to t is hence given by,

$$r_t(\theta) = r_{t-1}^f + (\,r_t^M - r_{t-1}^f\,)\theta_{t-1}. \tag{1}$$

Our aim is to find strategies that optimally exploit the information contained in the information set $\mathcal{G}_{t-1}$. However, instead of specifying θ_{t-1} conditionally period-by-period, we characterize the optimal strategy ex-ante as a function of the conditioning information.

Efficient Strategies

Denote by $\mu_{t-1} = E_{t-1}(r_t^M)$ the conditional expected return on the index portfolio. Traditional market timing strategies would simply switch between the index portfolio and the risk-free asset on the basis of the sign of $\mu_{t-1} - r_{t-1}^f$. In contrast, we consider here strategies that are *unconditionally* efficient in the sense that they minimize the unconditional variance of returns $r_t(\theta)$ for given unconditional mean. In other words, our strategies are designed to be optimal with respect to *ex-post* performance criteria. In particular, our strategies will attain the maximal achievable unconditional Sharpe ratio.

Let $\sigma_{t-1}^2 = E_{t-1}((r_t^M - \mu_{t-1})^2)$ denote the conditional variance of the risky asset return. Although in the empirical applications in this paper, we will assume σ_{t-1} to be constant, the theoretical results stated below hold also when the conditional variance is allowed to be time-varying. Obviously, the *conditional* Sharpe ratio for the period from time $t-1$ to time t can thus be written as $H_{t-1} = (\mu_{t-1} - r_{t-1}^f)/\sigma_{t-1}$. It can now be shown[1] that any unconditionally efficient strategy can be written as,

$$\theta_{t-1}^* = \frac{w - r_{t-1}^f}{1 + H_{t-1}^2} \cdot \frac{\mu_{t-1} - r_{t-1}^f}{\sigma_{t-1}^2}, \tag{2}$$

where $w \in \mathbb{R}$ is a constant, related to the unconditional expected return on the strategy. By choosing w appropriately, one can now construct efficient strategies to track a given target expected return or target variance.

2.1 Properties of Efficient Strategies

From (2), it is clear that the conditional Sharpe ratio H_{t-1} plays a key role in the behavior of the optimal strategy. Abhyankar, Basu, and Stremme (2005b) show that the maximum (squared) *unconditional* Sharpe ratio, attainable by optimally managed efficient portfolio

[1] See Abhyankar, Basu, and Stremme (2005b), or Ferson and Siegel (2001).

strategies, can be written as $\lambda_*^2 = E(H_{t-1}^2)$. In other words, the squared unconditional Sharpe ratio is given by the unconditional second moment of the conditional Sharpe ratio[2]. Consequently, time-variation in the conditional Sharpe ratio improves the ex-post risk-return trade-off for the mean-variance investor, a point also noted by Cochrane (1999). See also Section 3 and Figures 2 and 3 for an illustration.

For small values of $\mu_{t-1} - r_{t-1}^f$, the efficient weights in (2) respond almost linearly to changes in μ_{t-1}, shifting more money into the index portfolio the higher its expected return relative to the Treasury bill. However, for extreme values of μ_{t-1}, the behavior of the weights is dominated by the denominator $1 + H_{t-1}^2$, forcing the asset allocation back towards the risk-free asset. This creates a 'conservative response' to extreme signals, as observed also by Ferson and Siegel (2001). Note that the corresponding conditionally efficient strategy is missing the normalization factor $1 + H_{t-1}^2$ and thus tends to 'over-react' to extreme values of the conditioning instrument.

To shed additional light on the behavior of the efficient weights, consider for the moment an investor who chooses an optimal asset allocation such as to maximize conditional quadratic utility. The unconditionally efficient allocation (2) then corresponds to a conditional risk aversion coefficient that is proportional to $1 + H_{t-1}^2$. In other words, the unconditionally efficient market timing strategy corresponds to a conditionally optimal strategy for an investor with time-varying risk aversion. In particular, the implied conditional risk aversion co-efficient increases when the conditional expected return μ_{t-1} takes on extreme values, thus causing the strategy to respond more conservatively to extreme information. In contrast, the conditionally optimal strategy for constant risk aversion tends to 'over-react' to extreme signals. In other words, the portfolio weights of a conditionally efficient strategy tend to be more volatile than those of the corresponding unconditionally efficient strategy, an important consideration in particular in view of transaction costs (see also Section 3).

[2]This result holds even in the case of multiple risky assets. In the case of a single risky asset, this was shown by Jagannathan (1996).

2.2 Modeling Return Predictability

Although the theoretical results presented in the preceding sections are valid also in much more general settings, for our empirical analysis (see Section 3) we will restrict ourselves to a simple linear specification. More specifically, we assume that the return on the market index portfolio is described by a linear predictive model of the form,

$$r_t^M = \mu_0 + By_{t-1} + \varepsilon_t, \tag{3}$$

where y_{t-1} is a (vector of) lagged predictive instruments, and ε_t is an *iid* sequence of disturbances. In this setting, the conditional expectation of r_t^M is given by $\mu_{t-1} = \mu_0 + By_{t-1}$, and the conditional variance is constant, $\sigma_{t-1}^2 = \sigma^2(e_t)$ due to the *iid* assumption. For notational convenience, we normalize the instruments y_{t-1} to have zero mean, so that the unconditional expected return on the market index is $E(\, r_t^M \,) = \mu_0$.

To assess the economic value of optimal market timing, we measure the extent to which the optimal use of return predictability extends the unconditionally efficient frontier and thus the opportunity set available to the mean-variance investor. As a benchmark, denote by λ_0 the maximum Sharpe ratio of a buy-and-hold strategy, i.e.

$$\lambda_0 = \frac{E(\, r_t^M - r_{t-1}^f \,)}{\sigma(\, r_t^M \,)} \tag{4}$$

As noted above, the maximum (squared) unconditional Sharpe ratio, attainable by optimally managed strategies, is given by $\lambda_*^2 = E(\, H_{t-1}^2 \,)$. Thus, the economic gain of optimal market timing can be measured by the difference $\Omega := \lambda_*^2 - \lambda_0^2$ in squared Sharpe ratios with and without the optimal use of conditioning information. Our null hypothesis is that predictability has no effect, i.e. $\Omega = 0$.

Abhyankar, Basu, and Stremme (2005a) show that under the null hypothesis, the test statistic $T \cdot \Omega$ (where T is the number of time series observations) has an F-distribution in finite samples, and a χ^2 distribution (with one degree of freedom in the case of a single risky asset) asymptotically. This enables us to assess whether the increase in Sharpe ratio due to the optimal use of asset return predictability is statistically significant. Moreover, it is straight-forward to show that under the null hypothesis, we have

$$\Omega = \lambda_*^2 - \lambda_0^2 = \frac{R^2}{1 - R^2}, \tag{5}$$

where R^2 is the coefficient of determination in the predictive regression (3). In other words, the measure Ω of the *economic* value of predictability is directly associated with the statistical properties of the predictive regression.

3 Empirical Analysis

In this section, we briefly describe the empirical methodology and the data used, and discuss the results of our empirical analysis.

3.1 Data and Methodology

For our empirical analysis, we use monthly return data covering the period from January 1960 to December 2003. As the single risky asset r_t^M, we use the total return on the CRSP value-weighted market index. As the (conditionally) risk-free asset r_{t-1}^f, we use the return on the corresponding 1-month US Treasury bill.

We categorize the predictive instruments used into two groups, capturing (a) changes in the level and shape of the term structure of interest rates, and (b) macro-economic indicators. The former group consists of the current short rate (we use the 1-month Treasury bill rate as a proxy), the slope of the term structure (the yield spread between the 10-year bond and the 1-month bill), and a proxy for the convexity of the yield curve (the difference between twice the 5-year Treasury yield and the sum of the 1-month and 10-year rates). All data for this group are obtained from the Economic Database (FRED) at the Federal Reserve Bank of St. Louis[3]. In addition, we include the credit yield spread, defined as the difference in 10-year yield between AAA-rated corporate and government bonds, which we obtained from Datastream. The group of economic indicators includes inflation (the change in CPI over the period), and the rate of growth in aggregate consumption and the level of unemployment. All of these were constructed from data published by the Federal Reserve Bank of St. Louis.

For each group of instruments, we estimate the predictive regression (3), and compute the implied maximum fixed-weight and optimally managed Sharpe ratios using the expressions from Section 2. We then construct the corresponding conditionally and unconditionally efficient market-timing strategies using (2), and assess their performance using a variety of standard *ex-post* performance measures. For the out-of-sample analysis, we use the first 20 years of data to estimate (3), construct the weights of the efficient strategies based on the model estimates, and then assess their performance on the basis of their realized returns in the out-of-sample period.

3.2 In-Sample Results

We first analyze whether *any* market-timing strategy can work by estimating the model using all predictive instruments and computing the test statistic Ω defined in the preceding section. The results are summarized in the last column of Table 1. While the maximum Sharpe ratio without conditioning information ('fixed-weight') is only 0.37, this almost doubles to 0.73 when all predictive variables are used. The p-value of our test statistic is indistinguishable from zero, indicating that the result is significant at any level of confidence. Thus it is clear that our market timing strategy performs well in-sample. Even when the two groups of instruments are used separately (columns 2 and 3 in Table 1) the increase in Sharpe ratio (to 0.57 and 0.59, respectively), though less dramatic, is still significant at the 1% level. Interestingly, the economic gain from optimal market-timing is considerable (and

[3] `http://research/stlouisfed.org`

significant), even though the joint R^2 of the predictive regression is less than 5% in all cases.

Next we focus on the performance of the optimally managed market-timing strategy, designed to maximize average returns while tracking a target volatility of 15% annually (which approximately reflects the volatility of the benchmark index over the sample period). The results are shown in Table 2. When all instruments are used, the optimally managed maximum-return strategy achieves an (annualized) mean of 17.2%, while the corresponding efficient static (fixed-weight) portfolio barely matches the benchmark return (11.7%). Because the optimal strategy matches the target volatility quite closely, the Sharpe ratio approximately doubles from about 0.36 to 0.70. Overall, the *ex-post* performance of the efficient strategies comes very close to the (theoretical) Sharpe ratios implied by the model estimates (as shown in Table 1).

We also computed the *ex-post* Jensen's alphas, tracking errors and information ratios of our strategies, relative to the market benchmark. Unsurprisingly, the fixed-weight strategy tracks the benchmark very closely (with a tracking error of 1%), but also barely matches its performance. In contrast, our optimal strategies deviate substantially from the benchmark (with tracking errors of between 12 and 13.5%), but at the same time generate alphas of up to 8.5%. More importantly, our optimal market-timing strategies also provide a considerable amount of portfolio insurance (without options): while the optimal fixed-weight strategy actually performs less well in recessions, the unconditionally efficient strategy achieves recession alphas in excess of 20%. This is illustrated in Figure 1, which shows the cumulative return of the two types of strategies, relative to the benchmark: while the optimal strategy participates in the up-swings of the benchmark, it hardly suffers when the benchmark takes a dip (note that the scale of the graph is logarithmic).

When we constrain the strategy to have non-negative weights the mean drops to 12.07% but the volatility is reduced to 8.48%. Overall, the Sharpe ratio drops only slightly to 0.69 (from 0.73). This suggests that the unconstrained strategy does not take extreme long or short positions (this is also confirmed by Figure 4) and illustrates the conservative response of these portfolio weights to extreme values of the predictive variables, a property first observed in

Ferson and Siegel (2001).

Focusing on just the term-structure variables, our test statistic has a p-value of 0.38%, showing that the increase in Sharpe ratio due to optimal market-timing using these variables is highly significant, although the corresponding maximum-return strategy has a lower mean (but also lower volatility) than that using all predictive variables. Similarly, the p-value when only the macro-economic variables are used is 0.21%. In summary, both groups of indicators clearly have significant predictive ability, and are thus clearly useful for optimal market-timing.

3.3 How do the Strategies Work?

A closer analysis of the portfolio weights in (2) reveals that the conditional Sharpe ratio plays a key role in the performance of these strategies. As the squared unconditional Sharpe ratio λ_*^2 is the expectation of the squared conditional Sharpe ratio $H_{t-1}^2 = (\mu_{t-1} - r_{t-1}^f)^2/\sigma_{t-1}^2$, time-variation in the conditional Sharpe ratio is good for a mean-variance investor. Figure 2 shows the response of the efficient market-timing weights to changes in the conditional Sharpe ratio H_{t-1} and the conditional expected return μ_{t-1} of the market index. The graph shows that the magnitude of the optimal position in the risky asset is driven by H_{t-1}, while the *direction* of the position is determined by the expected return μ_{t-1}. This demonstrates that the optimally managed strategy exhibits a more measured behavior than the 'all-or-nothing' switching-strategies common in the market-timing literature. In other words, while the conditional expected index return μ_{t-1} tells the portfolios manager whether to move into our out of the market, the conditional Sharpe ratio H_{t-1} drives the optimal *risk control* for the strategy. The importance of the conditional Sharpe ratio in optimal asset allocation was also emphasized by Cochrane (1999).

Figure 3 compares the behavior of the optimal market-timing strategy in response to variations in conditional Sharpe ratio, for 'good' and 'bad' predictive instruments[4]. Evidently, a 'good' predictor generates more variability in the conditional Sharpe ratio, but the optimal market-timing weights respond more conservatively to these variations. In other words, when the instrument used does not possess significant predictive ability, the optimal strategy tends to 'over-react' to spurious signals, thus unduly increasing both risk and transaction costs. The choice of predictive instruments is therefore of great importance in the design of successful market-timing strategies.

3.4 Comparison with Conditionally Efficient Strategies

As we discussed in Section 2.1, our unconditionally efficient strategies are by construction also conditionally efficient, while the converse is not generally true. While our strategies are thus theoretically optimal, it is nonetheless important to compare their performance with that of the corresponding conditionally efficient strategies. The latter are constructed using the conditional mean and covariance matrix of the assets, based on the realized values of the predictive variables, so that their weights are not an ex-ante prescribed function of the instrument. In other words, while the unconditionally efficient strategy is truly dynamic, the conditionally efficient strategy is a concatenation of period-by-period static portfolios.

While the functional form of the weights of conditionally efficient strategies is similar to (2), it lacks the dynamic normalization by the conditional Sharpe ratio H_{t-1}. As a consequence, the conditionally efficient weights respond much more aggressively to changes in the predictive instruments (see Figure 4 for an illustration). In particular, for small changes in the instrument around its mean, the conditionally efficient weights tend to switch dramatically between long and short positions. In our empirical study, conditionally efficient strategies not only under-performed their unconditionally efficient counterparts, but also incurred sig-

[4]Based on our estimation results, we used the lagged market return as the 'bad' instrument, and the family of term structure variables as 'good' predictors.

nificantly higher transaction costs. Moreover, because of the extreme switching behavior, the performance of conditionally efficient strategies is much more sensitive to the imposition of short-sale constraints.

3.5 Out-of-Sample Analysis

Our analysis in the previous section focused on the statistical evidence for return predictability and its implications for optimal market-timing. In this section we analyze whether an investment manager could actually make money from this strategy. To that end we perform an out-of-sample analysis, estimating the parameters of the data-generating process using the first 20 years of data (1960-1980), constructing the optimal market-timing strategy on the basis of these estimates, and then assessing its performance on the basis of the returns realized in the out-of-sample period.

The results are reported in Table 3. While the optimal Sharpe ratios are slightly lower out-of-sample, they are still more than 80% higher than the fixed-weight Sharpe ratio. While the performance of the fixed-weight strategy barely matches the benchmark (with an alpha of 0.5%), our strategy beats the index by more than 13%, indicating that a portfolio manager following this strategy could have made significant gains over this period.

4 Conclusion

This paper provides both a statistical test to determine whether any market timing strategy using predictive variables will work as well as an optimal market timing strategy. Using the CRSP value-weighted index and commonly used macro-economic and interest rate variables we find that market timing over the 1960-2003 period could have lead to significant economic gains. Our optimal market timing strategy out-performed the benchmark both in and out of sample, suggesting that a portfolio manager following this strategy could have made considerable gains over this period.

References

ABHYANKAR, A., D. BASU, and A. STREMME (2005a): "The Optimal Use of Asset Return Predictability: An Empirical Analysis," working paper, Warwick Business School.

ABHYANKAR, A., D. BASU, and A. STREMME (2005b): "Portfolio Efficiency and Discount Factor Bounds with Conditioning Information: A Unified Approach," working paper, Warwick Business School.

COCHRANE, J. (1999): "Portfolio Advice for a Multi-Factor World," *Economic Perspectives, Federal Reserve Bank of Chicago*, 23(3), 59–78.

FERSON, W., and A. SIEGEL (2001): "The Efficient Use of Conditioning Information in Portfolios," *Journal of Finance*, 56(3), 967–982.

JAGANNATHAN, R. (1996): "Relation between the Slopes of the Conditional and Unconditional Mean-Standard Deviation Frontier of Asset Returns," in *Modern Portfolio Theory and its Applications: Inquiries into Asset Valuation Problems*, ed. by S. Saito et al. Center for Academic Societies, Osaka, Japan.

	Fixed-Weight	Optimally Managed		
		Interest Rates	Economic Indicators	All Instruments
Sharpe Ratio	0.3719	0.5675	0.5908	0.7285
p-Value		0.0038**	0.0021**	0.0000**
Instrument	Coefficient (R^2)			
Short Rate		0.4223		1.1751
		(0.0004)		(0.0035)
Term Spread		1.4441		1.8806
		(0.1159)		(0.1848)
Convexity		−0.0143		−0.0185
		(0.0345)		(0.0580)
Credit Yield Spread		−0.2903		−0.5691
		(0.0087)		(0.0334)
Inflation			−0.1282	−0.1437
			(0.0123)	(0.0157)
Consumption Growth			0.1804	0.1753
			(0.0005)	(0.0005)
Unemployment Growth			0.1417	0.1566
			(0.0088)	(0.0110)
Maximum R^2		0.0165	0.0187	0.0355

Table 1: Model Estimation Results (In-Sample)

This table reports the in-sample estimation results for the efficient fixed-weight strategy (first column), and the optimally managed market-timing strategies for 3 different sets of instruments. Annualized Sharpe ratios are computed using the explicit expressions developed in Section 2, based on the parameter of the predictive regression (3), estimated using the entire data sample. The corresponding p-values are obtained from the χ^2 distribution, as discussed in Section 2. The asterisks indicate significance at 5% ($*$) and 1% ($**$) level. Also reported are the regression coefficients for each instrument (the individual R^2 are shown in parentheses beneath each coefficient), and the aggregate R^2 of the entire regression.

	Benchmark	Fixed-Weight	Optimally Managed		
			Interest Rates	Economic Indicators	All Instruments
Mean Return	11.68%	11.69%	14.71%	14.99%	17.15%
Volatility	15.43%	14.92%	14.47%	15.24%	14.75%
Sharpe Ratio	0.3564	0.3695	0.5679	0.5568	0.7036
	Performance Relative to Benchmark				
Jensen's Alpha		0.21%	5.49%	5.44%	8.45%
Tracking Error		1.00%	12.03%	12.35%	13.34%
Information Ratio		0.2100	0.4559	0.4401	0.6334
	Business Cycle Performance				
Recession Alpha		0.12%	11.73%	15.98%	20.31%
Expansion Alpha		0.22%	4.46%	3.73%	6.54%

Table 2: Portfolio Performance (In-Sample)

This table reports the *ex-post* performance of the efficient fixed-weight strategy (first column), and the optimally managed market-timing strategies for 3 different sets of instruments. The strategies were constructed using (2), based on the parameters of the predictive regression (3), estimated using the entire data sample. The performance of the strategies was evaluated on the basis of their realized returns throughout the same sample period. Means, volatilities and Sharpe ratios are annualized. Jensen's alpha, tracking error and information ratio are obtained from a CAPM-style regression of portfolio returns on the benchmark returns. Recession and expansion alphas were calculated by replacing the constant in this regression by two dummy variables indicating recession and expansion periods, respectively.

	Fixed-Weight	Optimally Managed		
		Interest Rates	Economic Indicators	All Instruments
	In-Sample Estimates			
Sharpe Ratio	0.4574	0.6168	0.5329	0.6221
p-Value		0.0160**	0.1113	0.0141**
	Out-of-Sample Portfolio Performance			
Sharpe Ratio	0.3524	0.5902	0.4446	0.6039
Alpha	0.50%	12.48%	6.65%	13.53%

Table 3: Portfolio Performance (Out-of-Sample)

This table reports the *out-of-sample* performance of the efficient fixed-weight strategy (first column), and the optimally managed market-timing strategies for 3 different sets of instruments. The strategies were constructed using (2), based on the parameters of the predictive regression (3), estimated in-sample (using the first 20 years of data). The performance of the strategies was evaluated on the basis of their realized returns throughout the out-of-sample period. Means, volatilities and Sharpe ratios are annualized. Jensen's alpha, tracking error and information ratio are obtained from a CAPM-style regression of portfolio returns on the benchmark returns. Recession and expansion alphas were calculated by replacing the constant in this regression by two dummy variables indicating recession and expansion periods, respectively.

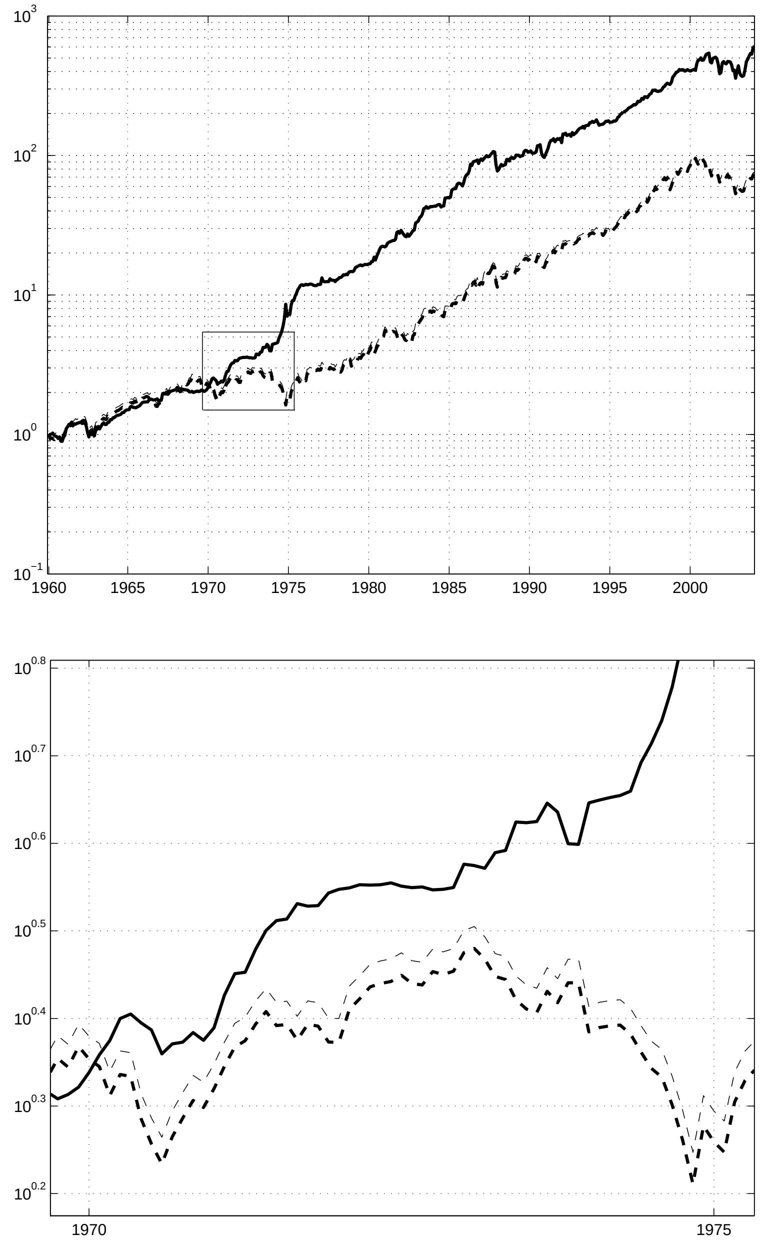

Figure 1: Cumulative Returns

This figure shows the cumulative return (future value of $1 invested at the beginning of the sample period) of the unconditionally efficient market-timing strategy (solid line), compared with the benchmark and the corresponding fixed-weight strategy (dashed lines). The top panel shows the returns over the entire sample period, while the bottom panel is a magnification of the 1970 to 1975 period as indicated in the top graph. Note that the scale is logarithmic.

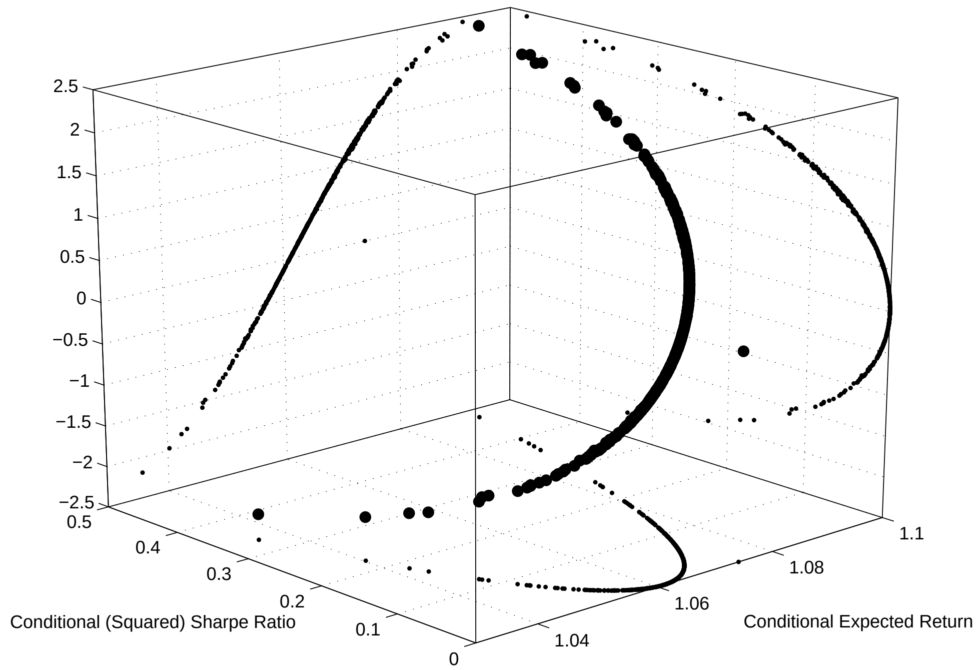

Figure 2: Efficient Weights

This figure shows the optimally managed weights θ_{t-1} on the risky asset, as a function of the conditional mean μ_{t-1} and the conditional Sharpe ratio H_{t-1}^2. The projections of the graph onto the 'walls' of the diagram show the relation between any two of the variables, respectively. The weights were constructed using (2), based on the parameters of the predictive regression (3), estimated using the entire data sample. All predictive variables are used as instruments.

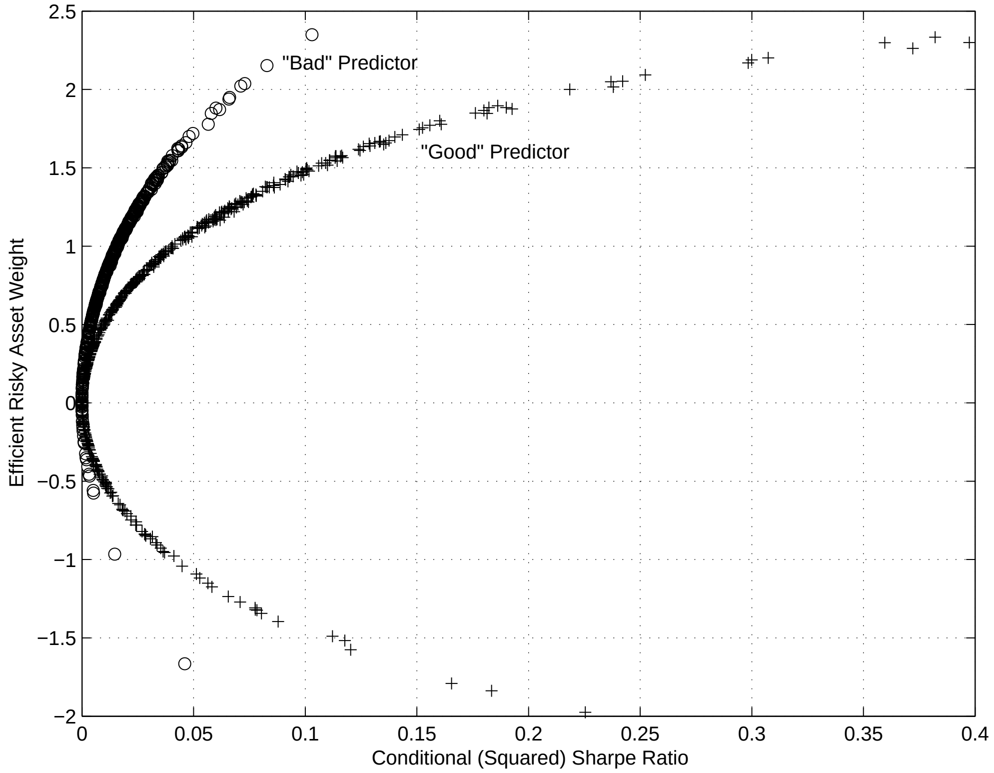

Figure 3: Efficient Weights ('Good' and 'Bad' Predictor

This figure shows the optimally managed weights θ_{t-1} on the risky asset, as a function of the conditional Sharpe ratio H^2_{t-1}, using either lagged market returns ('o'), or the set of term structure variables ('+') as predictive instruments. The weights were constructed using (2), based on the parameters of the predictive regression (3), estimated using the entire data sample.

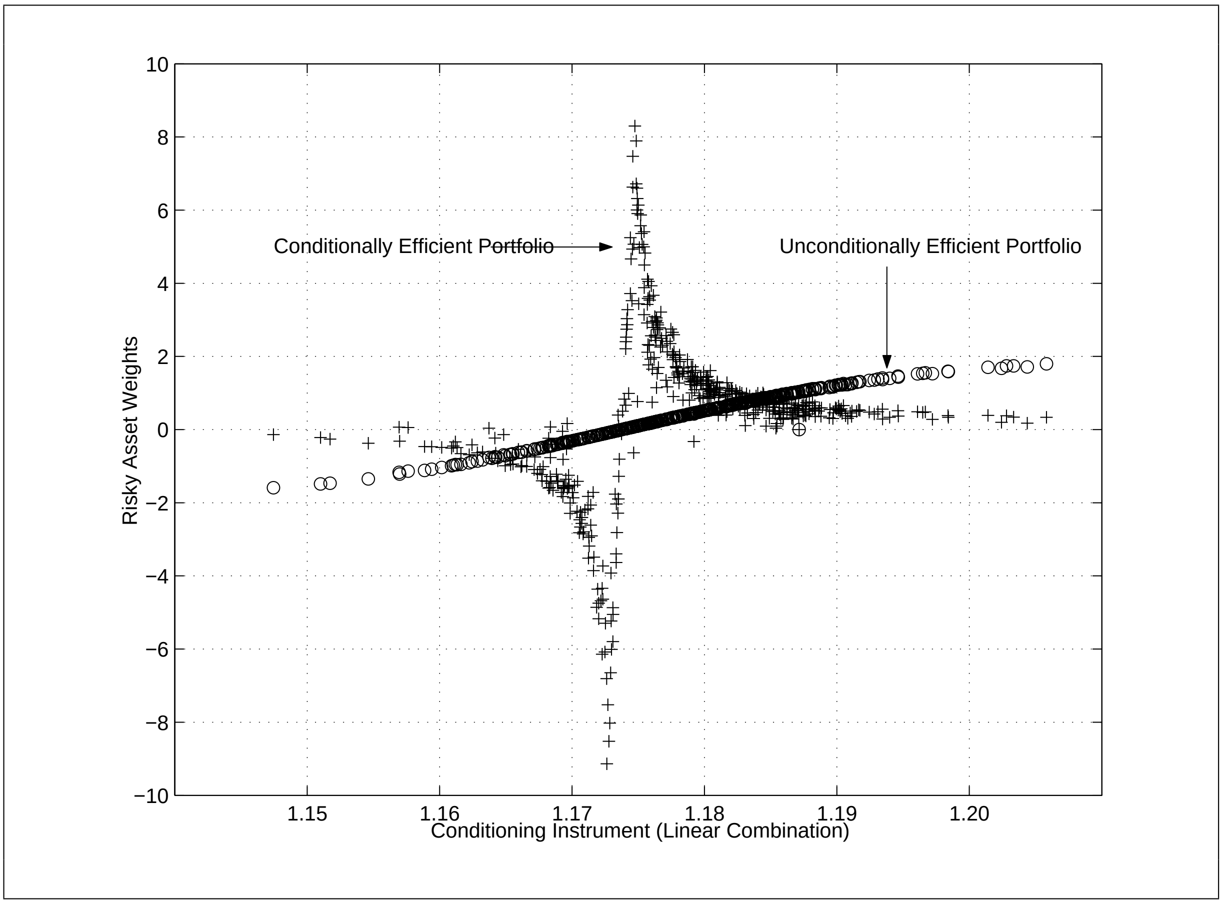

Figure 4: Efficient Weights (Conditional and Unconditional)

This figure shows the weights θ_{t-1} of the conditionally efficient ('+') and unconditionally efficient ('o') strategies, as functions of the linear combination By_{t-1} of the conditioning instruments. The weights were constructed using standard mean-variance theory for the conditionally efficient strategies, and (2) for the unconditionally efficient ones, based on the parameters of the predictive regression (3), estimated using the entire data sample.