Essays in Empirical Finance

by

Pasquale Della Corte

A thesis submitted in partial fulfilment of
the requirements for the degree of
Doctor of Philosophy in Finance

Warwick Business School
University of Warwick

October 2007

Contents

List of Tables

List of Figures

Acknowledgements

This thesis is the outcome of four years of extensive yet exciting and instructive research. I am grateful to many persons who have challenged my own thinking and thereby have influenced the aim and meaning of this work. Without their help, trust and encouragement, I would never have been able to complete this wonderful experience.

First of all, my deepest praise and most sincere gratitude go to my supervisors, Lucio Sarno and Ilias Tsiakas, for their wisdom, inspiration and endless enthusiasm throughout the whole work with this thesis. Their insightful comments and fruitful research discussions have shaped to a large degree my understanding and their guidance, mentorship and support have been invaluable. This thesis owes an enormous debt — personal and intellectual — to them.

My family deserves special mention for their unwavering support and for feeling proud of me. My father, Gennaro, is the person who planted the seed of my learning character ever since I was a child. My mother, Anna Maria, is the one who truly raised me with her caring and gentle love. They taught me the good things that really matter in life and convey to me the ambition to be always one step further. This thesis is dedicated to them. Moreover, I would like to deeply thank Sophia, for her patience and love and for brightening my mind from the first moment until the last day of this effort.

I also wish to thank Giorgio Valente and Dan Thornton for their generous advice. I greatly appreciate their support. Special thanks go to Gaetano Antinolfi and the Economics Department at Washington University in St. Louis where this thesis was partly written while I was visiting scholar in the last stage of my studies. My appreciation also extends to the Economics Department at the University of Brescia, especially to Franco Spinelli, Gianni Amisano and Carmine Trecroci for their support in the earlier period of this adventure.

I would also like to thank Andrea, Elvira, Francesco, Gastón, Giorgio, Kleopatra, Srikanth, Stefano, Suraj, and Wing for their friendship, advice and countless hours of conversations. I am also indebted to my colleagues at Warwick Business School and the Washington University in St. Louis for a good working atmosphere. I also need to mention Andrea, Carlo, Domenico and Vincenzo for their faithful friendship and encouragement whenever I turned to them.

I am also grateful for useful conversations or constructive comments to William Schwert (editor of *Journal of Financial Economics*), Joel Hasbrouck (editor of *Review of Financial Studies*), two anonymous referees, Abhay Abhyankar, David

Backus, Luc Bauwens, Peter Corvi, Frank Diebold, Massimo Guidolin, Jim Hamilton, Xing Jin, Rich Lyons, Marco Lyrio, Michael Moore, Roel Oomen, Carol Osler, Dagfinn Rime, Thomas Sargent, Nick Webber and Arnold Zellner, as well as to the organisers, Sid Chib and Ed Greenberg, and the participants at the seminar on Bayesian Inference in Econometrics and Statistics at the Olin Business School, the participants at the Finance Workshop at Warwick Business School, the 2006 European Science Foundation workshop on FX Markets at the University of Warwick, the 2006 European meeting of the Econometric Society, the EC^2 Conference on the Econometrics of Monetary Policy and Financial Decision Making at Rotterdam, the seminars at the European Central Bank and the International Monetary Fund, the Seventh Annual Missouri Conference at the University of Missouri-Columbia, the Policy Meeting and Macro System Committee Meeting at the Federal Reserve Bank of St. Louis, and the 2007 European meeting of the Econometric Society.

Finally, I gratefully acknowledge financial support from Warwick Business School, the Economic and Social Research Council (ESRC), and the University of Brescia.

Declaration

I declare that any material contained in this thesis has not been submitted for a degree to any other University. I further declare that one paper entitled: "The Expectation Hypothesis of the Term Structure of Very Short-Term Rates: Statistical Tests and Economic Value", drawn from Chapter One and Chapter Two of this thesis has been accepted for publication and is forthcoming in the *Journal of Financial Economics.* Furthermore, the paper "An Economic Evaluation of Empirical Exchange Rate Models: Robust Evidence of Predictability and Volatility Timing", drawn from Chapter Three is under review at the *Review of Financial Studies* (revise and resubmit).

Pasquale Della Corte

August 2007

Abstract

The aim of this thesis is to deepen our understanding of new empirical methods, results and implications in interest rate and foreign exchange markets. To this end, this thesis is organised in three chapters.

The first chapter tests the validity of the Expectation Hypothesis (EH) of the term structure using daily data for US repo rates spanning the 1991-2005 sample period and ranging in maturity from overnight to three months. We revisit a recent study by Longstaff (2000a) by implementing statistical tests designed to increase test power in this context. Specifically, we apply the Lagrange Multiplier and Distance Metric statistics to test a set of nonlinear cross-equation restrictions imposed by the EH on a vector autoregression model of the short- and long-term interest rates. We find that EH is rejected throughout the term structure examined on the basis of the statistical tests.

In the second chapter, we extend the study carried out in the first chapter in a different direction and assess the economic value of departures from the EH based on criteria of profitability and economic significance. In the context of a mean-variance framework, we compare the performance of a dynamic portfolio strategy consistent with EH to a dynamic portfolio strategy that exploits the departures from the EH. The results of our economic analysis are favourable to the EH, suggesting that the statistical rejections of the EH in the repo market are economically insignificant.

Finally, in the third chapter, we provide a comprehensive evaluation of the short-horizon predictive ability of economic fundamentals and forward premia on monthly exchange rate returns in a framework that allows for volatility timing. We implement Bayesian methods for estimation and ranking of a set of empirical exchange rate models, and construct combined forecasts based on Deterministic and Bayesian Model Averaging. More importantly, we assess the economic value of the in-sample and out-of-sample forecasting power of the empirical models, and find two key results: (i) a risk averse investor will pay a high performance fee to switch from a dynamic portfolio strategy based on the random walk model to one which conditions on the forward premium with stochastic volatility innovations; and (ii) strategies based on combined forecasts yield large economic gains over the random walk benchmark. These two results are robust to reasonably high transaction costs.

Overview

This thesis investigates new empirical methods, challenges established findings on classic issues of empirical finance and provides new results and implications in interest rate and foreign exchange markets.

The first chapter re-examines the validity of the Expectation Hypothesis (EH) of the term structure of the interest rates. Ever since Fisher (1896) postulated the Expectation Hypothesis (EH) of the term structure of interest rates, this simple and intuitively appealing theory has attracted an enormous amount of attention in financial economics. Many authors have argued that interest rates at different maturities move together because they are linked by the EH and a number of studies have addressed the empirical validity of this theory. However, this literature, using a variety of tests and data, generally rejects the EH (e.g. Roll, 1970; Fama, 1984b; Fama and Bliss, 1987; Frankel and Froot, 1987; Stambaugh, 1988; Froot, 1989; Campbell and Shiller, 1991; Bekaert, Hodrick and Marshall, 1997; Bekaert and Hodrick, 2001; Clarida, Sarno, Taylor and Valente, 2006; Sarno, Thornton and Valente, 2007).

An important exception is provided by Longstaff (2000a), who finds that the EH is supported by the data. Longstaff (2000a) presents the first tests of the EH at the extreme short end of the term structure, using repurchase (repo) rates with maturities measured in days or weeks. There are two reasons why Longstaff's study is important. First, if the EH cannot explain the term structure at this extreme short end, it seems unlikely that it can be of value at longer maturities. Second, the use of repo rates is especially appropriate for investigating the EH because repo rates represent the actual cost of holding riskless securities. Hence, repo rates provide potentially better measures of the short-term riskless term structure than other interest rates commonly used by the relevant literature, such as Treasury bill rates.

This chapter revisits the EH using an updated data set of repo rates from the same source as Longstaff (2000a). In fact, the literature on testing the EH has made much progress in recent years by developing increasingly sophisticated testing procedures that are particularly useful in this context. Given the statistical problems afflicting conventional tests of the EH, in this chapter we employ a test that was originally proposed in Campbell and Shiller (1987) and made operational in Bekaert and Hodrick (2001). Bekaert and Hodrick (2001) develop a procedure for testing the parameter restrictions that the EH imposes on a vector autoregression (VAR) of the short- and long-term interest rates. The procedure's size and power properties have been thoroughly investigated by Bekaert and Hodrick (2001) and Sarno, Thornton and Valente (2007). We apply this test to US repo rates ranging in maturity from overnight to three months over the sample period 1991-2005.

To anticipate the results of the first chapter, we find that the EH is statistically rejected for all pairs of repo rates in our sample throughout the maturity spectrum from overnight to three months. Our results differ from Longstaff's (2000a) presumably because the VAR test is more powerful and our sample period is somewhat longer than his.

In the second chapter, we extend the study carried out in the first chapter. We move beyond testing the validity of the EH from a purely statistical perspective and provide evidence on whether deviations from the EH are economically significant. Distinguishing between statistical analysis and economic evaluation is crucial for at least three reasons: in general statistical rejections of a hypothesis do not necessarily imply economic rejections (Leitch and Tanner, 1991); statistical VAR tests of the EH do not allow for transactions costs, which are critical for exploiting departures from the EH in real-world financial markets; and very powerful statistical tests may reject virtually any null hypothesis in large samples, without necessarily being informative about the size of departures from the hypothesis tested (Leamer, 1978). All these reasons suggest that an economic assessment of the deviations from the EH is desirable to complement the statistical tests.

In a mean-variance framework, we compare the performance of a dynamic portfolio strategy consistent with the EH to a dynamic portfolio strategy that exploits the departures from the EH. We use a utility-based performance criterion to compute the fee a risk-averse investor would be willing to pay to switch from the EH to a strategy that exploits departures from the EH to forecast interest rates. As an alternative economic measure, we also employ the risk-adjusted return of these two strategies. In short, we provide an economic test of the EH by evaluating the incremental profitability of an optimal (mean-variance efficient) strategy which relaxes the restrictions implied by the EH statement.

To anticipate the results of the second chapter, the economic analysis lend support to the EH as we find no tangible economic gain to an investor who exploits departures from the EH relative to an investor who allocates capital simply on the basis of the predictions of the EH. Specifically, the evidence in this chapter shows that the economic value of departures from the EH is modest and generally smaller than the costs that an investor would incur if he were to trade to exploit the mispricing implied by EH violations. Hence, despite the statistical rejections of the EH recorded in the previous chapter, we conclude that the EH provides a fairly reasonable approximation to the repo rates term structure, consistent with Longstaff's interpretation of the functioning of the repo market.

In the third chapter, we provide a comprehensive evaluation of the short-horizon predictive ability of economic fundamentals and forward premia on monthly exchange rate returns in a framework that allows for volatility timing. Forecasting exchange rates using models which condition on economically meaningful variables has long been at the top of the research agenda in international finance, and yet empirical success remains elusive. Starting with the seminal contribution of Meese and Rogoff (1983), a vast body of empirical research finds that models which condition on economic fundamentals cannot outperform a naive random walk model. Even though there is some evidence that exchange rates and fundamentals comove over long horizons (e.g. Mark, 1995; Mark and Sul, 2001), the prevailing view in international finance research is that exchange rates are not predictable, especially at short horizons.

A separate yet related literature finds that forward exchange rates contain valuable information for predicting spot exchange rates. In theory, the relation between spot and forward exchange rates is governed by the Uncovered Interest Parity (UIP) condition, which suggests that the forward premium must be perfectly positively related to future exchange rate changes. In practice, however, this is not the case as we empirically observe a negative relation. The result of the empirical failure of UIP is that conditioning on the forward premium often generates exchange rate predictability. For example, Backus, Gregory and Telmer (1993) and Backus, Foresi and Telmer (2001) explore this further and find evidence of predictability using the lagged forward premium as a predictive variable. Furthermore, Clarida, Sarno, Taylor and Valente (2003, 2006) and Boudoukh, Richardson and Whitelaw (2006) show that the term structure of forward exchange (and interest) rates contains valuable information for forecasting spot exchange rates.

On the methodology side, while there is extensive literature on statistical measures of the accuracy of exchange rate forecasts, there is little work assessing the economic value of exchange rate predictability. Relevant research to date comprises an early study by West, Edison and Cho (1993) which provides a utility-based evaluation of exchange rate volatility, and more recently, Abhyankar, Sarno and Valente (2005) who use a similar method for investigating long-horizon exchange rate predictability. However, in the context of dynamic asset allocation strategies, there is no study assessing the economic value of the predictive ability of empirical exchange rate models which condition on economic fundamentals or the forward premium while allowing for volatility timing.

Our empirical investigation attempts to fill this gap and connect the related literatures which examine the performance of empirical exchange rate models. We do this by employing a range of economic and Bayesian statistical criteria for performing a comprehensive assessment of the short-horizon, in-sample and out-of-sample, predictive ability of three sets of models for the conditional mean of monthly nominal exchange rate returns. These models include the naive random walk model, the monetary fundamentals model (in three variants), and the spot-forward regression model. Each of the models is studied under three volatility specifications: constant variance (standard linear regression), GARCH(1,1) and stochastic volatility (SV). In total, we evaluate the performance of 15 specifications, which encompass the most popular empirical exchange rate models studied in prior research. Our analysis employs monthly returns data ranging from January 1976 to December 2004 for three major US dollar exchange rates: the UK pound sterling, the Deutsch mark/euro, and the Japanese yen.

In addition to implementing Bayesian statistical methods for evaluating the models, an important contribution of our analysis is the use of economic criteria. Statistical evidence of exchange rate predictability in itself does not guarantee that an investor can earn profits from an asset allocation strategy that exploits this predictability. In practice, ranking models is useful to an investor only if it leads to tangible economic gains. Therefore, we assess the economic value of exchange rate predictability by evaluating the impact of predictable changes in the conditional foreign exchange (FX) returns and volatility on the performance of dynamic allocation strategies. We employ mean-variance analysis as a standard measure of portfolio performance and apply quadratic utility, which allows us to quantify how risk aversion affects the economic value of predictability, building on empirical studies of volatility timing in stock returns by Fleming, Kirby, and Ostdiek (2001) and Marquering and Verbeek (2004). Ultimately, we measure how much a risk averse investor is willing to pay for switching from a dynamic portfolio strategy based on the random walk model to one which conditions on either monetary fundamentals or forward premia and has a dynamic volatility specification.

Furthermore, we assess the statistical evidence on exchange rate predictability in a Bayesian framework. In particular, we rank the competing model specifications by computing the posterior probability of each model. The posterior probability is based on the marginal likelihood and hence it accounts for parameter uncertainty, while imposing a penalty for lack of parsimony (higher dimension). In the context of this Bayesian methodology, an alternative approach to determining the best model available is to form combined forecasts which exploit information from the entire universe of model specifications under consideration. Specifically, we implement the Bayesian Model Averaging (BMA) method, which weighs all conditional mean and volatility forecasts by the posterior probability of each model. We then compare the BMA results to those obtained from a Deterministic Model Averaging (DMA) strategy, which simply combines all model specifications with equal weights.

To preview the key results of the third chapter, we find strong economic and statistical evidence against the naive random walk benchmark with constant variance innovations. In particular, while conditioning on monetary fundamentals has no economic value either in-sample or out-of-sample, we establish that the predictive ability of forward exchange rate premia has substantial economic value in a dynamic portfolio allocation strategy, and that stochastic volatility significantly outperforms the constant variance and GARCH(1,1) models irrespective of the conditional mean specification. This leads to the conclusion that the best empirical exchange rate model is a model that exploits the information in the forward market for the prediction of conditional exchange rate returns and allows for stochastic volatility for the prediction of exchange rate volatility. We also provide evidence that combined forecasts which are formed using either DMA or BMA substantially outperform the random walk benchmark. These results are robust to reasonably high transaction costs and hold for all currencies both in-sample and out-of-sample. Finally, these findings have clear implications for international asset allocation strategies which are subject to FX risk.

1 A Statistical Evaluation of the Expectation Hypothesis of the Term Structure of Very Short-Term Rates

1.1 Introduction

Ever since Fisher (1896) postulated the Expectation Hypothesis (EH) of the term structure of interest rates, this simple and intuitively appealing theory has attracted an enormous amount of attention in financial economics. Many authors have argued that interest rates at different maturities move together because they are linked by the EH and a number of studies have addressed the empirical validity of this theory. However, this literature, using a variety of tests and data, generally rejects the EH (e.g. Roll, 1970; Fama, 1984b; Fama and Bliss, 1987; Frankel and Froot, 1987; Stambaugh, 1988; Froot, 1989; Campbell and Shiller, 1991; Bekaert, Hodrick and Marshall, 1997; Bekaert and Hodrick, 2001; Clarida, Sarno, Taylor and Valente, 2006; Sarno, Thornton and Valente, 2007).

An important exception is provided by Longstaff (2000a), who finds that the EH is supported by the data. Longstaff (2000a) presents the first tests of the EH at the extreme short end of the term structure, using repurchase (repo) rates with maturities measured in days or weeks. There are two reasons why Longstaff's study is important. First, if the EH cannot explain the term structure at this extreme short end, it seems unlikely that it can be of value at longer maturities. Second, the use of repo rates is especially appropriate for investigating the EH because repo rates represent the actual cost of holding riskless securities. Hence, repo rates provide potentially better measures of the short-term riskless term structure than other interest rates commonly used by the relevant literature, such as Treasury bill rates.

This chapter revisits the EH using an updated data set of repo rates from the same source as Longstaff (2000a). In fact, the literature on testing the EH has made much progress in recent years by developing increasingly sophisticated testing procedures that are particularly useful in this context. Given the statistical problems afflicting conventional tests of the EH, in this chapter we employ a test that was originally proposed in Campbell and Shiller (1987) and made operational in Bekaert and Hodrick (2001).[1] Bekaert and Hodrick (2001) develop a procedure for testing the parameter restrictions that the EH imposes on a vector autoregression (VAR) of the short- and long-term interest rates. The procedure's size and power properties have been thoroughly investigated by Bekaert and Hodrick (2001) and Sarno, Thornton and Valente (2007). We apply this test to US repo rates ranging in maturity from overnight to three months over the sample period 1991-2005.

To anticipate our results, we find that the EH is statistically rejected for all pairs of repo rates in our sample throughout the maturity spectrum from overnight to three months. Our results differ from Longstaff's (2000a) presumably because the VAR test is more powerful and our sample period is somewhat longer than his.

The outline of the chapter is as follows. Section 1.2 briefly describes the data and preliminary statistics on repo rates. Section 1.3 introduces the EH and the VAR framework within which the empirical work is carried out, with a description of the essential ingredients of the VAR testing procedure proposed by Bekaert and

[1] It is well known that tests that are commonly used to investigate the EH may generate paradoxical results due to finite sample biases, size distortions and power problems (e.g. see Campbell and Shiller, 1991; Bekaert, Hodrick and Marshall, 1997; Thornton, 2005, 2006).

Hodrick (2001). We report the results from the VAR tests of the EH in Section 1.4. The conclusions are presented in Section 1.5. Appendix A provides technical details on the VAR framework and estimation issues, in addition to further empirical results.

1.2 Data

The data set comprises daily observations of the closing overnight i_t, 1-week $i_t^{(1w)}$, 2-week $i_t^{(2w)}$, 3-week $i_t^{(3w)}$, 1-month $i_t^{(1m)}$, 2-month $i_t^{(2m)}$, and 3-month $i_t^{(3m)}$ general collateral government repo rates, from May 21, 1991 to December 9, 2005. The data are obtained from Bloomberg and the source of the data is Garban, a large Treasury securities broker. Repo rates are quoted on a 360-day basis and the rate quotations in Bloomberg are given in increments of basis points (*bps*). The total number of daily observations available is 3,625 and is essentially an update of the data set used by Longstaff (2000a).[2]

Table 1.1 reports the summary statistics for repo rates, in level and first difference. All variables are expressed in percentage points per annum. The data display similar properties to those described by Longstaff (2000a) for a shorter sample. The mean of the repo rates displays a mild smile effect across the term structure. In particular, the mean overnight rate of 3.9600 is slightly higher than the mean one-week rate of 3.9492, which turns out to be the lowest mean across the different maturities. The mean of three-month rate is 3.9924, which is approximately 3 *bps* higher

[2]Professor Longstaff kindly checked the consistency of our data set with the data used in Longstaff (2000a), which covered the sample from May 21 1991 to October 15 1999. Notice that only days for which a complete set of rates for all maturities are available are included in the sample. This resulted in 42 days being dropped from the sample. Finally, the period September 11, 2001 through September 30, 2001 is not available.

than the mean overnight rate. Table 1.1 also reports the mean repo rates for the different maturities by day of the week and shows a number of calendar regularities in the data. The mean repo tends to increase from Monday to Tuesday and to decrease afterwards, while the mean on Monday is always higher than the mean on Friday. For example, the mean overnight rate on Monday is 3.9718, which is about 5 *bps* higher than the mean overnight rate on Friday, equal to 3.9260. A similar pattern is observed for all other rates. However, it is important to note that these unconditional means are all very close to one another, and the differences are much smaller than the differences typically observed on other interest rates typically used in empirical research on the EH. For example, it is interesting to compare the means of repo rates to the means of Treasury bill (T-bill) rates. For comparison purposes, in Table 1.2 we report descriptive statistics on daily 1-month and 3-month US T-bill rates, also obtained from Bloomberg, both for a long sample from 1961 to 2005 and for the same sample as the repo rates data. The differences in the unconditional means between the 1–month and 3-month T-bill rates over the 1991-2005 sample are often about 15 *bps*, approximately five times larger than the maximum difference observed in repo markets for the same maturities. The differences in unconditional means for the full sample are even larger, up to 25 *bps*. Before embarking in our econometric analysis designed to test the EH, it is therefore worthwhile to note that the tiny differences in the unconditional means of repo rates at different maturities suggest that risk premia in repo markets are unlikely to be of particular economic importance. Put another way, these descriptive statistics are clearly indicative that the EH is more likely to hold on repo rates than T-bill rates.

We also report the standard deviations of daily changes in repo rates in Table 1.1. The overnight rate displays a standard deviation higher than the rates at other maturities. The standard deviation of daily changes in the overnight rate is about 18 *bps*, while the standard deviations for the other rates range from 5 to 6 *bps* per day. The standard deviations vary somewhat across days. The corresponding figures for T-bill rates, given in Table 1.2, indicate that changes in T-bill rates display a substantially higher dispersion than repo rates, with a standard deviation of about 16 *bps* for both 1-month and 3-month rates. However, it is worth mentioning that the standard deviation of the raw variables (annualised percentage returns) is not the standard deviation associated with an annual holding period. Therefore, we also report the annualised volatility $\sigma(a)$.[3] This battery of descriptive statistics confirms Longstaff's (2000a) argument that repo rates are smaller in magnitude and less volatile than T-bills.[4]

1.3 The Expectation Hypothesis

The EH of the term structure of interest rates relates a long-term n-period interest rate $i_t^{(n)}$ to a short-term m-period interest rate $i_t^{(m)}$. In the case of pure discount bonds, the EH can be stated as:

$$i_t^{(n)} = \frac{1}{k}\sum_{i=0}^{k-1} E_t[i_{t+mi}^{(m)}] + c^{(n,m)} \tag{1.1}$$

[3]Following Lo (2002), we compute the annualized volatility as $\sigma(a) = \sqrt{Var\left[i_t(a)\right]}$, where $i_t(a) = \sum_{k=0}^{a-1} i_{t-k}(d)$ is the sum of the daily returns, and $a = 250$ is the average number of trading days. Notice that the raw data are quoted on a 360-day basis and expressed in percentage points per annum. Hence, we determine the daily return as $i_t(d) = \frac{i_t}{360\times100}$ for a given raw repo rate i_t. We also report the product of the unconditional mean times the annualized volatility, $Mean \times \sigma(a)$, since this may be interpreted as the commonly used Black's volatility for caps under the assumption of log-normality.

[4]Notice also that the autocorrelation coefficients indicate a high level of persistence for all interest rates examined.

where $c^{(n,m)}$ is the term premium between the n- and m-period bonds (and may vary with the maturity of the rates); $k = n/m$ and is restricted to be an integer; and E_t denotes the mathematical expectation conditional on information set I_t available at time t.

In a market where expectations are formed rationally, an investor may either invest funds in a long-term n-period discount bond and hold it until maturity, or buy and roll over a sequence of short-term m-period discount bonds over the life of the long-term bond. Under the EH, these strategies should only differ by a constant term. As result, the long-term rate should be determined by a simple average of the current and expected future short-term rates plus a time-invariant term premium.[5] If the term premium $c^{(n,m)}$ is zero, the resulting form of the EH is often termed the 'pure' EH.

While much of the relevant literature relies on single equation tests of the EH, derived by reparameterising equation (1.1), a number of scholars reconsider the EH in a linear VAR framework and test the set of nonlinear restrictions which would make the VAR model consistent with the EH (Campbell and Shiller, 1991; Bekaert and Hodrick, 2001; Sarno, Thornton and Valente, 2007).[6] However, while the EH postulated in equation (1.1) is only a statement about how longer-term

[5]Fama (1984) derives equation (1.1) by assuming that the expected continuously compounded yields to maturity on all discount bonds are equal, up to a constant, while Shiller, Campbell, and Schoenholtz (1983) show that equation (1.1) is exact in some special cases and that it can be derived as a linear approximation to a number of nonlinear expectation theories of the term structure. For coupon bonds and consols with $n = \infty$, Shiller (1979) derives a similar linearized model where the long-term rate is a weighted average of expected future short-term rate plus a constant liquidity premium. Finally, note that, as showed by Longstaff (2000b), all traditional forms of the EH can be consistent with absence of arbitrage if markets are incomplete.

[6]The VAR methodology has been popular in the context of formulating and estimating dynamic linear rational expectations models since the 1970s, starting from Sargent (1977), Hansen and Sargent (1980), Sims (1980) and Wallis (1980).

rates are related to expected short term rates, the VAR setting further assumes a joint linear stochastic process for the dynamics of the long-term and short-term interest rates. This is a convenient assumption to extract predictions of future short-term rates by using current and past values of interest rates as information set. The VAR model is also inspired by the affine term structure literature in which conditional means are linear in a set of Markovian state variables (Duffie and Singleton, 1999; Dai and Singleton, 2000; Jagannathan, Kaplin and Sun, 2003; Ahn, Dittmar and Gallant, 2002; Bansal and Zhou, 2002; Clarida, Sarno, Taylor and Valente, 2006). This literature generally documents that affine specifications are unable to simultaneously match conditional means and conditional variances, leading to term premium puzzles.[7] Therefore, the linear VAR framework is rooted in a literature that has the potential to inherit some of the challenges faced by more traditional affine term structure models. This means that one may cannot rule out that the impact of these issues on EH tests based on the VAR framework may be substantial. For example, potential biases of the EH tests would arise if the interest rates data are generated by a process that is not encompassed within the VAR framework due to nonlinearities or time-varying covariances. In short, EH tests based on a VAR context are only valid under the maintained hypothesis that a linear VAR accurately describes the process of the short- and long-term interest rates and the relationship between them. This maintained assumption is questionable due to the well-documented limitations of affine specifications in matching the level and term premium in bonds simultaneously with the volatility of interest rates.

[7] Another stream of the literature also documents that affine structures cannot capture what is termed 'unspanned stochastic volatility' (e.g. Collin-Dufresne and Goldstein, 2002; Collin-Dufresne, Goldstein and Jones, 2007).

These caveats notwithstanding, in this chapter we rely on the VAR testing framework developed by Bekaert and Hodrick (2001) because of its desirable power properties in presence of highly nonlinear restrictions. Specifically, we implement the Generalised Method of Moments (GMM) to estimate a constrained VAR which forces the data to yield the relationship postulated by the EH and, then, test the validity of these restrictions by using the Lagrange Multiplier (LM) and Distance Metric (DM) statistics.[8]

1.3.1 The VAR Framework

Consider a bivariate VAR representation for the short- and long-term interest rates measured as deviations from their respective means:

$$i_t^{(m)} = a(L)i_{t-1}^{(m)} + b(L)i_{t-1}^{(n)} + u_{1,t} \tag{1.2}$$

$$i_t^{(n)} = c(L)i_{t-1}^{(m)} + d(L)i_{t-1}^{(n)} + u_{2,t} \tag{1.3}$$

where $a(L)$, $b(L)$, $c(L)$, and $d(L)$ are polynomials in the lag operator of order p, and $u_{1,t}$ and $u_{2,t}$ are error terms. For the sake of notational convenience and without loss of generality, we set $c^{(n,m)} = 0$ in equation (1.1) and use demeaned data in our analysis. This implies that we cannot discriminate between the standard formulation of the EH and the pure EH with a zero average term premium, but we focus on testing whether the term premium is constant over time.

The above formulation can be interpreted as a system where the forecasting

[8] A simple alternative would be to estimate the model without restrictions by least squares and to apply a Wald test. However, Bekaert and Hodrick (2001) provide simulation evidence that the Wald test has poor finite sample properties in presence on nonlinear restrictions relative to test statistics constrained under the null. Specifically, Bekaert and Hodrick (2001) show that the LM test has very satisfactory size properties and reasonable power. The DM test displays less satisfactory size and power properties than the LM test, whereas the Wald test shows the worst properties among these three test statistics.

equation (1.2) is used to generate the expected future short-term rate and the equation (1.3) determines the current long-term rate. Simultaneously, the system determines endogenously both sides of the EH statement given in equation (1.1), and allows joint estimation of the parameters. This improves efficiency by incorporating contemporaneous cross-correlation in the errors (Pagan, 1984; Mishkin, 1982).

The EH implies a set of nonlinear restrictions on the parameters of the above system. To define these restrictions, let us simplify the notation by translating the above p-order system into a first-order VAR companion form as

$$\begin{bmatrix} i_t^{(m)} \\ i_t^{(n)} \\ i_{t-1}^{(m)} \\ i_{t-1}^{(n)} \\ \vdots \\ i_{t-p+1}^{(m)} \\ i_{t-p+1}^{(n)} \end{bmatrix} = \begin{bmatrix} a_1 & b_1 & \cdots & a_{p-1} & b_{p-1} & a_p & b_p \\ c_1 & d_1 & \cdots & c_{p-1} & d_{p-1} & c_p & d_p \\ 1 & & & & & & \\ & 1 & & & & & \\ & & \ddots & & & & \\ & & & 1 & & & \\ & & & & 1 & & \end{bmatrix} \begin{bmatrix} i_{t-1}^{(m)} \\ i_{t-1}^{(n)} \\ i_{t-2}^{(m)} \\ i_{t-2}^{(n)} \\ \vdots \\ i_{t-p}^{(m)} \\ i_{t-p}^{(n)} \end{bmatrix} + \begin{bmatrix} u_{1,t} \\ u_{2,t} \\ \\ \\ \\ \\ \end{bmatrix} \tag{1.4}$$

where the blank elements are zeros. In compact form, this VAR can be expressed as

$$Y_t = \Gamma Y_{t-1} + \nu_t \tag{1.5}$$

where Y_t has $2p$ elements, Γ is a $2p$ square companion matrix, and v_t is the vector of innovations orthogonal to the information set available at time t, with zero mean and covariance matrix Σ_ν. Then, the EH subjects equation (1.5) to the following set of nonlinear cross-equation restrictions

$$e_2' = e_1' k^{-1} \left(I - \Gamma^m\right)^{-1} \left(I - \Gamma^n\right) \tag{1.6}$$

where $e_1 = (1, 0, \ldots, 0)'$ and $e_2 = (0, 1, 0, \ldots, 0)'$ are $2p$ dimensional indicator vectors.[9] Although equation (1.6) does not have a straightforward intuition, it gives a $2p$ dimensional vector of restrictions, nonlinear in the underlying parameters of Γ, such that the predictions of future short-term rates are consistent with the EH and the resulting constrained VAR collapses to equation (1.1). We can interpret these restrictions as a concise summary of the main implications stated by the theory. First, the constrained VAR defines the theoretical long-term rate we would observe in a world where expectations about future short-term rates are formed rationally. Second, under these restrictions the long-term rate contains all relevant information required by the market participants to predict future short-term rates. Put another way, the long-term rate provides optimal predictions of future short-term rates and deviations of the actual long-term rate from the theoretical long-term rate are unsystematic and unpredictable. Then, by rewriting the $2p$ dimensional vector of restrictions as

$$a(\theta) = e_2' - e_1' k^{-1} \left(I - \Gamma^m\right)^{-1} \left(I - \Gamma^n\right) \tag{1.7}$$

we can define the null hypothesis of rational expectations and constant term premium as

$$H_0 : a(\theta) = 0 \tag{1.8}$$

where θ is formed by collecting the relevant parameters of the companion matrix Γ.[10]

[9] Appendix A.1 provides further technical details on the restrictions implied by the EH in the VAR model.

[10] Specifically, the vector of parameters θ is defined as $\theta = (a_1, \cdots, a_p, b_1, \cdots, b_p, c_1, \cdots, c_p, d_1, \cdots, d_p)'$.

1.3.2 The VAR Tests

Bekaert and Hodrick (2001) propose a feasible method based on the GMM to estimate the VAR model under the hypothesis that the EH holds, defined by the nonlinear cross-equation restrictions on the parameters θ.[11]

Let $y_t \equiv [i_t^{(m)}, i_t^{(n)}]$ be the vector of data available at time t, u_t be the vector of orthogonal errors defined by the model, and x_{t-1} be the vector of instruments available at time $t-1$, formed by stacking lagged values of y_t (and possibly a constant term). Next, define the vector $z_t \equiv (y_t', x_{t-1}')'$, the vector-valued function of the data and the parameters $g(z_t, \theta) \equiv u_t \otimes x_{t-1}$, and the set of orthogonality conditions $E\left[g(z_t, \theta)\right] \equiv 0$. Using the corresponding sample moment conditions $g_T(\theta) \equiv T^{-1} \sum_{t=1}^{T} g(z_t, \theta)$ for a sample of size T, the parameters, θ, are estimated by minimizing the GMM criterion function

$$Q_T(\theta) \equiv g_T(\theta)' \Omega_T^{-1} g_T(\theta) \tag{1.9}$$

where Ω_T^{-1} is a positive semidefinite weighting matrix (Hansen, 1982).[12] To estimate the parameters, θ, subjected to the nonlinear restrictions defined by equation (1.6), we define the Lagrangian as

$$L(\theta, \gamma) = -\frac{1}{2} g_T(\theta)' \Omega_T^{-1} g_T(\theta) - a_T(\theta)' \gamma \tag{1.10}$$

where γ is a vector of Lagrange multipliers, and $a_T(\theta)$ is the sample counterpart of $a(\theta)$. While direct maximization of the Lagrangian is difficult as the constraints are

[11] Full maximum likelihood estimation of the restricted model requires restriction on the eigenvalues of the comapanion matrix Γ. Since the eigenvalues can be complex conjugates, direct estimation of the restricted VAR becomes quite complicated because the search must be conducted over potentially complex numbers (e.g. Bekaert and Hodrick, 2001; Melino, 2001).

[12] When Ω_T is chosen optimally, $\widehat{\theta}$ is asymptotically distributed as $\sqrt{T}(\widehat{\theta} - \theta_0) \rightarrow N(0, G_T' \Omega_T G_T)^{-1}$, where θ_0 denotes the true parameters, $\widehat{\theta}$ the parameter estimates, $G_T \equiv \nabla g_T\,(\theta)$ the gradient of the orthogonality conditions, and the symbol $\rightarrow$ denotes convergence in distribution.

nonlinear, Bekaert and Hodrick (2001) develop a recursive algorithm which extends the estimator proposed by Newey and McFadden (1994).[13]

If the restrictions have a significant impact on parameter estimation, then the value of the Lagrange multipliers is significantly different from zero and the null hypothesis that the EH holds is rejected. The hypothesis that the multipliers are jointly zero can be tested using the LM statistic

$$T\bar{\gamma}\left(A_T B_T^{-1} A_T'\right)\bar{\gamma} \longrightarrow \chi^2_{(2p)} \tag{1.11}$$

or the DM statistic

$$T g_T(\overline{\theta})' \Omega_T^{-1} g_T(\overline{\theta}) \longrightarrow \chi^2_{(2p)} \tag{1.12}$$

where $\overline{\theta}$ denotes the constrained estimates, and $2p$ is the number of restrictions implied by the EH.

1.3.3 Small Sample Properties

Tests of the EH null hypothesis have been known to suffer severely from problems related to finite sample bias estimation errors. In essence, the sampling distribution in finite sample may be significantly different from the asymptotic distribution (e.g. Bekaert, Hodrick and Marshall, 1997; Bekaert and Hodrick, 2001, Thornton, 2005, 2006). Thus, before estimating the unconstrained and constrained VARs, we follow Bekaert and Hodrick (2001) and use two different data generating processes (DGPs). Specifically, from the original data set, we simulate via bootstrap two bias-corrected data sets of 70,000 observations, with homoskedastic innovations and GARCH in-

[13]Notice that the GMM estimation is applied to the VAR defined in equations (1.2) and (1.3), whereas the companion VAR is exclusively used to simplify the derivation of the cross-equation restrictions. We refer to Appendix A.2 for further technical details on the GMM procedure.

novations, and use them throughout the econometric analysis. See Appendix A.3 for technical details on the procedure to account for small-sample bias in our analysis.

1.4 Empirical Results

In the empirical analysis, we obtain the unconstrained parameter estimate of θ, denoted $\widehat{\theta}$, by least squares and its constrained estimate $\overline{\theta}$ by the constrained GMM scheme for all possible pairwise combinations of short- and long-term rates such that $k = n/m$ is an integer. To take into account the day-of-the-week regularities in the short-term repo rates, documented in Table 1.1, we follow Longstaff (2000a) and set the VAR lag length to be $p = 5$.

Tables 1.3 and 1.4 report bias-corrected coefficients for the unconstrained VARs and the constrained VARs that satisfy the EH, respectively, when the DGP used to bias correct the parameters assumes homoskedastic innovations. Comparing the coefficients in Tables 1.3 and 1.4, we note that there are sharp differences in the constrained and unconstrained estimated dynamics. In particular, for each pairwise comparison, we find that the standard errors are quite large in the constrained VAR. Also, the absolute size of the constrained coefficients is much larger than the corresponding unconstrained ones, and, perhaps more importantly, the constrained coefficients measuring the response of the short-term rate to the long-term rate sometimes have a different sign from the corresponding unconstrained estimates. This is *prima facie* evidence that the EH restrictions may be inconsistent with the data, although this evidence does not constitute a formal statistical test.

For robustness, we also carry out estimation of the VAR-GARCH model, reported in Tables 1.5, 1.6 and 1.7. Table 1.5, Panel A reports the factor loadings, which are found to be statistically significant at standard significance levels, indicating the presence of GARCH effects. In Panel B, we also notice that the conditional variance turns out to be persistent for the overnight repo and moderately persistent for the spreads. Hence, departing from the assumption of homoskedasticity is likely to yield more accurate estimates of the VAR parameters and, consequently, more precise tests of the EH.

Tables 1.6 and 1.7 report bias-corrected coefficients for the unconstrained VARs and the constrained VARs that satisfy the EH, respectively, when the DGP used to bias correct the parameters assumes GARCH innovations. These results are quantitatively different from but qualitatively identical to the results for the VAR with homoskedastic innovations given in Tables 1.6-1.7. Specifically, the standard errors of parameters estimates in the constrained VAR are large, the absolute size of the constrained coefficients is larger than the corresponding unconstrained ones, and the constrained coefficients measuring the response of the short-term rate to the long-term rate have sometimes a different sign from the corresponding estimates in the unconstrained VAR.

The LM and DM tests results are presented in Table 1.8, where we report the p-values for the null hypothesis that the EH holds for all possible repo rates combinations of the integer $k = n/m$. The results in Table 1.8 indicate that the EH is rejected for each rate pair with p-values that are well below standard significance levels. Table 1.8 also reports the p-values from the J-test, which provides a specification test of the validity of the overidentifying moment conditions. The p-values are comfortably larger than conventional significance levels, validating the GMM estimation and, hence, the LM and DM tests.

1.5 Conclusions

The EH plays an important role in economics and finance and, not surprisingly, has been widely tested using a variety of tests and data. Much of the empirical literature has struggled to find evidence supporting the validity of the EH across a variety of data sets and countries, and employing increasingly sophisticated testing procedures. This chapter re-examines an important exception in this literature: the result that the EH appears to fit the behaviour of US repo rates at the shortest end of the term structure, measured at daily frequency from overnight to the 3-month maturity (Longstaff, 2000a). We extend this research by testing the restrictions implied by the EH on a VAR of the long- and short-term repo rate using the test proposed by Bekaert and Hodrick (2001). Our empirical investigation, in contrast to Longstaff (2000a), is not encouraging for the EH, which is statistically rejected across the term structure considered.

These findings differ from Longstaff (2000a), who does not reject the EH using conventional tests, because the VAR test is particularly powerful – and, hence, more likely to detect fine departures from the null hypothesis in finite sample – and because our sample is larger than Longstaff's (2000a). However, despite this statistical evidence, a legitimate and unanswered concern is whether the rejection of the EH may be due to small departures from the null hypothesis (or tiny data imperfections) which are not economically meaningful but appear statistically significant given the powerful test statistics and the very large sample size employed.[14]

[14]Leamer (1978, Chapter 4) points out that classical hypothesis testing will lead to rejection of any null hypothesis with a sufficiently large sample: 'Classical hypothesis testing at a fixed level

Moreover, the VAR tests are not designed to incorporate the fact that if one wanted to trade on departures from the EH – rather than assuming that the EH holds in a simple buy-and-hold allocation strategy – transactions costs create a wedge between returns from an active strategy exploiting departures from the EH and a simple buy-and-hold strategy. Finally, while the VAR tests rely on the ability of the VAR to capture the time-series properties of the term structure of repo rates, we are aware that the simple VAR tests, inspired by the literature on affine term structure models, is in fact unable to satisfactorily explain conditional means and volatility of interest rates. Hence, potential model misspecification and model uncertainty could play an important role in determining the rejection of the EH recorded in Table 1.8. In order to address these issues and to shed light on the economic significance of the statistical rejections of the EH recorded in this section, we proceed, in the next chapter, to an economic evaluation of the EH departures.

of significance increasingly distorts the interpretation of the data against a null hypothesis as the sample size grows. The significance level should consequently be a decreasing function of sample size' (p. 114).

Table 1.1

Descriptive Statistics for Daily Repo Rates

	Panel A: Percent Values							*Panel B: Percent Daily Changes*						
	i_t	$i_t^{(1w)}$	$i_t^{(2w)}$	$i_t^{(3w)}$	$i_t^{(1m)}$	$i_t^{(2m)}$	$i_t^{(3m)}$	Δi_t	$\Delta i_t^{(1w)}$	$\Delta i_t^{(2w)}$	$\Delta i_t^{(3w)}$	$\Delta i_t^{(1m)}$	$\Delta i_t^{(2m)}$	$\Delta i_t^{(3m)}$
$Mean$	3.9600	3.9492	3.9521	3.9544	3.9623	3.9752	3.9924	−0.0004	−0.0005	−0.0005	−0.0005	−0.0005	−0.0005	−0.0004
$Mean_{Mon}$	3.9718	3.9433	3.9420	3.9428	3.9483	3.9599	3.9764	−0.0360	−0.0106	−0.0057	−0.0051	−0.0016	−0.0004	0.0000
$Mean_{Tue}$	3.9728	3.9628	3.9657	3.9672	3.9757	3.9890	4.0051	−0.0040	−0.0064	−0.0037	−0.0029	−0.0036	−0.0035	−0.0030
$Mean_{Wed}$	3.9616	3.9496	3.9544	3.9571	3.9650	3.9784	3.9952	0.0036	−0.0022	−0.0049	−0.0039	−0.0037	−0.0022	0.0002
$Mean_{Thu}$	3.9683	3.9492	3.9526	3.9563	3.9642	3.9780	3.9969	−0.0330	−0.0028	−0.0006	−0.0010	0.0004	−0.0016	−0.0027
$Mean_{Fri}$	3.9260	3.9403	3.9445	3.9474	3.9565	3.9692	3.9866	0.0643	0.0191	0.0123	0.0103	0.0060	0.0055	0.0033
$Std\ Dev$	1.6998	1.6944	1.6973	1.6990	1.7003	1.7007	1.7013	0.1738	0.0648	0.0524	0.0517	0.0488	0.0507	0.0567
$Std\ Dev_{Mon}$	1.7039	1.7008	1.7017	1.7032	1.7023	1.7019	1.7010	0.1533	0.0672	0.0584	0.0621	0.0550	0.0517	0.0656
$Std\ Dev_{Tue}$	1.6951	1.6927	1.6959	1.6978	1.6999	1.6997	1.7009	0.1818	0.0640	0.0486	0.0498	0.0516	0.0540	0.0609
$Std\ Dev_{Wed}$	1.7115	1.7015	1.7054	1.7069	1.7072	1.7075	1.7074	0.2081	0.0621	0.0484	0.0448	0.0466	0.0484	0.0492
$Std\ Dev_{Thu}$	1.6975	1.6884	1.6927	1.6948	1.6978	1.6993	1.7009	0.1383	0.0447	0.0447	0.0463	0.0471	0.0518	0.0570
$Std\ Dev_{Fri}$	1.6953	1.6935	1.6956	1.6969	1.6991	1.6996	1.7006	0.1580	0.0779	0.0590	0.0535	0.0426	0.0470	0.0500
Min	0.8400	0.8900	0.8800	0.8700	0.8600	0.8300	0.8300	−1.5500	−0.8200	−0.8300	−0.8400	−0.8600	−0.8100	−0.8600
Max	6.7500	6.7000	6.5000	6.4900	6.4700	6.5000	6.5800	3.4000	1.1000	0.4100	0.6300	0.2900	0.3700	0.6200
ρ_1	0.9948	0.9993	0.9995	0.9995	0.9996	0.9996	0.9994	−0.3226	−0.0308	−0.1077	−0.1885	−0.1806	−0.2354	−0.2882
ρ_2	0.9929	0.9986	0.9991	0.9992	0.9993	0.9993	0.9992	−0.0921	−0.0150	0.0420	0.0399	0.0209	−0.0158	0.0467
ρ_3	0.9920	0.9979	0.9987	0.9989	0.9990	0.9991	0.9989	−0.0287	−0.0650	−0.0112	−0.0200	−0.0449	0.0123	−0.0345
ρ_4	0.9914	0.9973	0.9983	0.9986	0.9988	0.9989	0.9987	−0.0041	−0.1112	0.0101	0.0388	0.0491	0.0500	0.0494
ρ_5	0.9909	0.9969	0.9979	0.9983	0.9985	0.9986	0.9984	−0.0350	−0.0270	−0.0022	−0.0097	0.0276	−0.0225	0.0019
$\sigma(a)$	1.1640	1.1625	1.1654	1.1669	1.1681	1.1687	1.1691							
$Mean \times \sigma(a)$	4.6093	4.5909	4.6057	4.6142	4.6282	4.6460	4.6676							

The table summarizes the descriptive statistics for the daily repo rates (Panel A), and daily changes in repo rates (Panel B), from overnight to 3-month maturity. The data set consists of 3,625 daily observations of the indicated term government general collateral repo rates from May 21, 1991 to December 9, 2005, quoted on a 360-day basis and expressed in percentage points per annum. The period September 10, 2001 to September 30, 2001 is not included. The daily change in repo rate for the indicated weekday is measured from the indicated day to the next business day. ρ_i denotes the i-th order serial correlation coefficient. $\sigma(a) = \sqrt{Var[i_t(a)]}$ is the annualized volatility, where $i_t(a) = \sum_{k=0}^{a-1} i_{t-k}(d)$ is the sum of the daily returns, $a = 250$ is the average number of trading days, and $i_t(d) = \frac{i_t}{360 \times 100}$ is the daily return for a given raw repo rate i_t. All statistics are measured in percentage points per annum.

Table 1.2

Descriptive Statistics for Daily T-bill Rates

	Panel A: 1961-2005				*Panel B: 1991-2005*			
	$Tb_t^{(1m)}$	$Tb_t^{(3m)}$	$\Delta Tb_t^{(1m)}$	$\Delta Tb_t^{(3m)}$	$Tb_t^{(1m)}$	$Tb_t^{(3m)}$	$\Delta Tb_t^{(1m)}$	$\Delta Tb_t^{(3m)}$
$Mean$	5.5130	5.7597	0.0002	0.0001	3.6823	3.8358	−0.0005	−0.0006
$Mean_{Mon}$	5.5339	5.7754	0.0004	0.0018	3.7016	3.8508	−0.0034	−0.0057
$Mean_{Tue}$	5.5337	5.7798	−0.0044	−0.0103	3.7046	3.8584	−0.0057	−0.0102
$Mean_{Wed}$	5.5424	5.7864	−0.0176	−0.0079	3.6982	3.8483	−0.0111	−0.0048
$Mean_{Thu}$	5.5152	5.7694	0.0063	0.0012	3.6831	3.8405	0.0055	0.0029
$Mean_{Fri}$	5.4428	5.6900	0.0160	0.0158	3.6285	3.7851	0.0113	0.0137
$Std\ Dev$	2.7856	2.8567	0.1305	0.0933	1.5764	1.6112	0.0693	0.0416
$Std\ Dev_{Mon}$	2.8002	2.8709	0.1197	0.0805	1.5863	1.6217	0.0761	0.0466
$Std\ Dev_{Tue}$	2.7946	2.8591	0.1249	0.0725	1.5769	1.6200	0.0811	0.0386
$Std\ Dev_{Wed}$	2.7979	2.8678	0.1206	0.0818	1.5904	1.6219	0.0643	0.0386
$Std\ Dev_{Thu}$	2.7693	2.8501	0.1248	0.0953	1.5797	1.6124	0.0627	0.0431
$Std\ Dev_{Fri}$	2.7685	2.8382	0.1555	0.1238	1.5536	1.5855	0.0585	0.0369
Min	0.7360	0.7900	−1.8830	−1.3080	0.7360	0.7900	−1.1120	−0.8570
Max	17.926	17.682	2.0760	1.5090	6.4290	6.2970	0.9880	0.4490
ρ_1	0.9989	0.9995	0.0449	0.2000	0.9990	0.9997	−0.0239	0.0917
ρ_2	0.9977	0.9987	0.0344	0.0606	0.9981	0.9993	−0.0485	−0.0130
ρ_3	0.9964	0.9979	0.0187	0.0180	0.9973	0.9989	−0.0023	−0.0125
ρ_4	0.9951	0.9971	0.0270	0.0598	0.9965	0.9985	0.0026	0.0492
ρ_5	0.9938	0.9962	0.0718	0.0556	0.9956	0.9981	0.0305	0.0361
$\sigma(a)$	1.8488	1.9144			1.0682	1.0988		
$Mean \times \sigma(a)$	10.194	11.027			3.9310	4.2120		

The table summarizes the descriptive statistics for daily T-bill rates, Tb_t, and daily changes in T-bill rates, ΔTb_t, for the 1-month ($1m$) and 3-month ($3m$) maturity, respectively. The data are measured in percentage points per annum. Panel A reports the statistics for the period June 14, 1961 to December 30, 2005 and consists of 11110 daily observations. Panel B reports the statistics for the period May 21, 1991 to December 9, 2005 and consists of 3,568 daily observations. The daily change in the T-bill rate for the indicated maturity is measured from the indicated day to the next business day. ρ_i denotes the i-th order serial correlation coefficient. $\sigma(a) = \sqrt{Var[i_t(a)]}$ is the annualized volatility, where $i_t(a) = \sum_{k=0}^{a-1} i_{t-k}(d)$ is the sum of the daily returns, $a = 250$ is the average number of trading days, and $i_t(d) = \frac{i_t}{360\times100}$ is the daily return for a given raw repo rate i_t. All statistics are measured in percentage points per annum.

Table 1.3

Unconstrained VAR Dynamics with Homoskedastic Innovations

Panel A: overnight $i_t^{(m)}$ *vs. 1-week* $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.2662	0.7015	−0.0455	0.0347	−0.0146	0.0423	−0.0093	−0.0219	−0.0087	0.0548
	(0.0040)	(0.0086)	(0.0041)	(0.0115)	(0.0041)	(0.0115)	(0.0041)	(0.0115)	(0.0040)	(0.0087)
$i_t^{(n)}$	0.0462	0.9267	−0.0238	0.0305	−0.0034	−0.0438	−0.0128	−0.0427	0.0008	0.1219
	(0.0018)	(0.0040)	(0.0019)	(0.0053)	(0.0019)	(0.0053)	(0.0019)	(0.0053)	(0.0018)	(0.0040)

Panel B: overnight $i_t^{(m)}$ *vs. 2-week* $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.3258	0.4357	−0.0210	0.2842	−0.0094	−0.0262	−0.0048	−0.0774	−0.0223	0.1135
	(0.0039)	(0.0104)	(0.0041)	(0.0137)	(0.0041)	(0.0138)	(0.0041)	(0.0137)	(0.0039)	(0.0106)
$i_t^{(n)}$	0.0241	0.8714	−0.0194	0.1566	−0.0176	−0.0164	−0.0071	0.0026	0.0095	−0.0042
	(0.0014)	(0.0039)	(0.0015)	(0.0051)	(0.0015)	(0.0051)	(0.0015)	(0.0051)	(0.0014)	(0.0040)

Panel C: overnight $i_t^{(m)}$ *vs. 3-week* $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.3544	0.3889	0.0036	0.2738	0.0122	−0.0763	0.0060	−0.0514	−0.0023	0.0884
	(0.0039)	(0.0108)	(0.0041)	(0.0137)	(0.0041)	(0.0139)	(0.0041)	(0.0138)	(0.0038)	(0.0111)
$i_t^{(n)}$	0.0126	0.7984	−0.0162	0.1985	−0.0190	0.0083	−0.0096	0.0466	−0.0024	−0.0177
	(0.0014)	(0.0039)	(0.0015)	(0.0049)	(0.0015)	(0.0050)	(0.0015)	(0.0049)	(0.0014)	(0.0040)

Panel D: overnight $i_t^{(m)}$ *vs. 1-month* $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.4106	0.3061	0.0346	0.1944	0.0362	−0.1804	0.0221	0.1004	0.0123	0.0611
	(0.0038)	(0.0116)	(0.0041)	(0.0149)	(0.0041)	(0.0150)	(0.0041)	(0.0149)	(0.0038)	(0.0119)
$i_t^{(n)}$	0.0179	0.8146	−0.0164	0.1630	−0.0118	−0.0162	−0.0085	0.0848	−0.0038	−0.0240
	(0.0013)	(0.0038)	(0.0014)	(0.0049)	(0.0014)	(0.0049)	(0.0013)	(0.0049)	(0.0012)	(0.0039)

Panel E: overnight $i_t^{(m)}$ *vs. 2-month* $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.4539	0.2259	0.0589	0.1228	0.0596	−0.1476	0.0455	0.1073	0.0451	0.0262
	(0.0038)	(0.0116)	(0.0042)	(0.0143)	(0.0042)	(0.0144)	(0.0042)	(0.0144)	(0.0038)	(0.0117)
$i_t^{(n)}$	0.0293	0.7349	−0.0302	0.1841	−0.0095	0.0742	−0.0095	0.0686	−0.0124	−0.0301
	(0.0012)	(0.0038)	(0.0014)	(0.0047)	(0.0014)	(0.0047)	(0.0014)	(0.0047)	(0.0012)	(0.0039)

(continued)

Table 1.3 *(continued)*

Panel F: overnight $i_t^{(m)}$ vs. 3-month $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.4780 (0.0038)	0.2409 (0.0106)	0.0786 (0.0042)	0.0645 (0.0129)	0.0743 (0.0042)	−0.1137 (0.0131)	0.0592 (0.0042)	−0.0243 (0.0129)	0.0706 (0.0037)	0.0699 (0.0108)
$i_t^{(n)}$	0.0226 (0.0014)	0.6935 (0.0038)	−0.0184 (0.0015)	0.2493 (0.0046)	−0.0104 (0.0015)	0.0231 (0.0047)	−0.0068 (0.0015)	0.0756 (0.0046)	−0.0160 (0.0013)	−0.0131 (0.0039)

Panel G: 1-week $i_t^{(m)}$ vs. 2-week $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.6103 (0.0048)	0.3270 (0.0055)	−0.0389 (0.0056)	0.0793 (0.0067)	−0.1263 (0.0056)	0.1010 (0.0067)	−0.1026 (0.0056)	0.0647 (0.0067)	0.0132 (0.0047)	0.0706 (0.0059)
$i_t^{(n)}$	0.0377 (0.0043)	0.8525 (0.0049)	−0.0320 (0.0050)	0.1683 (0.0059)	−0.0916 (0.0050)	0.0476 (0.0060)	−0.0220 (0.0050)	0.0249 (0.0060)	0.0311 (0.0042)	−0.0171 (0.0052)

Panel H: 1-week $i_t^{(m)}$ vs. 3-week $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.7264 (0.0046)	0.1871 (0.0054)	−0.0187 (0.0056)	0.0822 (0.0066)	−0.1119 (0.0056)	0.1230 (0.0067)	−0.0749 (0.0056)	0.0284 (0.0066)	0.0732 (0.0045)	−0.0164 (0.0058)
$i_t^{(n)}$	0.0201 (0.0039)	0.7837 (0.0046)	−0.0437 (0.0048)	0.2176 (0.0056)	−0.0392 (0.0048)	0.0331 (0.0057)	−0.0629 (0.0048)	0.0876 (0.0056)	0.0167 (0.0038)	−0.0138 (0.0049)

Panel I: 1-month $i_t^{(m)}$ vs. 2-month $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.6411 (0.0054)	0.1791 (0.0052)	0.1533 (0.0061)	−0.0186 (0.0058)	−0.0528 (0.0062)	0.0345 (0.0058)	0.0745 (0.0061)	−0.0047 (0.0058)	−0.0159 (0.0053)	0.0090 (0.0054)
$i_t^{(n)}$	0.1690 (0.0057)	0.6119 (0.0054)	0.0047 (0.0064)	0.1651 (0.0060)	−0.1572 (0.0064)	0.1768 (0.0061)	−0.0625 (0.0064)	0.1200 (0.0061)	−0.0500 (0.0055)	0.0219 (0.0056)

Panel J: 1-month $i_t^{(m)}$ vs. 3-month $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.6952 (0.0047)	0.1253 (0.0041)	0.1603 (0.0055)	−0.0171 (0.0047)	−0.0267 (0.0056)	−0.0066 (0.0048)	0.0688 (0.0055)	0.0136 (0.0047)	−0.0051 (0.0045)	−0.0080 (0.0041)
$i_t^{(n)}$	0.1163 (0.0053)	0.6336 (0.0047)	0.0127 (0.0063)	0.2296 (0.0053)	−0.0956 (0.0064)	0.0671 (0.0054)	0.0210 (0.0063)	0.0745 (0.0053)	−0.0977 (0.0052)	0.0382 (0.0048)

The table presents the unconstrained VAR parameter estimates adjusted for small-sample bias. The data generating process (DGP) used for the bias-correction assumes homoskedastic innovations. $i_t^{(n)}$ is the n-period (long-term) rate and $i_t^{(m)}$ is the m-period (short-term) rate. Each panel reports different combinations of short-term and long-term repo rates such that $k = n/m$ is an integer. Standard errors are reported in parenthesis.

Table 1.4

Constrained VAR Dynamics with Homoskedastic Innovations

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
Panel A: overnight $i_t^{(m)}$ vs. 1-week $i_t^{(n)}$										
$i_t^{(m)}$	0.2286 (0.1751)	0.8225 (0.3164)	−0.5215 (0.1746)	0.6936 (0.4348)	−0.2054 (0.1817)	−0.1770 (0.4436)	−0.2616 (0.1711)	−0.0859 (0.4316)	0.1196 (0.1773)	0.3863 (0.3267)
$i_t^{(n)}$	−0.1347 (0.0894)	1.2090 (0.1305)	−0.0385 (0.0926)	0.0089 (0.1884)	−0.0320 (0.0987)	−0.0419 (0.1886)	−0.0141 (0.0959)	0.0079 (0.1871)	0.0084 (0.0842)	0.0272 (0.1456)
Panel B: overnight $i_t^{(m)}$ vs. 2-week $i_t^{(n)}$										
$i_t^{(m)}$	0.7535 (0.2055)	0.9498 (0.5970)	−0.5332 (0.2291)	2.1630 (0.8370)	−0.0431 (0.2345)	−2.7820 (0.8645)	0.2823 (0.2224)	−0.5580 (0.9049)	0.3245 (0.2017)	0.4431 (0.6167)
$i_t^{(n)}$	−0.0633 (0.0546)	1.2590 (0.1538)	−0.0745 (0.0576)	0.0563 (0.2237)	−0.0104 (0.0587)	−0.2213 (0.2347)	0.0280 (0.0587)	−0.0241 (0.2277)	0.0214 (0.0566)	0.0293 (0.1723)
Panel C: overnight $i_t^{(m)}$ vs. 3-week $i_t^{(n)}$										
$i_t^{(m)}$	0.5064 (0.2210)	2.1260 (0.6373)	0.4321 (0.2362)	−0.5844 (0.9172)	0.1015 (0.2336)	−0.7700 (0.8275)	0.2523 (0.2351)	−1.2270 (0.8132)	0.1721 (0.2133)	−0.0069 (0.6715)
$i_t^{(n)}$	−0.0898 (0.0543)	0.9670 (0.1397)	−0.0293 (0.0566)	0.0775 (0.1853)	−0.0160 (0.0592)	0.0661 (0.1886)	−0.0145 (0.0570)	0.0446 (0.1921)	−0.0062 (0.0497)	0.0002 (0.1507)
Panel D: overnight $i_t^{(m)}$ vs. 1-month $i_t^{(n)}$										
$i_t^{(m)}$	0.9455 (0.1866)	1.2760 (0.5981)	−0.3140 (0.2179)	1.1240 (0.7717)	0.5022 (0.2246)	−2.8020 (0.8165)	0.0288 (0.2080)	0.8825 (0.8894)	0.1478 (0.1735)	−0.7888 (0.6722)
$i_t^{(n)}$	−0.0707 (0.0444)	0.9640 (0.1205)	−0.0082 (0.0480)	0.0393 (0.1619)	−0.0211 (0.0476)	0.0867 (0.1657)	−0.0052 (0.0480)	−0.0062 (0.1691)	−0.0049 (0.0402)	0.0260 (0.1396)
Panel E: overnight $i_t^{(m)}$ vs. 2-month $i_t^{(n)}$										
$i_t^{(m)}$	0.7475 (0.1841)	0.9542 (0.5065)	0.1531 (0.2052)	0.3933 (0.5847)	0.4843 (0.1991)	−0.3941 (0.5977)	−0.1387 (0.2048)	−0.4929 (0.6083)	−0.0651 (0.1945)	−0.6399 (0.4921)
$i_t^{(n)}$	−0.0483 (0.0452)	0.9502 (0.1608)	−0.0123 (0.0468)	0.0237 (0.1838)	−0.0087 (0.0494)	0.0387 (0.1776)	0.0056 (0.0469)	0.0305 (0.1737)	0.0019 (0.0447)	0.0183 (0.1550)

(continued)

Table 1.4 *(continued)*

Panel F: overnight $i_t^{(m)}$ vs. 3-month $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.7324	0.7945	0.1271	0.2549	0.4525	0.2989	−0.1255	−0.9856	−0.0764	−0.4717
	(0.1788)	(0.5197)	(0.1994)	(0.5724)	(0.1871)	(0.6151)	(0.2020)	(0.5764)	(0.1812)	(0.4872)
$i_t^{(n)}$	−0.0310	0.9698	−0.0070	0.0133	−0.0050	0.0193	0.0037	0.0265	0.0014	0.0088
	(0.0512)	(0.1415)	(0.0538)	(0.1703)	(0.0558)	(0.1651)	(0.0544)	(0.1673)	(0.0489)	(0.1429)

Panel G: 1-week $i_t^{(m)}$ vs. 2-week $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.6383	0.7791	−0.7206	0.8372	0.9606	−1.0300	−1.2080	0.4947	0.6610	−0.4139
	(0.2936)	(0.3515)	(0.3955)	(0.4998)	(0.3996)	(0.5078)	(0.3849)	(0.4794)	(0.2983)	(0.4074)
$i_t^{(n)}$	−0.2584	1.4060	0.0007	0.1942	0.1562	−0.3139	−0.4770	0.1818	0.2953	−0.1849
	(0.1946)	(0.2209)	(0.2292)	(0.2792)	(0.2279)	(0.2946)	(0.2244)	(0.2747)	(0.1950)	(0.2586)

Panel H: 1-week $i_t^{(m)}$ vs. 3-week $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.6635	0.3950	−0.2646	0.3161	0.3033	−0.2055	−0.6291	0.5706	0.1752	−0.3269
	(0.2598)	(0.2815)	(0.3411)	(0.3520)	(0.3318)	(0.3629)	(0.3085)	(0.3373)	(0.2474)	(0.3081)
$i_t^{(n)}$	−0.0064	1.0090	−0.1196	0.1536	0.0726	−0.0245	−0.2311	0.2025	0.0699	−0.1272
	(0.2071)	(0.2002)	(0.2615)	(0.2228)	(0.2481)	(0.2341)	(0.2513)	(0.2137)	(0.1665)	(0.1855)

Panel I: 1-month $i_t^{(m)}$ vs. 2-month $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.7047	0.1758	0.1700	−0.0967	−0.5617	0.3737	0.4135	−0.1273	−0.2036	0.1514
	(0.2284)	(0.1910)	(0.2531)	(0.2134)	(0.2592)	(0.2085)	(0.2603)	(0.2124)	(0.2246)	(0.2008)
$i_t^{(n)}$	−0.1482	1.0880	0.0859	−0.0489	−0.2815	0.1872	0.2070	−0.0638	−0.1019	0.0758
	(0.2451)	(0.1911)	(0.2794)	(0.2258)	(0.2925)	(0.2083)	(0.2933)	(0.2144)	(0.2214)	(0.2041)

Panel J: 1-month $i_t^{(m)}$ vs. 3-month $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.6712	0.2068	0.2236	−0.1122	−0.4692	0.2558	0.3072	−0.0090	−0.0269	−0.0472
	(0.1905)	(0.1495)	(0.2260)	(0.1677)	(0.2270)	(0.1647)	(0.2292)	(0.1733)	(0.1762)	(0.1594)
$i_t^{(n)}$	−0.1101	1.0690	0.0747	−0.0375	−0.1565	0.0853	0.1024	−0.0030	−0.0090	−0.0157
	(0.2277)	(0.1742)	(0.2659)	(0.1925)	(0.2753)	(0.1930)	(0.2715)	(0.2000)	(0.2165)	(0.1809)

The table presents the constrained VAR parameter estimates adjusted for small-sample bias. The data generating process (DGP) used for the bias-correction assumes homoskedatsic innovations. $i_t^{(n)}$ is the n-period (long-term) rate and $i_t^{(m)}$ is the m-period (short-term) rate. Each panel reports different combinations of short-term and long-term repo rates such that $k = n/m$ is an integer. Standard errors are reported in parenthesis.

Table 1.5
Factor GARCH Model

Panel A: Factor Loadings							
Maturity	f_1	f_2	f_3	f_4	f_5	f_6	f_7
i_t	1	0	0	0	0	0	0
$S_t^{(1w)}$	−0.8436 (0.00693)	1	0	0	0	0	0.2447 (0.01446)
$S_t^{(2w)}$	−0.8933 (0.00638)	0	1	0	0	0	0.3058 (0.01298)
$S_t^{(3w)}$	−0.9116 (0.00616)	0	0	1	0	0	0.3786 (0.01231)
$S_t^{(1m)}$	−0.9463 (0.00615)	0	0	0	1	0	0.4794 (0.01172)
$S_t^{(2m)}$	−0.9559 (0.00635)	0	0	0	0	1	0.6792 (0.00924)
$S_t^{(3m)}$	−0.9653 (0.00718)	0	0	0	0	0	1

Panel B: GARCH (1,1)							
Parameter	h_{1t}	h_{2t}	h_{3t}	h_{4t}	h_{5t}	h_{6t}	h_{7t}
ω_j	0.0001 (0.00001)	0.0006 (0.00004)	0.0003 (0.00003)	0.0004 (0.00004)	0.0005 (0.00012)	0.0001 (0.00003)	0.0008 (0.00006)
β_j	0.9014 (0.00888)	0.1995 (0.03783)	0.5305 (0.04272)	0.2212 (0.05774)	0.1121 (0.0614)	0.4940 (0.06973)	0.1427 (0.04294)
α_j	0.0886 (0.00888)	0.0795 (0.00773)	0.0240 (0.00340)	0.0300 (0.00411)	0.0144 (0.00298)	0.0091 (0.00159)	0.0650 (0.00833)

The table reports the volatility parameters, estimated by quasi-maximum likelihood, for the Factor-GARCH model. i_t denotes overnight repo rate and $S_t^{(j)}$ the spread between the j-period repo rate $i_t^{(j)}$ and overnight repo rate; F is the factor loadings matrix governing the structure given in equation (16) with overnight repo rate and 3-month spread as the factors, denoted as f_1 and f_7 respectively. The idiosyncratic innovations are assumed to follow an Augmented GARCH (1,1) process $h_{jt} = \omega_j \sqrt{i_{t-1}} + \beta_j h_{kt-1} + \alpha_j v_{j,t-1}$ with $j \in \{1, 2, \ldots, 7\}$. Asymptotic standard errors are reported below the parameter estimates.

Table 1.6

Unconstrained VAR Dynamics with GARCH Innovations

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
Panel A: overnight $i_t^{(m)}$ *vs. 1-week* $i_t^{(n)}$										
$i_t^{(m)}$	0.2754 (0.0044)	0.6997 (0.0128)	−0.0456 (0.0045)	0.0491 (0.0168)	−0.0179 (0.0045)	0.0445 (0.0168)	−0.0285 (0.0045)	−0.0368 (0.0168)	−0.0182 (0.0044)	0.0763 (0.0119)
$i_t^{(n)}$	0.0301 (0.0015)	0.9498 (0.0044)	−0.0332 (0.0015)	0.0416 (0.0058)	−0.0110 (0.0016)	−0.0308 (0.0058)	−0.0206 (0.0016)	−0.0348 (0.0058)	−0.0108 (0.0015)	0.1192 (0.0041)
Panel B: overnight $i_t^{(m)}$ *vs. 2-week* $i_t^{(n)}$										
$i_t^{(m)}$	0.3455 (0.0041)	0.4422 (0.0144)	−0.0186 (0.0043)	0.2476 (0.0188)	−0.0036 (0.0043)	−0.0154 (0.0189)	−0.0182 (0.0043)	−0.0881 (0.0189)	−0.0252 (0.0041)	0.1302 (0.0139)
$i_t^{(n)}$	0.0247 (0.0012)	0.8858 (0.0041)	−0.0197 (0.0012)	0.1562 (0.0054)	−0.0196 (0.0012)	−0.0212 (0.0054)	−0.0139 (0.0012)	0.0081 (0.0054)	0.0031 (0.0012)	−0.0040 (0.0040)
Panel C: overnight $i_t^{(m)}$ *vs. 3-week* $i_t^{(n)}$										
$i_t^{(m)}$	0.3719 (0.0040)	0.3720 (0.0143)	0.0038 (0.0042)	0.2369 (0.0181)	0.0160 (0.0042)	−0.0575 (0.0182)	−0.0092 (0.0042)	−0.0512 (0.0181)	−0.0119 (0.0040)	0.1248 (0.0140)
$i_t^{(n)}$	0.0218 (0.0011)	0.8077 (0.0040)	−0.0140 (0.0012)	0.1948 (0.0051)	−0.0187 (0.0012)	0.0039 (0.0051)	−0.0143 (0.0012)	0.0504 (0.0051)	−0.0033 (0.0011)	−0.0290 (0.0039)
Panel D: overnight $i_t^{(m)}$ *vs. 1-month* $i_t^{(n)}$										
$i_t^{(m)}$	0.4114 (0.0039)	0.2689 (0.0142)	0.0329 (0.0042)	0.1777 (0.0182)	0.0383 (0.0042)	−0.1292 (0.0184)	0.0009 (0.0042)	0.0825 (0.0183)	0.0053 (0.0039)	0.1068 (0.0142)
$i_t^{(n)}$	0.0240 (0.0010)	0.8122 (0.0039)	−0.0103 (0.0011)	0.1622 (0.0049)	−0.0118 (0.0011)	−0.0133 (0.0050)	−0.0083 (0.0011)	0.0787 (0.0049)	−0.0041 (0.0011)	−0.0297 (0.0038)
Panel E: overnight $i_t^{(m)}$ *vs. 2-month* $i_t^{(n)}$										
$i_t^{(m)}$	0.4470 (0.0038)	0.1867 (0.0150)	0.0565 (0.0042)	0.1016 (0.0188)	0.0606 (0.0042)	−0.1241 (0.0190)	0.0228 (0.0042)	0.0967 (0.0188)	0.0344 (0.0039)	0.1136 (0.0149)
$i_t^{(n)}$	0.0336 (0.0010)	0.7628 (0.0038)	−0.0259 (0.0011)	0.1750 (0.0048)	−0.0097 (0.0011)	0.0664 (0.0048)	−0.0074 (0.0011)	0.0621 (0.0048)	−0.0126 (0.0010)	−0.0449 (0.0038)

(continued)

Table 1.6 *(continued)*

Panel F: overnight $i_t^{(m)}$ vs. 3-month $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.4692 (0.0038)	0.1991 (0.0136)	0.0739 (0.0042)	0.0364 (0.0167)	0.0744 (0.0042)	−0.0822 (0.0170)	0.0385 (0.0042)	−0.0270 (0.0168)	0.0624 (0.0038)	0.1512 (0.0137)
$i_t^{(n)}$	0.0268 (0.0011)	0.7143 (0.0038)	−0.0158 (0.0012)	0.2416 (0.0047)	−0.0133 (0.0012)	0.0174 (0.0047)	−0.0022 (0.0012)	0.0712 (0.0047)	−0.0184 (0.0011)	−0.0224 (0.0038)

Panel G: 1-week $i_t^{(m)}$ vs. 2-week $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.6555 (0.0054)	0.3337 (0.0061)	−0.0672 (0.0064)	0.0616 (0.0074)	−0.1200 (0.0064)	0.0931 (0.0074)	−0.1038 (0.0063)	0.0540 (0.0074)	0.0168 (0.0053)	0.0743 (0.0064)
$i_t^{(n)}$	0.0500 (0.0048)	0.8743 (0.0054)	−0.0608 (0.0057)	0.1650 (0.0066)	−0.1000 (0.0057)	0.0472 (0.0066)	−0.0226 (0.0057)	0.0309 (0.0066)	0.0272 (0.0048)	−0.0118 (0.0057)

Panel H: 1-week $i_t^{(m)}$ vs. 3-week $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.7871 (0.0048)	0.1764 (0.0055)	−0.0575 (0.0060)	0.0738 (0.0067)	−0.1066 (0.0060)	0.1144 (0.0068)	−0.0790 (0.0060)	0.0231 (0.0067)	0.0840 (0.0047)	−0.0175 (0.0058)
$i_t^{(n)}$	0.0672 (0.0042)	0.7766 (0.0048)	−0.0786 (0.0052)	0.2205 (0.0059)	−0.0517 (0.0052)	0.0352 (0.0059)	−0.0513 (0.0052)	0.0830 (0.0059)	0.0124 (0.0041)	−0.0142 (0.0051)

Panel I: 1-month $i_t^{(m)}$ vs. 2-month $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.6449 (0.0054)	0.2023 (0.0055)	0.1449 (0.0061)	−0.0272 (0.0062)	−0.0467 (0.0061)	0.0257 (0.0062)	0.0751 (0.0061)	−0.0102 (0.0062)	−0.0098 (0.0053)	0.0004 (0.0056)
$i_t^{(n)}$	0.1784 (0.0053)	0.6350 (0.0054)	−0.0067 (0.0059)	0.1523 (0.0060)	−0.1542 (0.0059)	0.1733 (0.0060)	−0.0631 (0.0059)	0.1177 (0.0060)	−0.0470 (0.0052)	0.0139 (0.0055)

Panel J: 1-month $i_t^{(m)}$ vs. 3-month $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.7057 (0.0046)	0.1360 (0.0043)	0.1508 (0.0055)	−0.0221 (0.0049)	−0.0258 (0.0055)	−0.0107 (0.0050)	0.0687 (0.0055)	0.0120 (0.0049)	−0.0046 (0.0045)	−0.0107 (0.0044)
$i_t^{(n)}$	0.1337 (0.0049)	0.6479 (0.0046)	−0.0023 (0.0059)	0.2213 (0.0053)	−0.1025 (0.0059)	0.0650 (0.0054)	0.0203 (0.0059)	0.0745 (0.0053)	−0.0923 (0.0048)	0.0338 (0.0047)

The table presents the unconstrained VAR parameter estimates adjusted for small-sample bias. The data generating process (DGP) used for the bias-correction assumes GARCH innovations. $i_t^{(n)}$ is the n-period (long-term) rate and $i_t^{(m)}$ is the m-period (short-term) rate. Each panel reports different combinations of short-term and long-term repo rates such that $k = n/m$ is an integer. Standard errors are reported in parenthesis.

Table 1.7

Constrained VAR Dynamics with GARCH Innovations

Panel A: overnight $i_t^{(m)}$ *vs. 1-week* $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.3251 (0.2008)	1.7610 (0.7283)	−0.4489 (0.2135)	1.2170 (0.9923)	−0.6072 (0.2034)	−0.4146 (0.9415)	0.4183 (0.1991)	−2.3480 (0.9386)	0.2903 (0.1912)	0.8060 (0.5594)
$i_t^{(n)}$	−0.0542 (0.0407)	1.1950 (0.1287)	−0.1091 (0.0421)	0.2705 (0.1788)	−0.0905 (0.0430)	−0.1230 (0.1801)	0.0480 (0.0425)	−0.2598 (0.1865)	0.0325 (0.0427)	0.0903 (0.1129)

Panel B: overnight $i_t^{(m)}$ *vs. 2-week* $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.3725 (0.2266)	4.2410 (1.1830)	−0.5280 (0.2552)	−0.1566 (1.4590)	−0.1095 (0.2489)	−3.3550 (1.4720)	0.9648 (0.2654)	−3.5840 (1.5570)	0.2606 (0.2092)	2.8940 (1.0480)
$i_t^{(n)}$	−0.0026 (0.0300)	0.8769 (0.1282)	−0.0323 (0.0340)	0.3286 (0.1724)	−0.0610 (0.0340)	0.0505 (0.1791)	−0.0046 (0.0323)	−0.1853 (0.1742)	0.0024 (0.0321)	0.0273 (0.1151)

Panel C: overnight $i_t^{(m)}$ *vs. 3-week* $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	1.0480 (0.2449)	1.5620 (1.1850)	−0.1600 (0.2659)	1.5160 (1.4880)	−0.7218 (0.2751)	−1.3830 (1.5010)	0.6094 (0.2651)	−2.6810 (1.5000)	0.5887 (0.2298)	0.6252 (1.0660)
$i_t^{(n)}$	−0.0036 (0.0280)	0.9236 (0.1101)	−0.0326 (0.0301)	0.1892 (0.1367)	−0.0622 (0.0307)	0.0377 (0.1476)	−0.0116 (0.0299)	−0.0574 (0.1424)	0.0080 (0.0298)	0.0085 (0.1042)

Panel D: overnight $i_t^{(m)}$ *vs. 1-month* $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.9087 (0.1952)	2.6450 (1.1030)	−0.1655 (0.2189)	0.3231 (1.3950)	−0.5085 (0.2154)	−1.2960 (1.4340)	0.5701 (0.2073)	−2.7540 (1.5210)	0.5065 (0.1961)	0.7747 (1.0260)
$i_t^{(n)}$	−0.0137 (0.0281)	0.8960 (0.1044)	−0.0093 (0.0329)	0.1264 (0.1266)	−0.0336 (0.0317)	0.0706 (0.1290)	−0.0198 (0.0325)	−0.0099 (0.1352)	−0.0029 (0.0300)	−0.0043 (0.1052)

Panel E: overnight $i_t^{(m)}$ *vs. 2-month* $i_t^{(n)}$

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
$i_t^{(m)}$	0.9254 (0.1606)	1.1460 (0.9667)	0.1275 (0.1857)	0.2946 (1.1410)	−0.4429 (0.1850)	0.9320 (1.1700)	0.3702 (0.1725)	−2.5520 (1.2540)	0.2059 (0.1631)	−0.0044 (0.8935)
$i_t^{(n)}$	−0.0339 (0.0300)	0.9697 (0.1347)	−0.0032 (0.0355)	0.0165 (0.1697)	−0.0012 (0.0365)	0.0229 (0.1665)	−0.0090 (0.0358)	0.0409 (0.1675)	−0.0033 (0.0308)	0.0001 (0.1296)

(continued)

Table 1.7 *(continued)*

	$i_{t-1}^{(m)}$	$i_{t-1}^{(n)}$	$i_{t-2}^{(m)}$	$i_{t-2}^{(n)}$	$i_{t-3}^{(m)}$	$i_{t-3}^{(n)}$	$i_{t-4}^{(m)}$	$i_{t-4}^{(n)}$	$i_{t-5}^{(m)}$	$i_{t-5}^{(n)}$
Panel F: overnight $i_t^{(m)}$ vs. 3-month $i_t^{(n)}$										
$i_t^{(m)}$	0.8504 (0.1403)	1.3090 (0.9041)	0.1023 (0.1688)	−0.2667 (1.1020)	−0.3512 (0.1670)	1.1180 (1.0980)	0.3024 (0.1566)	−2.2630 (1.1480)	0.1693 (0.1445)	0.0312 (0.7819)
$i_t^{(n)}$	−0.0186 (0.0307)	0.9923 (0.1233)	−0.0014 (0.0378)	0.0093 (0.1327)	−0.0007 (0.0371)	0.0076 (0.1366)	−0.0034 (0.0369)	0.0162 (0.1329)	−0.0012 (0.0319)	−0.0002 (0.1137)
Panel G: 1-week $i_t^{(m)}$ vs. 2-week $i_t^{(n)}$										
$i_t^{(m)}$	0.9011 (0.2267)	0.8882 (0.2760)	−0.1610 (0.2827)	−0.2609 (0.3567)	−0.3622 (0.2752)	0.4408 (0.3508)	−0.0055 (0.2945)	−0.4077 (0.3640)	0.5327 (0.2282)	−0.5652 (0.3030)
$i_t^{(n)}$	0.0226 (0.1540)	1.2010 (0.1775)	−0.1065 (0.1839)	−0.0195 (0.2253)	−0.2151 (0.1815)	0.2979 (0.2253)	−0.0620 (0.1785)	−0.1045 (0.2116)	0.2194 (0.1470)	−0.2328 (0.1995)
Panel H: 1-week $i_t^{(m)}$ vs. 3-week $i_t^{(n)}$										
$i_t^{(m)}$	1.4440 (0.2134)	−0.4138 (0.2671)	−0.8129 (0.3113)	0.7940 (0.3196)	−0.1750 (0.3124)	0.4348 (0.3207)	−0.4230 (0.3037)	0.0078 (0.3227)	0.5080 (0.2319)	−0.3659 (0.2912)
$i_t^{(n)}$	−0.2060 (0.1840)	1.2570 (0.1944)	−0.2410 (0.2428)	0.1966 (0.2277)	0.0251 (0.2354)	0.0161 (0.2305)	−0.0307 (0.2398)	−0.0559 (0.2304)	0.1335 (0.1784)	−0.0961 (0.1962)
Panel I: 1-month $i_t^{(m)}$ vs. 2-month $i_t^{(n)}$										
$i_t^{(m)}$	0.6018 (0.1682)	0.2211 (0.1672)	−0.0800 (0.1766)	0.1783 (0.1894)	−0.2035 (0.1694)	0.0155 (0.1823)	0.2014 (0.1722)	0.1004 (0.1891)	−0.0995 (0.1535)	0.0638 (0.1722)
$i_t^{(n)}$	−0.1992 (0.1764)	1.1110 (0.1637)	−0.0400 (0.1897)	0.0892 (0.1801)	−0.1018 (0.1851)	0.0077 (0.1704)	0.1007 (0.1915)	0.0502 (0.1840)	−0.0498 (0.1558)	0.0319 (0.1647)
Panel J: 1-month $i_t^{(m)}$ vs. 3-month $i_t^{(n)}$										
$i_t^{(m)}$	0.7011 (0.1649)	0.2237 (0.1458)	0.0444 (0.1933)	0.0875 (0.1680)	−0.2645 (0.1896)	0.0844 (0.1686)	0.1491 (0.1885)	−0.0607 (0.1718)	−0.0083 (0.1522)	0.0431 (0.1574)
$i_t^{(n)}$	−0.0997 (0.1701)	1.0750 (0.1377)	0.0148 (0.1994)	0.0292 (0.1690)	−0.0882 (0.2064)	0.0281 (0.1697)	0.0497 (0.2123)	−0.0202 (0.1659)	−0.0028 (0.1646)	0.0144 (0.1461)

The table presents the constrained VAR parameter estimates adjusted for small-sample bias. The data generating process (DGP) used for the bias-correction assumes GARCH innovations. $i_t^{(n)}$ is the n-period (long-term) rate and $i_t^{(m)}$ is the m-period (short-term) rate. Each panel reports different combinations of short-term and long-term repo rates such that $k = n/m$ is an integer. Standard errors are reported in parenthesis.

Table 1.8
Test Statistics

				Panel A: Bias-Correction with Homoskedastic Innovations						
$i^{(n)}/i^{(m)}$	$i_t^{(1w)}/i_t$	$i_t^{(2w)}/i_t$	$i_t^{(3w)}/i_t$	$i_t^{(1m)}/i_t$	$i_t^{(2m)}/i_t$	$i_t^{(3m)}/i_t$	$i_t^{(2w)}/i_t^{(1w)}$	$i_t^{(3w)}/i_t^{(1w)}$	$i_t^{(2m)}/i_t^{(1m)}$	$i_t^{(3m)}/i_t^{(1m)}$
LM	0.0001	0	0	0	0	0.0055	0	0.0001	0.0037	0.0023
DM	0	0	0	0	0	0	0	0	0	0
$J-Test$	0.34	0.40	0.18	0.18	0.59	0.81	0.16	0.16	0.48	0.78
				Panel B: Bias-Correction with GARCH Innovations						
$i^{(n)}/i^{(m)}$	$i_t^{(1w)}/i_t$	$i_t^{(2w)}/i_t$	$i_t^{(3w)}/i_t$	$i_t^{(1m)}/i_t$	$i_t^{(2m)}/i_t$	$i_t^{(3m)}/i_t$	$i_t^{(2w)}/i_t^{(1w)}$	$i_t^{(3w)}/i_t^{(1w)}$	$i_t^{(2m)}/i_t^{(1m)}$	$i_t^{(3m)}/i_t^{(1m)}$
LM	0	0	0	0	0	0	0	0.0001	0.0001	0.0005
DM	0	0	0	0	0	0	0	0	0	0
$J-Test$	0.63	0.37	0.38	0.30	0.91	0.71	0.16	0.58	0.43	0.96

The table reports the p-values for the Lagrange Multiplier (LM) and distance metric (DM) statistics under the null hypothesis that the EH is validated by the data for each pairwise combination of short-term and long-term repo rates such that $k = n/m$ is an integer. $i_t^{(n)}$ is the n-period (long-term) rate and $i_t^{(m)}$ is the m-period (short-term) rate. The p-values are calculated by bootstrap as described in the text. Panel A reports the results when the data generating process (DGP) used for bias-correction assumes homoskedastic innovations. Panel B reports the results when the DGP used for bias-correction assumes GARCH innovations. 0 denotes p-values below 10^{-5}. The $J-Test$ is the test for the overidentifying moment conditions in the GMM estimation, and figures reported are p-values.

2 An Economic Evaluation of the Expectation Hypothesis of the Term Structure of Very Short-Term Rates

2.1 Introduction

In this chapter we move beyond testing the validity of the EH from a purely statistical perspective and provide evidence on whether deviations from the EH are economically significant. Distinguishing between statistical analysis and economic evaluation is crucial for at least three reasons: in general statistical rejections of a hypothesis do not necessarily imply economic rejections (Leitch and Tanner, 1991); statistical VAR tests of the EH do not allow for transactions costs, which are critical for exploiting departures from the EH in real-world financial markets; and very powerful statistical tests may reject virtually any null hypothesis in large samples, without necessarily being informative about the size of departures from the hypothesis tested (Leamer, 1978). All these reasons suggest that an economic assessment of the deviations from the EH is desirable to complement the statistical tests.

In a mean-variance framework, we compare the performance of a dynamic portfolio strategy consistent with the EH to a dynamic portfolio strategy that exploits the departures from the EH. We use a utility-based performance criterion to compute the fee a risk-averse investor would be willing to pay to switch from the EH to a strategy that exploits departures from the EH to forecast interest rates. As an alternative economic measure, we also employ the risk-adjusted return of these two strategies. In short, we provide an economic test of the EH by evaluating the incremental profitability of an optimal (mean-variance efficient) strategy which relaxes the restrictions implied by the EH statement.

To anticipate our results, the results of our economic analysis lend support to the EH as we find no tangible economic gain to an investor who exploits departures from the EH relative to an investor who allocates capital simply on the basis of the predictions of the EH. Specifically, the evidence in this chapter shows that the economic value of departures from the EH is modest and generally smaller than the costs that an investor would incur if he were to trade to exploit the mispricing implied by EH violations. Hence, despite the statistical rejections of the EH recorded in the previous chapter, we conclude that the EH provides a fairly reasonable approximation to the repo rates term structure, consistent with Longstaff's interpretation of the functioning of the repo market.

The remainder of the chapter is as follows. In the next section we briefly review the framework for measuring the economic value of departures from the EH. Section 2.3 lays out the mean-variance setting and describes the performance measures used to assess the economic significance of EH violations. Section 2.4 reports the results on the validity of the EH using economic value measures. The conclusions are presented in Section 2.5.

2.2 Measuring the Economic Value of Deviations from the EH

We wish to measure whether departures from the EH provide information that is economically valuable, regardless of whether or not they are statistically significant on the basis of econometric tests. This section discusses the framework we use to evaluate the impact of allowing for deviations from the EH on the performance of dynamic allocation strategies in the repo market. We employ mean-variance analysis as a standard measure of portfolio performance assuming quadratic utility. Ultimately, we aim at measuring how much an investor is willing to pay for switching from a strategy that assumes that the EH holds ($\mathcal{EH}$ strategy) to a dynamic strategy which conditions on departures from the EH ($\mathcal{DEH}$ strategy). The $\mathcal{EH}$ strategy uses the outcome from the constrained VAR to determine the portfolio allocation, whereas the $\mathcal{DEH}$ strategy is based on the unconstrained VAR. The allocation strategy we consider is simple and intuitive. It consists of taking a position (either long or short) in a long-term repo, and then hedging it with an offsetting rolling position in a series of short-maturity repos. If the EH governs the relation between the long-term and short-term rates and an investor takes long positions in long-term repos and short rolling positions in short-term repos, then following this strategy over time allows the investor to earn the unconditional term premium, denoted as $c^{(n,m)}$ in equation (1.1). However, if one thinks of all repo rates in deviations from their unconditional mean (i.e. setting $c^{(n,m)} = 0$), as we do in our setting below, then this strategy should earn a return of zero before costs.

Regardless of the EH rejections recorded in Table 1.8, the tiny differences in unconditional means of repo rates at different maturities observed in Table 1.1 suggest the possibility that the economic value of trading on deviations from the EH in the repo market may not be as appealing as the statistical rejections from the VAR tests may imply. The investor using the constrained VAR is effectively using the simple strategy described above based upon the belief that there is no difference in the returns from investing in the longer repo rate and from investing in a series of shorter repo rates. However, if the investor does not believe in the EH and hence uses the unconstrained VAR, the resulting allocation strategy will be the outcome of the predictions of the model with respect to whether the longer-term rate is under- or over-valued relative to the series of shorter repo rates over the maturity of the longer rate. This may be seen as the implementation of the popular carry trade strategy that attempts to exploit mispricing along the term structure of interest rates. In other words, using the unconstrained VAR is tantamount to exploiting the deviations from the EH which we have recorded in the earlier statistical analysis. If the unconstrained VAR model gives predictions of short-term repo rates consistent with the EH, the results from the $\mathcal{EH}$ strategy should be equal to the results from the $\mathcal{DEH}$ strategy.[15] From this setting we can calculate directly a variety of common performance measures, in the form of performance fees $\mathcal{F}$ (Fleming, Kirby and Ostdiek, 2001) and risk-adjusted abnormal returns $\mathcal{M}$ (Modigliani and Modigliani, 1997).

We realise that a portfolio consisting only of repo rates is unlikely to be a realistic portfolio managed by a US investor. The repurchase agreements involving US Treasury securities are mainly used by banks in order to manage the quantity of reserves on a short-term basis and, hence, play an important role in the Federal Reserve's implementation of monetary policy. Moreover, the repo market plays a fundamental role in dealers' hedging activities and repos are used by investment managers who hedge the interest rate risk related to the activity of short-selling Treasury securities. Our main objective is not to design a realistic (executable) asset

[15]Nevertheless, when incorporating transactions costs, this equality will not hold exactly, and therefore incorporating transactions costs is a further relevant issue in the construction of a measure of economic value.

allocation strategy, but rather to measure the economic significance of deviations from the EH. Our measures of economic value complement the LM and DM tests for statistical significance of the EH by showing whether the constraints imposed on the VAR by the EH have economic value. On the one hand, departures from the EH may be statistically insignificant, and yet provide considerable value to an investor. On the other hand, the departures might be statistically significant, but be of little or no economic value to a repo market investor.[16] This economic evaluation is easier to carry out and assess by focusing exclusively on a VAR where the only assets being modeled are repo rates at various maturities, because the only source of risk in the resulting repo portfolio is interest rate risk.

2.3 The EH in a Dynamic Mean-Variance Framework

In mean-variance analysis, the maximum expected return strategy leads to a portfolio allocation on the efficient frontier. Specifically, consider the trading strategy of an investor who has a k-period horizon and constructs a daily dynamically rebalanced portfolio that maximises the conditional expected return subject to achieving a target conditional volatility. Computing the time-varying weights of this portfolio requires predictions of the k-period ahead forecast of the conditional mean and the conditional variance-covariance matrix.

Let r_{t+k} denote the $N \times 1$ vector of risky asset returns; $\mu_{t+k|t} = E_t\left[r_{t+k}\right]$ is the conditional expectation of r_{t+k}; and $\Sigma_{t+k|t} = E_t[(r_{t+k} - \mu_{t+k|t})(r_{t+k} - \mu_{t+k|t})']$ is the conditional variance-covariance matrix of r_{t+k}.[17] At each period t, the investor

[16]See Leitch and Tanner (1991) for an early treatment of the relationship between statistical significance and economic value.

[17]We use the subscript $t+k$ to indicate an investment horizon of k periods ahead, where $k = n/m$ is an integer which depends on the long- and short-term interest rates.

solves the following problem:

$$\max_{w_t}\{\mu_{p,t+k} = w_t'\mu_{t+k|t} + (1 - w_t'\iota)\, r_f\} \tag{2.1}$$
$$\text{s.t.}\ \left(\sigma_p^*\right)^2 = w_t'\Sigma_{t+k|t}w_t$$

where w_t is the $N \times 1$ vector of portfolio weights on the risky assets, $\mu_{p,t+k}$ is the conditional expected return of the portfolio, σ_p^* is the target conditional volatility of the portfolio returns, and r_f is the return on the riskless asset.[18] The solution to this optimization problem delivers the following risky asset weights:

$$w_t = \frac{\sigma_p^*}{\sqrt{C_t}}\Sigma_{t+k|t}^{-1}(\mu_{t+k|t} - \iota r_f) \tag{2.2}$$

where $C_t = (\mu_{t+k|t} - \iota r_f)'\Sigma_{t+k|t}^{-1}(\mu_{t+k|t} - \iota r_f)$. The weight on the riskless asset is $1 - w_t'\iota$.

By design, in this setting the optimal weights will vary across models only to the extent that predictions of the conditional moments will vary, which is precisely what the empirical models provide. In our setting, we carry out the economic value analysis comparing the outcome from the $\mathcal{DEH}$ strategy – a strategy that exploits deviations from the EH – to the $\mathcal{EH}$ strategy which assumes that the EH holds. We compute the calculations for both cases with homoskedastic and GARCH innovations in the bias-correction DGPs. In short, our objective is to determine whether there is economic value in using the unconstrained VAR which relaxes the constraints imposed by the EH.

[18]For simplicity, we drop the subscript t from the riskless return r_f.

2.3.1 Quadratic Utility

We rank the performance of the competing repo rate models using the West, Edison, and Cho (1993) methodology, which is based on mean-variance analysis with quadratic utility. The investor's realised utility in period $t+k$ can be written as:

$$U\left(W_{t+k}\right)=W_{t+k}-\frac{\lambda}{2}W_{t+k}^{2}=W_{t}R_{p,t+k}-\frac{\lambda W_{t}^{2}}{2}R_{p,t+k}^{2} \tag{2.3}$$

where W_{t+k} is the investor's wealth at $t+k$, λ determines his risk preference, and

$$R_{p,t+k}=1+r_{p,t+k}=1+\left(1-w_{t}^{\prime}1\right)r_{f}+w_{t}^{\prime}r_{t+k} \tag{2.4}$$

is the period $t+k$ gross return on his portfolio.

We quantify the economic value of deviations from the EH by setting the investor's degree of relative risk aversion (RRA), $\delta_t = \lambda W_t/\left(1-\lambda W_t\right)$, equal to a constant value δ. In this case, West, Edison, and Cho (1993) demonstrate that one can use the average realised utility, $\overline{U}\left(\cdot\right)$, to consistently estimate the expected utility generated by a given level of initial wealth. Specifically, the average utility for an investor with initial wealth W_0 is equal to:

$$\overline{U}\left(\cdot\right)=W_{0}\sum_{t=0}^{T-1}\left\{R_{p,t+k}-\frac{\delta}{2\left(1+\delta\right)}R_{p,t+k}^{2}\right\}. \tag{2.5}$$

We standardise the investor problem by assuming he allocates \$1 in every time period. Average utility depends on taste for risk. In the absence of restrictions on δ, quadratic utility exhibits increasing degree of RRA. This is counterintuitive since, for instance, an investor with increasing RRA becomes more averse to a percentage loss in wealth when his wealth increases. As in West, Edison and Cho

(1993) and Fleming, Kirby and Ostdiek (2001), fixing the degree of RRA, δ, implies that expected utility is linearly homogeneous in wealth: double wealth and expected utility doubles. Furthermore, by fixing δ rather than λ, we are implicitly interpreting quadratic utility as an approximation to a non-quadratic utility function, with the approximating choice of λ dependent on wealth. The estimate of expected quadratic utility given in Equation (2.5) is used to implement the Fleming, Kirby and Ostdiek (2001) framework for assessing the economic value of the $\mathcal{DEH}$ and $\mathcal{EH}$ strategies.[19]

2.3.2 Performance Measures

At any point in time, one set of estimates of the conditional moments is better than a second set if investment decisions based on the first set lead to higher average realised utility, $\overline{U}$. Alternatively, a better model requires less wealth to yield a given level of $\overline{U}$ than the alternative model. Following Fleming, Kirby, and Ostdiek (2001) we measure the economic value of the interest rate strategies by equating the average utilities for selected pairs of portfolios. Suppose, for example, that holding a portfolio constructed using the optimal weights based on the $\mathcal{EH}$ strategy yields the same average utility as holding the portfolio implied by the $\mathcal{DEH}$ strategy. The latter portfolio is subject to daily management expenses $\mathcal{F}$, expressed as a fraction of wealth invested in the portfolio. Since the investor would be indifferent between

[19] A critical aspect of mean-variance analysis is that it applies exactly only when the return distribution is normal *or* the utility function is quadratic. Hence, the use of quadratic utility is not necessary to justify mean-variance optimization. For instance, one could instead consider using utility functions belonging to the constant relative risk aversion (CRRA) class, such as power or log utility. However, quadratic utility is an attractive assumption because it provides a high degree of analytical tractability. Quadratic utility may also be viewed as a second order Taylor series approximation to expected utility. In an investigation of the empirical robustness of the quadratic approximation, Hlawitschka (1994) finds that a two-moment Taylor series expansion "may provide an excellent approximation" (p. 713) to expected utility and concludes that the ranking of common stock portfolios based on two-moment Taylor series is "almost exactly the same" (p. 714) as the ranking based on a wide range of utility functions.

these two strategies, we interpret $\mathcal{F}$ as the maximum performance fee the investor would be willing to pay to switch from the $\mathcal{EH}$ to the $\mathcal{DEH}$ strategy. In general, this utility-based criterion measures how much an investor with a mean-variance utility function is willing to pay for conditioning on the deviations from the EH, as modeled in the unconstrained VAR model.[20]

The performance fee depends on the investor's degree of risk aversion and is a measure of the economic significance of violations of the EH. To estimate the fee, we find the value of $\mathcal{F}$ that satisfies

$$\sum_{t=0}^{T-1} \left\{ \left(R_{p,t+k}^{\mathcal{DEH}} - \mathcal{F}\right) - \frac{\delta}{2(1+\delta)} \left(R_{p,t+k}^{\mathcal{DEH}} - \mathcal{F}\right)^2 \right\} = \sum_{t=0}^{T-1} \left\{ R_{p,t+k}^{\mathcal{EH}} - \frac{\delta}{2(1+\delta)} \left(R_{p,t+k}^{\mathcal{EH}}\right)^2 \right\} \tag{2.6}$$

where $R_{p,t+k}^{\mathcal{DEH}}$ denotes the gross portfolio return constructed using the predictions from the unconstrained VAR model, and $R_{p,t+k}^{\mathcal{EH}}$ is the gross portfolio return implied by the constrained VAR model. In the absence of transactions costs, under the EH $\mathcal{F} = 0$, while if the EH is violated $\mathcal{F} > 0$. However, when allowing for transactions costs, it is also possible that $\mathcal{F} < 0$ if the positive gain from trading on the information provided by the EH violation is lower than the loss incurred by the more costly dynamic rebalancing of the $\mathcal{DEH}$ strategy.

We also consider the Modigliani and Modigliani (1997) measure $\mathcal{M}$, which defines the abnormal return that the $\mathcal{DEH}$ strategy would have earned over the $\mathcal{EH}$ strategy if it had the same risk as the $\mathcal{EH}$ strategy

$$\mathcal{M} = \sigma[\varkappa^{\mathcal{EH}}](\mathcal{SR}^{\mathcal{DEH}} - \mathcal{SR}^{\mathcal{EH}}) \tag{2.7}$$

where $\mathcal{SR} = E[\varkappa]/\sigma[\varkappa]$ is the Sharpe Ratio, and $E[\varkappa]$ and $\sigma[\varkappa]$ are the expected

[20]For studies following this approach see also Fleming, Kirby and Ostdiek (2003), Marquering and Verbeek (2004) and Han (2006).

value and standard deviations of the excess return, $\varkappa$, of a selected strategy. The $\mathcal{DEH}$ strategy is leveraged downwards or upwards, so that it has the same volatility as the $\mathcal{EH}$ strategy. Therefore, the risk-adjusted abnormal return, $\mathcal{M}$, measures the outperformance of the $\mathcal{DEH}$ strategy with respect to the $\mathcal{EH}$ strategy while matching the same level of risk.[21]

2.3.3 Dynamic Strategies, Transaction Costs and Short Selling

Consider a US investor who allocates his wealth between a long–term n-period discount bond and a sequence of k short-term m-period discount bonds. The long-term bond price is known with certainty and implies a riskless return, whereas the rolling combination of short-term bonds generates a risky return, since $k-1$ future short-term bond prices are not known. Hence, on the basis of riskless return, r_f, and the forecasts of the conditional moments of risky return, $r_{t+k|t}$, the investor will define his portfolio optimization problem at time t.

We consider two alternative trading strategies. The $\mathcal{EH}$ strategy assumes that EH holds exactly, and hence the investor takes a position using forecasts based on the constrained VAR. In this case, the investor effectively trades assuming that equation (1.1) holds and, in the absence of transactions costs, he is indifferent between investing in the long rate or a series of short rates. However, if transactions costs are positive and equal for short- and long-rates, the investor will prefer investing in the long rate as this minimises costs. The $\mathcal{DEH}$ strategy uses the forecasts based on the unconstrained VAR. Specifically, each strategy comprises two steps at

[21] We also compute a measure that allows for downside risk. However, since the results are qualitatively identical to the performance fees and risk-adjusted abnormal returns, we do not report them here to conserve space.

time t. First, the investor uses the selected VAR model to generate the conditional moments, $\mu_{t+k|t}$ and $\Sigma_{t+k|t}$. Second, conditional on the predictions of this model and given the riskless return r_f, he dynamically rebalances his portfolio by computing optimal weights. He repeats this process every day until the end of the sample period.[22]

This setup determines whether using one particular conditional specification affects the performance of a short-horizon allocation strategy in an economically meaningful way. The predictions are all in-sample predictions, since our focus is not to provide forecasting models of the repo term structure but to evaluate the measured departures from the EH as determined by the unconstrained VAR model.

With daily rebalancing, transaction costs play an important role in evaluating the relative performance of different strategies. In particular, we assume that transaction costs at time t equal a fixed proportion τ of the value traded in long-term and short-term repos (Marquering and Verbeek, 2004; Han, 2006). We also assume that the costs are the same for trading short and long rates. This is consistent with the fact that the bid-ask spread is fairly constant across maturities in the repo market, in the order of 2 to 5 *bps*. We report results both with and without transactions costs, and also study the impact of short selling constraints. In the case of limited short selling we constrain the portfolio weights to be bounded between -1 and 2 (assuming that the investor can borrow no more than 100% of his wealth), while in the case of no short selling, the portfolio weights are constrained between 0 and 1.

[22] Since we consider a single risky return, $\Sigma_{t+k|t}$ simply reduces to a variance term. Notice that parameter estimates are based on the full sample information.

2.4 The Economic Value of EH Departures

Given the parameter estimates reported in Tables 1.3-1.4 and 1.6-1.7, we assume that a US investor dynamically updates his portfolio weights daily after reestimating the VAR model with the latest available data. The key question is whether the dynamic strategy that allows for departures from the EH generates economic gains relative to a benchmark dynamic strategy that assumes that the EH holds. We assess the economic value of conditioning on departures from the EH by analyzing the performance of the dynamically rebalanced portfolio constructed using pairwise combinations of repo rates.[23]

We compute the performance fee $\mathcal{F}$ and the risk-adjusted abnormal return $\mathcal{M}$ for (i) two target annualised portfolio volatilities, $\sigma_p^* = \{1\%, 2\%\}$, which are in a range that includes the observed annualised standard deviation of the data reported in Table 1.1; (ii) a degree of relative risk aversion $\delta = 5$;[24] (iii) for each pair of repo maturities where the long maturity is an exact multiple of the short maturity; (iv) two different DGPs for the parameter estimates, with homoskedastic and heteroskedastic innovations. Furthermore, we also exploit the impact of transaction costs and short selling by considering four different scenarios. In *case* 1 transaction costs are ignored and the weights are unrestricted; *in case* 2 the weights are unrestricted but we introduce transaction costs with $\tau = 4$ *bps*, a realistic cost on the basis of the observed bid-ask spread in the repo market; in *case* 3 we also add a limited short selling constraint by restricting the weights to be between -1 and 2;

[23] For weekends and holidays we consider the rate on the previous business day for which a rate was reported.

[24] We investigated different values of δ in the range between 2 and 10 but found no qualitative difference in our results.

and finally in *case* 4 we do not allow short selling so that the weights are between 0 and 1. The performance measures, $\mathcal{F}$ and $\mathcal{M}$, are reported in annualised basis points.[25]

2.4.1 Performance Measures

Table 2.1 presents the in-sample performance fees $\mathcal{F}$ and the risk-adjusted abnormal returns $\mathcal{M}$ for the $\mathcal{DEH}$ strategy against the $\mathcal{EH}$ strategy when the bootstrap experiment for bias correction assumes homoskedastic innovations. Panel A reports the results for a target volatility $\sigma_p^* = 1\%$, and Panel B for $\sigma_p^* = 2\%$.

The results in Table 2.1 suggest that the performance fees for switching from a model that assumes the EH holds to a model that exploits departures from the EH is generally fairly modest when we do not consider transaction costs and the portfolio weights are unrestricted (*case* 1). For example, if we set the target volatility at $\sigma_p^* = 1\%$, the annual performance fee a risk-averse investor would be willing to pay to switch from the $\mathcal{EH}$ strategy to the $\mathcal{DEH}$ strategy is at most 1.34 *bps*. If we calibrate the target volatility to be $\sigma_p^* = 2\%$, the largest annual performance fee reaches 2.70 *bps* and occurs when the overnight repo rate is the short-term rate and the 1-week repo rate is the long-term rate.

However, when we introduce transaction costs (*case* 2), the performance fees $\mathcal{F}$ become even smaller and are slightly negative at the shorter end of the maturity spectrum. For instance, given $\sigma_p^* = 1\%$ and the overnight repo rate versus the 3-week repo rate, the $\mathcal{DEH}$ strategy has a negative annual performance fee of about 3

[25] We experimented with slightly different values of transactions costs in the range between 2 and 5 *bps*, and found qualitatively similar results. Note that the transactions costs are virtual identical across maturities in the repo market, possibly only slightly smaller on one-day repos by some 0.5 *bps*.

bps. This suggests that the higher transactions costs incurred in the $\mathcal{DEH}$ strategy outweigh the benefit of conditioning on EH violations, with the performance fee generally decreasing in $k = m/n$ due to the larger number of trades needed in the rolling strategy. In other words, the EH violations are not economically significant after costs are taken into account.

When we move at the longer spectrum of the maturity and consider 1-month versus 3-month repo rates for $\sigma_p^* = 1\%$, we notice a performance fee of 0.49 *bps*. Interestingly, when we combine transaction costs and limited short-selling (*case* 3), the performance measures remain virtually the same as in *case 2*, suggesting that the weights are in the range from -1 and 2. In the fourth scenario, we consider dynamic strategies without short selling and with transaction costs (*case* 4). In this case the fees decrease moderately in absolute values confirming that the short selling constraints are now binding on the profitability of the strategies but their impact is modest. The risk-adjusted abnormal returns $\mathcal{M}$, are of very similar magnitude as (in some columns identical to) the performance fees $\mathcal{F}$, leading therefore to the same conclusions.

For robustness purposes, Table 2.2 reports the same performance criteria, $\mathcal{F}$ and $\mathcal{M}$, when we assume GARCH innovations for the bias correction procedure. The results are qualitatively identical to the case of the VAR with homoskedastic errors discussed in Table 2.1, providing evidence that EH violations are economically unimportant. However, quantitatively the results in Table 2.2 provide evidence of even smaller gains from the $\mathcal{DEH}$ strategy, with the performance fee $\mathcal{F}$ never reaching 2 *bps*.

2.5 Conclusions

This chapter re-examines an important exception in this literature: the result that the EH appears to fit the behaviour of US repo rates at the shortest end of the term structure, measured at daily frequency from overnight to the 3-month maturity (Longstaff, 2000a). In the first chapter we showed how Longstaff's results are overturned when using a longer sample period and more powerful statistical tests.

We innovate in this chapter by moving beyond statistical tests and providing complementary evidence on the validity of the EH using some economic value calculations. We assess the economic value of exploiting departures from the EH – i.e. using empirical models which condition on information contained in EH deviations – relative to the economic value of using a model that assumes the EH holds. The empirical results indicate that the economic value of departures from the EH is modest and generally smaller than the costs that an investor would incur to exploit the mispricing implied by EH violations. These findings are consistent with the thrust of Longstaff's (2000a) original conclusion.

The results from economic value calculations are in contrast with the results from VAR tests reported earlier. This difference confirms that statistical rejections of a hypothesis do not always imply economic rejections and raises doubts about the ability of the simple linear VAR framework to capture the relationship between repo rates at different maturities. Activities in the repo market at maturities of days or weeks are largely driven by liquidity considerations and by the attempts of banks to manage the quantity of reserves and to hedge interest rate risk on a short-term basis, rather than to speculate in search of excess returns. Hence, it seems unlikely that investors would be actively exploiting EH departures on a very short-term basis. Our main conclusion is that, even though the EH may be rejected statistically, it still provides a very reasonable approximation to the term structure of repo rates and constitutes a useful theory for practitioners in the repo market.

Table 2.1
Economic Value Results with Homoskedastic Innovations

Panel A: $\sigma_p^* = 1\%$

$i_t^{(m)} - i_t^{(n)}$	*Case* 1 $\mathcal{F}$	*Case* 1 $\mathcal{M}$	*Case* 2 $\mathcal{F}$	*Case* 2 $\mathcal{M}$	*Case* 3 $\mathcal{F}$	*Case* 3 $\mathcal{M}$	*Case* 4 $\mathcal{F}$	*Case* 4 $\mathcal{M}$
$i_t - i_t^{(1w)}$	1.34	1.34	−1.01	−0.95	−1.01	−0.95	−0.57	−0.52
$i_t - i_t^{(2w)}$	0.47	0.47	−2.62	−2.50	−2.62	−2.50	−1.41	−1.28
$i_t - i_t^{(3w)}$	0.20	0.20	−3.33	−3.15	−3.33	−3.15	−1.77	−1.56
$i_t - i_t^{(1m)}$	0.44	0.44	−4.70	−4.46	−4.70	−4.46	−2.79	−2.45
$i_t - i_t^{(2m)}$	0.92	0.92	−7.19	−7.23	−7.19	−7.23	−4.63	−4.11
$i_t - i_t^{(3m)}$	1.51	1.51	−12.29	−12.34	−12.29	−12.34	−6.19	−6.40
$i_t^{(1w)} - i_t^{(2w)}$	0.34	0.34	0.31	0.31	0.31	0.31	0.05	0.05
$i_t^{(1w)} - i_t^{(3w)}$	0.49	0.49	0.47	0.47	0.47	0.47	0.11	0.11
$i_t^{(1m)} - i_t^{(2m)}$	0.40	0.40	0.39	0.39	0.39	0.39	0.20	0.21
$i_t^{(1m)} - i_t^{(3m)}$	0.60	0.60	0.49	0.50	0.49	0.50	0.32	0.32

Panel B: $\sigma_p^* = 2\%$

$i_t^{(m)} - i_t^{(n)}$	*Case* 1 $\mathcal{F}$	*Case* 1 $\mathcal{M}$	*Case* 2 $\mathcal{F}$	*Case* 2 $\mathcal{M}$	*Case* 3 $\mathcal{F}$	*Case* 3 $\mathcal{M}$	*Case* 4 $\mathcal{F}$	*Case* 4 $\mathcal{M}$
$i_t - i_t^{(1w)}$	2.70	2.67	−1.11	−1.17	−1.11	−1.17	−0.54	−0.41
$i_t - i_t^{(2w)}$	0.95	0.94	−3.32	−3.04	−3.32	−3.04	−2.82	−2.39
$i_t - i_t^{(3w)}$	0.39	0.39	−7.73	−7.35	−7.73	−7.35	−3.55	−3.80
$i_t - i_t^{(1m)}$	0.88	0.88	−9.52	−9.00	−9.52	−9.00	−4.58	−4.39
$i_t - i_t^{(2m)}$	1.83	1.83	−17.29	−17.32	−17.29	−17.32	−8.27	−9.33
$i_t - i_t^{(3m)}$	3.02	3.02	−22.49	−22.54	−22.49	−22.54	−11.40	−11.37
$i_t^{(1w)} - i_t^{(2w)}$	0.68	0.68	0.62	0.63	0.62	0.63	0.10	0.10
$i_t^{(1w)} - i_t^{(3w)}$	0.99	0.99	0.95	0.96	0.95	0.96	0.21	0.21
$i_t^{(1m)} - i_t^{(2m)}$	0.80	0.80	0.76	0.76	0.76	0.76	0.41	0.41
$i_t^{(1m)} - i_t^{(3m)}$	1.20	1.20	1.10	1.10	1.10	1.10	0.63	0.63

The table reports the in-sample performance fees $\mathcal{F}$ and the risk-adjusted abnormal returns $\mathcal{M}$ for the $\mathcal{DEH}$ strategy against the $\mathcal{EH}$ strategy when the data generating process used for bias-correction assumes homoskedastic innovations. Panel A (B) reports the performance measures when the target portfolio volatility is set to 1% (2%) for all pairwise combinations of short-term $i_t^{(m)}$ and long-term $i_t^{(n)}$repo rates such that $k = n/m$ is an integer. Each strategy is consistent with an optimizing investor allocating capital in two assets: the long-term repo rate, known with certainty at the time of trading, and a risky return generated by rolling the short-term asset for k periods. The $\mathcal{EH}$ strategy assumes that he EH holds exactly and uses the conditional forecasts implied by the constrained VAR. The $\mathcal{DEH}$ strategy conditions on the departures from the EH and uses the conditional forecasts implied by the unconstrained VAR. The performance fees $\mathcal{F}$ denote the amount an investor with quadratic utility and a degree of relative risk aversion equal to 5 is willing to pay for switching from the benchmark strategy $\mathcal{EH}$ to the alternative strategy $\mathcal{DEH}$. The risk-adjusted abnormal return, $\mathcal{M}$, defines the outperformance of the $\mathcal{DEH}$ strategy over the $\mathcal{EH}$ strategy if they had the same level of risk. We consider four different scenarios: *case* 1 (zero transaction costs and no short selling constraints); *case* 2 (non-zero transaction costs and no short selling constraints); *case* 3 (non-zero transaction costs and limited short-selling between -1 and 2); and *case* 4 (non-zero transaction costs and no short-selling). All the performance measures are reported in annual basis points.

Table 2.2
Economic Value Results with GARCH Innovations

Panel A: $\sigma_p^* = 1\%$								
	Case 1		*Case* 2		*Case* 3		*Case* 4	
$i_t^{(m)} - i_t^{(n)}$	$\mathcal{F}$	$\mathcal{M}$	$\mathcal{F}$	$\mathcal{M}$	$\mathcal{F}$	$\mathcal{M}$	$\mathcal{F}$	$\mathcal{M}$
$i_t - i_t^{(1w)}$	0.55	0.55	−1.35	−1.35	−1.35	−1.35	−0.94	−0.92
$i_t - i_t^{(2w)}$	0.02	0.02	−2.44	−2.42	−2.44	−2.42	−1.20	−1.11
$i_t - i_t^{(3w)}$	0.02	0.02	−3.81	−3.63	−3.81	−3.63	−1.89	−1.68
$i_t - i_t^{(1m)}$	0.52	0.52	−5.76	−5.50	−5.76	−5.50	−2.17	−2.82
$i_t - i_t^{(2m)}$	0.57	0.57	−8.66	−8.68	−8.66	−8.68	−3.06	−3.02
$i_t - i_t^{(3m)}$	0.86	0.86	−11.87	−11.91	−11.87	−11.91	−5.54	−5.65
$i_t^{(1w)} - i_t^{(2w)}$	0.23	0.23	0.21	0.21	0.21	0.21	0.02	0.02
$i_t^{(1w)} - i_t^{(3w)}$	0.26	0.26	0.23	0.24	0.23	0.24	0.05	0.05
$i_t^{(1m)} - i_t^{(2m)}$	0.23	0.23	0.20	0.19	0.20	0.19	0.12	0.12
$i_t^{(1m)} - i_t^{(3m)}$	0.26	0.26	0.23	0.24	0.23	0.24	0.14	0.14
Panel B: $\sigma_p^* = 2\%$								
	Case 1		*Case* 2		*Case* 3		*Case* 4	
$i_t^{(m)} - i_t^{(n)}$	$\mathcal{F}$	$\mathcal{M}$	$\mathcal{F}$	$\mathcal{M}$	$\mathcal{F}$	$\mathcal{M}$	$\mathcal{F}$	$\mathcal{M}$
$i_t - i_t^{(1w)}$	1.10	1.10	−3.78	−3.79	−3.78	−3.79	−1.44	−1.38
$i_t - i_t^{(2w)}$	0.04	0.04	−6.26	−6.27	−6.26	−6.27	−2.39	−2.11
$i_t - i_t^{(3w)}$	0.03	0.03	−7.79	−7.39	−7.79	−7.39	−3.78	−3.07
$i_t - i_t^{(1m)}$	1.03	1.03	−11.67	−11.10	−11.67	−11.10	−5.35	−5.12
$i_t - i_t^{(2m)}$	1.14	1.14	−17.28	−17.27	−17.28	−17.27	−6.12	−6.15
$i_t - i_t^{(3m)}$	1.72	1.71	−21.71	−21.69	−21.71	−21.69	−11.07	−11.84
$i_t^{(1w)} - i_t^{(2w)}$	0.45	0.45	0.39	0.40	0.39	0.40	0.03	0.03
$i_t^{(1w)} - i_t^{(3w)}$	0.52	0.52	0.48	0.48	0.48	0.48	0.10	0.10
$i_t^{(1m)} - i_t^{(2m)}$	0.46	0.46	0.39	0.39	0.39	0.39	0.24	0.24
$i_t^{(1m)} - i_t^{(3m)}$	0.51	0.51	0.46	0.46	0.46	0.46	0.28	0.28

The table reports the in-sample performance fees $\mathcal{F}$ and the risk-adjusted abnormal returns $\mathcal{M}$ for the $\mathcal{DEH}$ strategy against the $\mathcal{EH}$ strategy when the data generating process used for bias-correction assumes GARCH innovations. Panel A (B) reports the performance measures when the target portfolio volatility is set to 1% (2%) for all pairwise combinations of short-term $i_t^{(m)}$ and long-term $i_t^{(n)}$repo rates such that $k = n/m$ is an integer. Each strategy is consistent with an optimizing investor allocating capital in two assets: the long-term repo rate, known with certainty at the time of trading, and a risky return generated by rolling the short-term asset for k periods. The $\mathcal{EH}$ strategy assumes that he EH holds exactly and uses the conditional forecasts implied by the constrained VAR. The $\mathcal{DEH}$ strategy conditions on the departures from the EH and uses the conditional forecasts implied by the unconstrained VAR. The performance fees $\mathcal{F}$ denote the amount an investor with quadratic utility and a degree of relative risk aversion equal to 5 is willing to pay for switching from the benchmark strategy $\mathcal{EH}$ to the alternative strategy $\mathcal{DEH}$. The risk-adjusted abnormal return, $\mathcal{M}$, defines the outperformance of the $\mathcal{DEH}$ strategy over the $\mathcal{EH}$ strategy if they had the same level of risk. We consider four different scenarios: *case* 1 (zero transaction costs and no short selling constraints); *case* 2 (non-zero transaction costs and no short selling constraints); *case* 3 (non-zero transaction costs and limited short-selling between -1 and 2); and *case* 4 (non-zero transaction costs and no short-selling). All the performance measures are reported in annual basis points.

3 An Economic Evaluation of Empirical Exchange Rate Models: Robust Evidence of Predictability and Volatility Timing

3.1 Introduction

Forecasting exchange rates using models which condition on economically meaningful variables has long been at the top of the research agenda in international finance, and yet empirical success remains elusive. Starting with the seminal contribution of Meese and Rogoff (1983), a vast body of empirical research finds that models which condition on economic fundamentals cannot outperform a naive random walk model. Even though there is some evidence that exchange rates and fundamentals comove over long horizons (e.g. Mark, 1995; Mark and Sul, 2001), the prevailing view in international finance research is that exchange rates are not predictable, especially at short horizons.

A separate yet related literature finds that forward exchange rates contain valuable information for predicting spot exchange rates. In theory, the relation between spot and forward exchange rates is governed by the Uncovered Interest Parity (UIP) condition, which suggests that the forward premium must be perfectly positively related to future exchange rate changes. In practice, however, this is not the case as we empirically observe a negative relation.[26] The result of the empirical failure of UIP is that conditioning on the forward premium often generates exchange rate predictability. For example, Backus, Gregory and Telmer (1993) and Backus, Foresi and Telmer (2001) explore this further and find evidence of predictability using the

[26]See, for example, Bilson (1981), Fama (1984), Froot and Thaler (1990), and Backus, Foresi and Telmer (2001). For a survey of this literature, see Lewis (1995), Engel (1996) and the references therein.

lagged forward premium as a predictive variable. Furthermore, Clarida, Sarno, Taylor and Valente (2003, 2006) and Boudoukh, Richardson and Whitelaw (2006) show that the term structure of forward exchange (and interest) rates contains valuable information for forecasting spot exchange rates.

On the methodology side, while there is extensive literature on statistical measures of the accuracy of exchange rate forecasts, there is little work assessing the economic value of exchange rate predictability. Relevant research to date comprises an early study by West, Edison and Cho (1993) which provides a utility-based evaluation of exchange rate volatility, and more recently, Abhyankar, Sarno and Valente (2005) who use a similar method for investigating long-horizon exchange rate predictability. However, in the context of dynamic asset allocation strategies, there is no study assessing the economic value of the predictive ability of empirical exchange rate models which condition on economic fundamentals or the forward premium while allowing for volatility timing.

Our empirical investigation attempts to fill this gap and connect the related literatures which examine the performance of empirical exchange rate models. We do this by employing a range of economic and Bayesian statistical criteria for performing a comprehensive assessment of the short-horizon, in-sample and out-of-sample, predictive ability of three sets of models for the conditional mean of monthly nominal exchange rate returns. These models include the naive random walk model, the monetary fundamentals model (in three variants), and the spot-forward regression model. Each of the models is studied under three volatility specifications: constant variance (standard linear regression), GARCH(1,1) and stochastic volatility (SV).

In total, we evaluate the performance of 15 specifications, which encompass the most popular empirical exchange rate models studied in prior research. Our analysis employs monthly returns data ranging from January 1976 to December 2004 for three major US dollar exchange rates: the UK pound sterling, the Deutsch mark/euro, and the Japanese yen.

In addition to implementing Bayesian statistical methods for evaluating the models, an important contribution of our analysis is the use of economic criteria. Statistical evidence of exchange rate predictability in itself does not guarantee that an investor can earn profits from an asset allocation strategy that exploits this predictability. In practice, ranking models is useful to an investor only if it leads to tangible economic gains. Therefore, we assess the economic value of exchange rate predictability by evaluating the impact of predictable changes in the conditional foreign exchange (FX) returns and volatility on the performance of dynamic allocation strategies. We employ mean-variance analysis as a standard measure of portfolio performance and apply quadratic utility, which allows us to quantify how risk aversion affects the economic value of predictability, building on empirical studies of volatility timing in stock returns by Fleming, Kirby, and Ostdiek (2001) and Marquering and Verbeek (2004).[27] Ultimately, we measure how much a risk averse investor is willing to pay for switching from a dynamic portfolio strategy based on the random walk model to one which conditions on either monetary fundamentals or forward premia and has a dynamic volatility specification.

[27] For studies of asset return predictability following this approach see also Kandell and Stambaugh (1996), Barberis (2000), Baks, Metrick and Wachter (2001), Bauer (2001), Shanken and Tamayo (2001), Avramov (2002), and Cremers (2002). Karolyi and Stulz (2003) provide a survey of asset allocation in an international context.

Furthermore, we assess the statistical evidence on exchange rate predictability in a Bayesian framework. In particular, we rank the competing model specifications by computing the posterior probability of each model. The posterior probability is based on the marginal likelihood and hence it accounts for parameter uncertainty, while imposing a penalty for lack of parsimony (higher dimension). In the context of this Bayesian methodology, an alternative approach to determining the best model available is to form combined forecasts which exploit information from the entire universe of model specifications under consideration. Specifically, we implement the Bayesian Model Averaging (BMA) method, which weighs all conditional mean and volatility forecasts by the posterior probability of each model. We then compare the BMA results to those obtained from a Deterministic Model Averaging (DMA) strategy, which simply combines all model specifications with equal weights.

To preview our key results, we find strong economic and statistical evidence against the naive random walk benchmark with constant variance innovations. In particular, while conditioning on monetary fundamentals has no economic value either in-sample or out-of-sample, we establish that the predictive ability of forward exchange rate premia has substantial economic value in a dynamic portfolio allocation strategy, and that stochastic volatility significantly outperforms the constant variance and GARCH(1,1) models irrespective of the conditional mean specification. This leads to the conclusion that the best empirical exchange rate model is a model that exploits the information in the forward market for the prediction of conditional exchange rate returns and allows for stochastic volatility for the prediction of exchange rate volatility. We also provide evidence that combined forecasts which are formed using either DMA or BMA substantially outperform the random walk benchmark. These results are robust to reasonably high transaction costs and hold for all currencies both in-sample and out-of-sample. Finally, these findings have clear implications for international asset allocation strategies which are subject to FX risk.

The remainder of the chapter is organised as follows. In the next section we briefly review the relevant literature on exchange rate predictability using either fundamentals or forward exchange premia as conditioning information. Section 3.3 lays out the competing empirical models for the conditional mean and volatility of exchange rate returns. Section 3.4 describes the data, whereas Section 3.5 discusses the framework for assessing the economic value of exchange rate predictability for a risk averse investor with a dynamic portfolio allocation strategy. Section 3.6 provides a sketch of the Bayesian estimation tools, discusses the approach to model selection, and explains the construction of combined forecasts using methods such as BMA. Our empirical results are reported in Section 3.7, followed by robustness checks in Section 3.8. Finally, Section 3.9 concludes.

3.2 Stylised Facts on Exchange Rate Predictability

In this section we briefly review the theoretical and empirical research that motivates our conditioning on lagged monetary fundamentals and forward premia in the set of empirical exchange rate models.

3.2.1 Exchange Rates and Monetary Fundamentals

There is extensive literature in international finance which studies the relation between nominal exchange rates and monetary fundamentals and focuses on the following predictive variable, x_t:

$$x_t = z_t - s_t \tag{3.1}$$

$$z_t = (m_t - m_t^*) - \rho\,(y_t - y_t^*) \tag{3.2}$$

where s_t is the log of the nominal exchange rate (defined as the domestic price of foreign currency); m_t is the log of the money supply; y_t is the log of national income; asterisks denote variables of the foreign country; note that long-run money neutrality is imposed (as the coefficient on $m_t - m_t^*$ is unity as predicted by conventional theories of exchange rate determination) and ρ is a scalar that is common across countries.

Theories of exchange rate determination view z_t as the core set of economic fundamentals that determine the long-run equilibrium exchange rate. These theories include traditional models based on aggregate demand functions (e.g. Mark, 1995, and the references therein), and representative-agent general equilibrium models (e.g. Lucas, 1982; Obstfeld and Rogoff, 1995). The relation between exchange rates and fundamentals defined in Equations 3.1 and 3.2 suggests that a deviation of the nominal exchange rate, s_t, from its long-run equilibrium level determined by the fundamentals, z_t (i.e. $x_t \neq 0$), requires the exchange rate to move in the future so as to converge towards its long-run equilibrium. In other words, the deviation x_t has predictive power on future realizations of the exchange rate.[28]

In Equation 3.2 it is often assumed for simplicity that $\rho = 1$.[29] This implies

[28] Engel and West (2005) show that x_t will not have predictive power if the discount factor of future fundamentals in the exchange rate pricing condition is close to unity. This condition can be written as $s_t = (1-b)\Sigma_{i=0}^{\infty} b^i E_t z_{t+i} = (1-b)E_t z_t + bE_t \Delta s_{t+1}$, and it implies a predictive regression of the form $\Delta s_{t+1} = \frac{1-b}{b}(s_t - z_t) + \varepsilon_{t+1}$, where $\varepsilon_{t+1} \equiv (1-b)\Sigma_{i=0}^{\infty} b^i (E_{t+1} - E_t) z_{t+1+i}$ and it is assumed that $E_t z_t \equiv z_t$. If $b \approx 1$ and z_t is nonstationary, then the exchange rate predictability to be detected empirically will be low even if the fundamentals model is correct.

[29] See Mark and Sul (2001) for a detailed discussion of the pros and cons of assuming $\rho = 1$, and

that bilateral differences in real income are equally important to monetary factors in predicting exchange rates. In monetary models of exchange rate determination, both under flexible and sticky prices, ρ is interpreted as the income elasticity of money demand, and hence $0 \leq \rho \leq 1$ (Sarno and Taylor, 2003, Ch. 4). In general equilibrium models (e.g. Lucas, 1982), ρ depends on preference parameters. In these models, some utility functions can imply a negative value for ρ in very special cases, but the upper bound of ρ remains at unity. More importantly, in assessing exchange rate predictability Mark (1995) experiments with a range of values for ρ and finds that the results are very similar to the case of $\rho = 1$. Therefore, following Mark (1995) and the vast majority of papers in this literature, we set $\rho = 1$ throughout this chapter.

Despite the appeal of the theoretical relation between exchange rates and fundamentals, the empirical evidence is mixed. On the one hand, short-run exchange rate variability appears to be disconnected from the underlying fundamentals (Mark, 1995) in what is commonly referred to as the "exchange rate disconnect puzzle". On the other hand, some recent empirical research finds that fundamentals and nominal exchange rates move together *in the long run* (Groen, 2000; Berkowitz and Giorgianni, 2001; Mark and Sul, 2001; Rapach and Wohar, 2002). Either way, our study contributes to the empirical literature on the predictive ability of monetary fundamentals on exchange rates by providing an economic evaluation of the in-sample and out-of-sample forecasting power of fundamentals at a *short* (one-month ahead) horizon.

for further citations of papers providing support for this assumption.

3.2.2 The Spot-Forward Exchange Rate Relation

Assuming risk neutrality and rational expectations, Uncovered Interest Parity (UIP) is the cornerstone condition for FX market efficiency. For a one-period horizon, UIP is represented by the following equation:

$$E_{t-1}\Delta s_t = i_{t-1} - i^*_{t-1} \tag{3.3}$$

where i_{t-1} and i^*_{t-1} are the one-period domestic and foreign nominal interest rates respectively; and $\Delta s_t \equiv s_t - s_{t-1}$.

In the absence of riskless arbitrage, Covered Interest Parity (CIP) holds and implies:

$$f_{t-1} - s_{t-1} = i_{t-1} - i^*_{t-1} \tag{3.4}$$

where f_{t-1} is the log of the one-period forward exchange rate (i.e. the rate agreed now for an exchange of currencies in one period). Substituting the interest rate differential $i_{t-1} - i^*_{t-1}$ in Equation 3.3 by the forward premium (or forward discount) $f_{t-1} - s_{t-1}$, we can estimate the following regression, which is commonly referred to as the "Fama regression" (Fama, 1984a):

$$\Delta s_t = \alpha + \beta \left(f_{t-1} - s_{t-1} \right) + u_t \tag{3.5}$$

where u_t is a disturbance term.

If UIP holds, we should find that $\alpha = 0$, $\beta = 1$, and the disturbance term u_t is uncorrelated with information available at time $t-1$. Despite the increasing sophistication of the econometric techniques implemented and the improving quality of the data sets utilised, empirical studies estimating the Fama regression consistently reject the UIP condition (Hodrick, 1987; Lewis, 1995; Engel, 1996). As a result, it is now a stylised fact that estimates of β tend to be closer to minus unity than plus unity (Froot and Thaler, 1990). The negative value of β is the defining feature of what is commonly referred to as the "forward bias puzzle," namely the tendency of high-interest currencies to appreciate when UIP would predict them to depreciate.[30]

Attempts to explain the forward bias puzzle using models of risk premia have met with limited or mixed success, especially for plausible degrees of risk aversion (e.g. Engel, 1996, and the references therein). Moreover, it has proved difficult to explain the rejection of UIP by resorting to a range of proposed explanations, including learning, peso problems and bubbles (e.g. Lewis, 1995); consumption-based asset pricing theories, which allow for departures from both time-additive preferences (Backus, Gregory and Telmer, 1993; Bansal, Gallant, Hussey and Tauchen, 1995; Bekaert, 1996) and from expected utility (Bekaert, Hodrick and Marshall, 1997); and using popular models of the term structure of interest rates adapted to a multi-currency setting (Backus, Foresi and Telmer, 2001). In conclusion, even with the benefit of twenty years of hindsight, the forward bias has not been convincingly explained and remains a puzzle in international finance research.

In this context, the objective of this chapter is neither to find a novel resolution to the forward bias puzzle nor to discriminate among competing explanations.

[30] Exceptions to this puzzle include Bansal (1997), who finds that the forward bias is related to the sign of the interest rate differential; Bansal and Dahlquist (2000), who document that the forward bias is largely confined to developed economies and countries where the interest rate is lower than the US; and Bekaert and Hodrick (2001), who provide a "partial rehabilitation" of UIP by accounting for small-sample distortions. See also Lustig and Verdelhan (2007) for a more recent attempt to explain the forward bias puzzle focusing on the cross-sectional properties of foreign currency risk premia.

Instead, we focus on predicting short-horizon exchange rate returns when conditioning on the lagged forward premium, thus empirically exploiting the forward bias reported in the strand of literature stemming from Bilson (1981), Fama (1984a), Bekaert and Hodrick (1993) and Backus, Gregory and Telmer (1993). For example, Bilson (1981) argues that regressions conditioning on the forward premium can potentially yield substantial economic returns, whereas arguments based on limits to speculation would suggest otherwise (Lyons, 2001; Sarno, Valente and Leon, 2006). Furthermore, term structure models that exploit departures from UIP often yield accurate out-of-sample forecasts (e.g. Clarida and Taylor, 1997; Clarida, Sarno, Taylor and Valente, 2003; Boudoukh, Richardson and Whitelaw, 2006). However, little attention has been given to the question of whether the statistical rejection of UIP and the forward bias resulting from the negative estimate of β offers economic value to an international investor facing FX risk. Our chapter fills this void in the literature by assessing the economic value of the predictive ability of empirical exchange rate models which condition on the forward premium in the context of dynamic asset allocation strategies.

3.3 Modeling FX Returns and Volatility

In this section we present the candidate models applied to monthly exchange rate returns in our study of short-horizon exchange rate predictability. We use a set of specifications for the dynamics of both the conditional mean and volatility, which are set against the naive random walk benchmark. In short, we estimate five conditional mean and three conditional volatility specifications yielding a total of 15 models for each of the three dollar exchange rates under consideration.

3.3.1 The Conditional Mean

We examine five conditional mean specifications in which the dynamics of exchange rate returns are driven by the following regression:

$$\Delta s_t = \alpha + \beta x_{t-1} + u_t, \quad u_t = v_t \varepsilon_t, \quad \varepsilon_t \sim NID\,(0,1)\,. \tag{3.6}$$

Our first specification is the naive random walk (RW) model, which sets $\beta = 0$. This model is the standard benchmark in the literature on exchange rate predictability since the seminal work of Meese and Rogoff (1983).

The next three model specifications condition on monetary fundamentals (MF). Specifically, MF_1 uses the canonical version $x_t = z_t - s_t$ as defined in Equations 3.1 and 3.2 assuming $\rho = 1$. This is the most common formulation of the monetary fundamentals model since Mark (1995). The second variant of the monetary fundamentals model, MF_2, corrects for the deterministic component in the deviation of the exchange rate from monetary fundamentals by allowing for an intercept and a slope parameter; in other words, we run the ordinary least squares (OLS) regression $s_t = \kappa_0 + \kappa_1 z_t + \zeta_t$, and set $x_t = -\widehat{\zeta}_t$. The third variant, MF_3, further corrects for the time trend in fundamentals deviations; in this case, we run the OLS regression $s_t = \kappa_0 + \kappa_1 z_t + \kappa_2 t + \zeta_t$, where t is a simple time trend, and again we set $x_t = -\widehat{\zeta}_t$. The motivation behind the MF_2 and MF_3 variants derives from empirical evidence that cointegration between s_t and z_t will sometimes be established only by correcting for the deterministic components (either a constant or a constant and a time trend) in the cointegrating residual (e.g. Rapach and Wohar, 2002). Note, however, that in the out-of-sample exercise in this chapter we estimate the deterministic component recursively as we move through the data sample, and hence our results do not suffer from "look-ahead bias". We do not report cointegration tests nor search for the best possible specification of the long-run relation between exchange rates and monetary fundamentals since the focus of this chapter is on measuring the economic value of predictability due to monetary fundamentals, not on understanding the determinants of the long-run exchange rate equilibrium. In short, we make no prior assumptions on the best formulation of the MF model by considering all three specifications proposed in the literature.

Finally, the fifth conditional mean specification is the forward premium (FP) model, which sets $x_t = f_t - s_t$ as in Equation 3.5 resulting in the Fama (1984a) regression. The FP model stems directly from the spot-forward exchange rate relation derived from UIP. Hence it constitutes the empirical model which exploits the forward bias and allows us to assess the economic value of conditioning on the forward premium in the context of dynamic asset allocation strategies. The forward bias (a negative estimate of the β coefficient in the FP model) implies that the more the foreign currency is at a premium in the forward market, the less the home currency is expected to depreciate. Equivalently, the more domestic interest rates exceed foreign interest rates, the more the domestic currency tends to appreciate over the holding period.

3.3.2 The Conditional Variance

We model the dynamics of the conditional variance by implementing three models: the simple linear regression (LR), the GARCH(1,1) model, and the stochastic volatility (SV) model. The linear regression framework simply assumes that the conditional variance of FX return innovations is constant over time ($v_t^2 = v^2$), and therefore presents the benchmark against which models with time-varying conditional variance will be evaluated.

The benchmark GARCH(1,1) model of Bollerslev (1986) is defined as:

$$v_t^2 = \omega + \gamma_1 v_{t-1}^2 + \gamma_2 u_{t-1}^2. \tag{3.7}$$

In this formulation, $u_t \mid \mathcal{F}_{t-1} \sim N(0, v_t^2)$, and therefore the conditional variance v_t^2 is time-varying but deterministic given the information set $\mathcal{F}_{t-1}$. Furthermore, the stationarity and positive variance restrictions impose the following conditions: $\omega > 0$, $\gamma_1 \geq 0$, $\gamma_2 \geq 0$, and $\gamma_1 + \gamma_2 < 1$. The main contribution of this simple GARCH specification is that it models volatility clustering by allowing for a persistent, and hence predictable conditional variance. Our motivation for studying the simple GARCH(1,1) model is based on the early study of West, Edison and Cho (1993), which performs a utility-based evaluation of exchange rate volatility and finds that GARCH(1,1) is the best performing model.

Stochastic volatility models are similar to the GARCH process in that they capture the persistent and hence predictable component of volatility. Unlike GARCH models, however, the assumption of a stochastic second moment introduces an additional source of risk that cannot be perfectly hedged using $t-1$ information. A GARCH specification describes the conditional distribution of returns as being exclusively a function of past information. In contrast, the SV model specifies the joint conditional distribution of both the return and the volatility process.[31] Intuitively,

[31] For details on SV models see Kim, Shephard and Chib (1998) and Chib, Nardari and Shephard (2002). For an application of SV models to exchange rates, see Harvey, Ruiz and Shephard (1994). Finally, for a comparison of GARCH and SV models see Fleming and Kirby (2003).

SV allows for the possibility of random contemporaneous volatility shocks due to news events and policy changes – in other words, unobserved contemporaneous variables that may affect the volatility process.[32]

According to the plain vanilla SV model, the persistence of the conditional volatility v_t is captured by the dynamics of the Gaussian stochastic log-variance process h_t:

$$v_t = \exp\left(h_t/2\right) \tag{3.8}$$

$$h_t = \mu + \phi\left(h_{t-1} - \mu\right) + \sigma\eta_t, \quad \eta_t \sim NID\left(0, 1\right). \tag{3.9}$$

In the SV model, return and volatility innovations are independent: $\{\varepsilon_t\} \perp \{\eta_t\}$. Furthermore, the model assumes (and the estimation algorithm imposes) $|\phi| < 1$ so that the log-variance is a stationary process.

3.4 FX Data and Descriptive Statistics

The data sample consists of 348 monthly observations ranging from January 1976 to December 2004, and focuses on three exchange rates relative to the US dollar: the UK pound sterling (USD/GBP), Deutsch mark/euro (USD/DEM-EURO), and Japanese yen (USD/JPY). The spot and one-month forward exchange rates are taken from *Datastream* for the period of January 1985 onwards, whereas for the period ranging from January 1976 to December 1984 they are taken from Hai, Mark and Wu (1997). After the introduction of the euro in January 1999, we use the euro

[32]In fact, market microstructure theories of speculative trading (e.g. Tauchen and Pitts, 1983; Andersen, 1996) provide rigorous arguments for modeling volatility as stochastic.

exchange rate to replace the Deutsch mark rate.

Data on money supply and income are from the International Monetary Fund's *International Financial Statistics* database. Specifically, we define the money supply as the sum of money (line code 34) and quasi-money (line code 35) for Germany and Japan, whereas for the UK we use $M0$ (line code 19). Since German exchange rate data are only available until December 1998, we use the money and quasi-money data of the Euro Area for the remaining period (January 1999 to December 2004). The US data is obtained from the aggregate $M2$ of the Board of Governors of the Federal Reserve System. Furthermore, we use the monthly industrial production index (line code 66) as a proxy for national income rather than the gross domestic product (GDP), because the latter is available only at the quarterly frequency.[33] We deseasonalise the money and industrial production indices following the procedure of Gomez and Maravall (2000). Note that we ignore the complication arising from the fact that the data we use on monetary fundamentals may not be available in real time and may not suffer from the measurement errors that characterise real-time macroeconomic data (Faust, Rogers and Wright, 2003). This issue will not affect our main findings on the predictive ability of the forward premium and stochastic volatility.

We take logarithmic transformations of the raw data to yield time series for s_t, f_t, m_t, m_t^*, y_t, and y_t^*. The monetary fundamentals series, z_t, is constructed as in Equation 3.2 imposing $\rho = 1$; s_t is taken as the natural logarithm of the domestic price of foreign currency, the US being the domestic country; f_t is the natural

[33]For all countries, the correlation coefficient between the quarterly industrial production index and GDP over our sample period is higher than 0.95.

logarithm of the US dollar price of a one-month forward contract issued at time t for delivery of one unit of foreign currency at time $t+1$. Finally, in our economic evaluation of the set of candidate exchange rate models, the proxy for the riskless domestic and foreign bonds is the end-of-month Euromarket interest rate with one month maturity, obtained from *Datastream*.[34]

Table 3.1 reports the descriptive statistics for the monthly percent FX returns, Δs_t, the three monetary fundamentals predictors, MF_1, MF_2, and MF_3, also expressed in percent, and the percent forward premium, $f_t - s_t$. For our sample period, the sample means of the FX returns are -0.012% for USD/GBP, 0.165% for USD/DEM-EURO, and 0.309% for USD/JPY. The FX return standard deviations are similar across the three exchange rates at about 3% per month. Finally, the exchange rate return sample autocorrelations are approximately 0.10 but decay rapidly.

The three specifications of monetary fundamentals predictors display very high volatility and persistence. For instance, the standard deviation of MF_1 is about 20% for the UK, 30% for Germany and 40% for Japan. However, the standard deviation of MF_3 (which is corrected for both the deterministic and the time trend component) is approximately half the value of the canonical monetary fundamentals MF_1. The three monetary fundamentals predictors exhibit little skewness and excess kurtosis. The sample autocorrelation coefficient is very high for all three specifications and decreasing slowly.

Finally, the average forward premium is negative for the UK, but positive for

[34]We use the Eurocurrency deposit rate as a proxy for the riskless rate because these deposits are comparable across countries in all respects (such as issuer, credit risk and maturity) except for currency of denomination; see Levich (1985).

Germany and Japan. The standard deviation of $f_t - s_t$ is low across all exchange rates (in fact, about 100 times smaller than MF_1), but the forward premium exhibits high kurtosis and its sample autocorrelation is high and decreasing slowly.

3.5 Measuring the Economic Value of Exchange Rate Predictability

This section discusses the framework we use in order to evaluate the impact of predictable changes in both exchange rate returns and volatility on the performance of dynamic allocation strategies. We employ mean-variance analysis as a standard measure of portfolio performance and apply quadratic utility, which allows us to quantify how risk aversion affects economic value. Ultimately, we aim at measuring how much an investor is willing to pay for switching from the naive random walk strategy that assumes no predictability in exchange rates to a dynamic strategy which conditions on monetary fundamentals or the forward premium and allows for time-varying volatility.

3.5.1 FX Models in a Dynamic Mean-Variance Framework

In mean-variance analysis, the maximum expected return strategy leads to a portfolio allocation on the efficient frontier. Specifically, consider an investor who has a one-month horizon and constructs a dynamically rebalanced portfolio that maximises the conditional expected return subject to achieving a target conditional volatility. Computing the time-varying weights of this portfolio requires one-step ahead forecasts of the conditional mean and the conditional variance-covariance matrix. Let r_{t+1} denote the $K \times 1$ vector of risky asset returns; $\mu_{t+1|t} = E_t[r_{t+1}]$ is the conditional expectation of r_{t+1}; and $\Sigma_{t+1|t} = E_t\left[\left(r_{t+1} - \mu_{t+1|t}\right)\left(r_{t+1} - \mu_{t+1|t}\right)'\right]$

is the conditional variance-covariance matrix of r_{t+1}. At each period t, the investor solves the following problem:

$$\max_{w_t} \left\{\mu_{p,t+1} = w_t'\mu_{t+1|t} + \left(1 - w_t'\iota\right) r_f\right\}$$
$$\text{s.t. } \left(\sigma_p^*\right)^2 = w_t'\Sigma_{t+1|t}w_t \tag{3.10}$$

where w_t is the $K \times 1$ vector of portfolio weights on the risky assets; ι is a $K \times 1$ vector of ones; $\mu_{p,t+1}$ is the conditional expected return of the portfolio; σ_p^* is the target conditional volatility of the portfolio returns; and r_f is the return on the riskless asset. The solution to this optimization problem delivers the risky asset weights:

$$w_t = \frac{\sigma_p^*}{\sqrt{C_t}}\Sigma_{t+1|t}^{-1}\left(\mu_{t+1|t} - \iota r_f\right) \tag{3.11}$$

where $C_t = \left(\mu_{t+1|t} - \iota r_f\right)'\Sigma_{t+1|t}^{-1}\left(\mu_{t+1|t} - \iota r_f\right)$. The weight on the riskless asset is $1 - w_t'\iota$.

Constructing the optimal portfolio weights requires estimates of the conditional expected returns, variances and covariances. We consider five conditional mean strategies (RW, MF_1, MF_2, MF_3, and FP) and three conditional volatility strategies (LR, $GARCH$, and SV) for a total of 15 sets of one-step ahead conditional expected return and volatility forecasts. The conditional covariances are computed using the constant conditional correlation (CCC) model of Bollerslev (1990), in which the dynamics of covariances are driven by the time-variation in the conditional volatilities.[35] By design, in this setting the optimal weights will vary across

[35] In notation local to this footnote, the CCC model of Bollerslev (1990) specifies the covariances as follows: $\sigma_{ij,t} = \sigma_{i,t}\sigma_{j,t}\rho_{ij}$, where $\sigma_{i,t}$ and $\sigma_{j,t}$ are the conditional volatilities implied by either

models only to the extent that forecasts of the conditional mean and volatility will vary, which is precisely what the empirical models provide. The benchmark against which we compare the model specifications is the random walk model with constant variance (RW^{LR}). In short, our objective is to determine whether there is economic value in (i) conditioning on lagged monetary fundamentals and, if so, which of the three specifications works best, (ii) conditioning on the lagged forward premium, (iii) using a GARCH volatility specification, and (iv) implementing an SV process for the monthly FX innovations.

3.5.2 Quadratic Utility

Mean-variance analysis is a natural framework for assessing the economic value of strategies which exploit predictability in the mean and variance. In particular, we rank the performance of the competing FX models using the West, Edison and Cho (1993) methodology, which is based on mean-variance analysis with quadratic utility. The investor's realised utility in period $t+1$ can be written as:

$$U\left(W_{t+1}\right) = W_{t+1} - \frac{\lambda}{2} W_{t+1}^2 = W_t R_{p,t+1} - \frac{\lambda W_t^2}{2} R_{p,t+1}^2 \tag{3.12}$$

where W_{t+1} is the investor's wealth at $t+1$, λ determines his risk preference, and

$$R_{p,t+1} = 1 + r_{p,t+1} = 1 + \left(1 - w_t' \iota\right) r_f + w_t' r_{t+1} \tag{3.13}$$

is the period $t+1$ gross return on his portfolio.

the GARCH(1,1) or the SV process, and ρ_{ij} is the constant sample correlation coefficient. Note that for the out-of-sample results we use a rolling correlation estimate updated every time a new observation is added. From a numerical standpoint, implementing the CCC model is attractive because it eliminates the possibility of $\Sigma_{t+1|t}$ not being positive definite.

We quantify the economic value of exchange rate predictability by setting the investor's degree of relative risk aversion (RRA) $\delta_t = \lambda W_t / (1 - \lambda W_t)$ equal to a constant value δ. In this case, West, Edison and Cho (1993) demonstrate that one can use the average realised utility, $\overline{U}(\cdot)$, to consistently estimate the expected utility generated by a given level of initial wealth. Specifically, the average utility for an investor with initial wealth W_0 is equal to:

$$\overline{U}(\cdot) = W_0 \sum_{t=0}^{T-1} \left\{ R_{p,t+1} - \frac{\delta}{2(1+\delta)} R_{p,t+1}^2 \right\}. \tag{3.14}$$

We standardise the investor problem by assuming he allocates \$1 in every time period.

Average utility depends on taste for risk. In the absence of restrictions on δ, quadratic utility exhibits increasing RRA. This is counterintuitive since, for instance, an investor with increasing RRA becomes more averse to a percentage loss in wealth when his wealth increases. As in West, Edison and Cho (1993) and Fleming, Kirby and Ostdiek (2001), fixing the degree of RRA, δ, implies that expected utility is linearly homogeneous in wealth: double wealth and expected utility doubles. Furthermore, by fixing δ rather than λ, we are implicitly interpreting quadratic utility as an approximation to a non-quadratic utility function, with the approximating choice of λ dependent on wealth. The estimate of expected quadratic utility given in Equation 3.14 is used to implement the Fleming, Kirby and Ostdiek (2001) framework for assessing the economic value of our FX strategies in the context of dynamic asset allocation.

A critical aspect of mean-variance analysis is that it applies exactly only when the return distribution is normal *or* the utility function is quadratic. Hence, the use of quadratic utility is not necessary to justify mean-variance optimization. For instance, one could instead consider using utility functions belonging to the constant relative risk aversion (CRRA) class, such as power or log utility. However, quadratic utility is an attractive assumption because it allows us to consider non-normal distributions of returns, while remaining within the mean-variance framework as well as providing a high degree of analytical tractability.[36]

Additionally, quadratic utility may be viewed as a second order Taylor series approximation to expected utility. In an investigation of the empirical robustness of the quadratic approximation, Hlawitschka (1994) finds that a two-moment Taylor series expansion "may provide an excellent approximation" (p. 713) to expected utility and concludes that the ranking of common stock portfolios based on two-moment Taylor series is "almost exactly the same" (p. 714) as the ranking based on a wide range of utility functions.

3.5.3 Performance Measures

At any point in time, one set of estimates of the conditional mean and variance is better than a second set if investment decisions based on the first set lead to higher average realised utility, $\overline{U}$. Alternatively, the optimal model requires less wealth to yield a given level of $\overline{U}$ than a suboptimal model. Following Fleming, Kirby and Ostdiek (2001) we measure the economic value of our FX strategies by equating the average utilities for selected pairs of portfolios. Suppose, for example, that holding a

[36]In fact, assuming quadratic utility allows us to use the Fleming, Kirby and Ostdiek (2001) framework (also based on quadratic utility) for evaluating the performance of fat-tailed volatility specifications, such as the $tGARCH$ model of Bollerslev (1987).

portfolio constructed using the optimal weights based on the Random Walk/Linear Regression (RW^{LR}) model yields the same average utility as holding the Forward Premium/Stochastic Volatility (FP^{SV}) optimal portfolio that is subject to monthly expenses Φ, expressed as a fraction of wealth invested in the portfolio. Since the investor would be indifferent between these two strategies, we interpret Φ as the maximum performance fee he will pay to switch from the RW^{LR} to the FP^{SV} strategy. In other words, this utility-based criterion measures how much a mean-variance investor is willing to pay for conditioning on the lagged forward premium under stochastic volatility innovations. The performance fee will depend on the investor's degree of risk aversion. To estimate the fee, we find the value of Φ that satisfies:

$$\sum_{t=0}^{T-1} \left\{ \left(R^*_{p,t+1} - \Phi\right) - \frac{\delta}{2\,(1+\delta)} \left(R^*_{p,t+1} - \Phi\right)^2 \right\} = \sum_{t=0}^{T-1} \left\{ R_{p,t+1} - \frac{\delta}{2\,(1+\delta)} R^2_{p,t+1} \right\} \tag{3.15}$$

where $R^*_{p,t+1}$ is the gross portfolio return constructed using the expected return and volatility forecasts from the FP^{SV} model, and $R_{p,t+1}$ is the gross portfolio return implied by the benchmark RW^{LR} model.

In the context of mean-variance analysis, a commonly used measure of economic value is the Sharpe ratio. However, as suggested by Marquering and Verbeek (2004) and Han (2006), the Sharpe ratio can be misleading because it severely underestimates the performance of dynamic strategies. Specifically, the realised Sharpe ratio is computed using the sample standard deviation of the realised portfolio returns and hence it overestimates the conditional risk an investor faces at each point in time. Furthermore, the Sharpe ratio cannot quantify the exact economic gains of the dynamic strategies over the static random walk strategy in the direct way of the performance fees. Therefore, our economic analysis of short-horizon exchange rate predictability focuses primarily on performance fees, while Sharpe ratios of selected models are reported in the robustness section.[37]

3.5.4 The Dynamic FX Strategies

In this mean-variance quadratic-utility framework, we design the following global strategy. Consider a US investor who builds a portfolio by allocating his wealth between four bonds: one domestic (US), and three foreign bonds (UK, Germany and Japan). At the beginning of each month, the four bonds yield a riskless return in local currency. Hence the only risk the US investor is exposed to is FX risk. Each month the investor takes two steps. First, he uses each of the 15 models to forecast the one-month ahead conditional mean and volatility of the exchange rate returns. Second, conditional on the forecasts of each model, he dynamically rebalances his portfolio by computing the new optimal weights for the maximum return strategy. This setup is designed to inform us whether using one particular conditional mean and volatility specification affects the performance of a short-horizon allocation strategy in an economically meaningful way. The yields of the riskless bonds are proxied by monthly Eurodeposit rates.

In the context of this maximum return dynamic strategy we compute both the in-sample and the out-of-sample performance fee, Φ, where the out-of-sample period

[37]The annualized Sharpe ratios reported in Table 3.10 are adjusted for the serial correlation in the monthly portfolio returns generated by the dynamic strategies. Specifically, following Lo (2002), we multiply the monthly Sharpe ratios by the adjustment factor $\frac{12}{\sqrt{12+2\sum_{k=1}^{11}(12-k)\rho_k}}$, where ρ_k is the autocorrelation coefficient of portfolio returns at lag k.

starts in January 1990 and ends in December 2004. Furthermore, we compare the performance fees for the combinations corresponding to the following cases: (i) three sets of target annualised portfolio volatilities ($\sigma_p^* = \{8\%, 10\%, 12\%\}$); (ii) all pairs of 15 models (for example MF_1^{LR} vs. RW^{LR}; or FP^{SV} vs. RW^{SV}); and (iii) degrees of RRA $\delta = \{2, 6\}$. We report the estimates of Φ as annualised fees in basis points.[38]

3.5.5 Transaction Costs

The impact of transaction costs is an essential consideration in assessing the profitability of trading strategies. This is especially true in our case because the trading strategy based on the random walk benchmark is static (independent of state variables), whereas the remaining empirical models generate dynamic strategies.[39] Furthermore, making an accurate determination of the size of transaction costs is difficult because it involves three factors: (i) the type of investor (e.g. individual vs. institutional investor), (ii) the value of the transaction, and (iii) the nature of the broker (e.g. brokerage firm vs. direct internet trading). This difficulty is reflected in the wide range of estimates used in empirical studies. For example, Marquering and Verbeek (2004) consider three levels of transaction costs, 0.1%, 0.5% and 1%, to represent low, medium and high costs.

Our approach avoids these concerns by calculating the break-even transaction cost, τ^{BE}, that renders investors indifferent between two strategies (e.g. Han, 2006). In particular, we assume that transaction costs equal a fixed proportion (τ) of the

[38]Note that, due to lack of data for the Japanese eurocurrency interest rate, the in-sample period in our economic value results starts in January 1979. In contrast, for the statistical analysis the in-sample period starts in January 1976.

[39]The random walk model (RW^{LR}) is the only empirical model that assumes constant mean and variance. Therefore, the in-sample optimal weights for the RW^{LR} trading strategy remain constant over time. However, the out-of-sample optimal weights will vary because every month we re-estimate the drift and variance of the RW^{LR} model.

value traded in each bond: $\tau|w_t - w_{t-1}\frac{1+r_t}{1+r_{p,t}}|$. In comparing a dynamic strategy with the static (random walk) strategy, an investor who pays transaction costs lower than τ^{BE} will prefer the dynamic strategy. We report τ^{BE} in monthly basis points.[40]

3.6 Estimation and Forecasting

3.6.1 Bayesian Markov Chain Monte Carlo Estimation

Stochastic volatility models are generally less popular in empirical applications than GARCH despite their parsimonious structure, intuitive appeal and popularity in theoretical option pricing. This is primarily due to the numerical difficulty associated with estimating SV models using conventional classical econometric methods. Specifically, discrete-time SV models cannot be estimated with standard likelihood-based methods because the likelihood function is not available analytically. Bayesian estimation offers a substantial computational advantage over any classical approach because it avoids tackling difficult numerical optimization procedures. In this context, we estimate all three volatility frameworks (LR, GARCH and SV) using similar Bayesian Markov Chain Monte Carlo (MCMC) estimation algorithms. This is a crucial aspect of our econometric analysis because it renders the posterior mean estimates directly comparable across the three volatility structures. It also allows us to use the same model risk diagnostics for all model specifications. Finally, a distinct advantage of Bayesian inference is that it provides the posterior distribution of a regression coefficient conditional on the data, which holds for finite samples and regardless of whether exchange rates (and fundamentals) are (co)integrated (e.g. Sims, 1988). This is not the case in classical inference, where the small samples typically

[40]In contrast to Φ, which is reported in annual basis points, τ^{BE} is reported in monthly basis points because τ^{BE} is a proportional cost paid every month when the portfolio is rebalanced.

employed in the study of exchange rate predictability combined with the assumption that exchange rates and fundamentals are cointegrated can have a critical impact in overstating predictability (e.g. Berkowitz and Giorgianni, 2001).

We estimate the parameters of the SV model using the Bayesian MCMC algorithm of Chib, Nardari, and Shephard (2002), which builds on the procedures developed by Kim, Shephard, and Chib (1998). The algorithm constructs a Markov chain whose limiting distribution is the target posterior density of the SV parameters. The Markov chain is a Gibbs sampler in which all parameters are drawn sequentially from their full conditional posterior distribution. The Gibbs sampler is iterated 5000 times and the sampled draws, beyond a burn-in period of 1000 iterations, are treated as variates from the target posterior distribution. We design a similar Bayesian MCMC algorithm for estimating the GARCH(1,1) parameters, which also draws from the insights of Vrontos, Dellaportas and Politis (2000). The Bayesian Linear Regression algorithm implements a simple MCMC assuming an independent Normal-Gamma prior distribution (for details see Koop, 2003). The MCMC algorithm for each of the three volatility models is summarised in the Appendix B. Each algorithm produces estimates of the posterior means $\theta = \{\theta_1, \theta_2\}$, where $\theta_1 = \{\alpha, \beta\}$ are the parameters of the return equation, and θ_2 are the parameters of the volatility specification: $\theta_2 = \{v^{-2}\}$ for the Linear Regression, $\theta_2 = \{\omega, \gamma_1, \gamma_2\}$ for the $GARCH(1,1)$ specification, and $\theta_2 = \{\mu, \phi, \sigma^2\}$ for the SV model. All θ parameters are time invariant.

The mean of the MCMC parameter draws is an asymptotically efficient estimator of the posterior mean of θ (see Geweke, 1989). The Numerical Standard Error (NSE) is the square root of the asymptotic variance of the MCMC estimator:

$$NSE = \sqrt{\frac{1}{I}\left\{\widehat{\psi}_0 + 2\sum_{j=1}^{B_I} K(z)\,\widehat{\psi}_j\right\}} \tag{3.16}$$

where $I = 5000$ is the number of iterations (beyond the initial burn-in of 1000 iterations), $j = 1, ..., B_I = 500$ lags is the set bandwidth, $z = \frac{j}{B_I}$, and $\widehat{\psi}_j$ is the sample autocovariance of the MCMC draws for each estimated parameter cut according to the Parzen kernel $K(z)$. The NSE diagnostic is distinct from the MCMC standard deviation. The latter is simply a measure of the variation in the MCMC parameter draws. In contrast, the NSE is a measure of the variation in the posterior mean estimate across many MCMCs we can potentially run. In other words, the NSE measures how much difference we should expect in the estimate of the posterior mean if estimation were to be repeated, and therefore provides a measure of convergence in the Markov chain.

The likelihood function of the SV models is not available analytically, and hence must be simulated. Specifically, the log-likelihood function is evaluated under the predictive density as:

$$\log \widehat{L} = \sum_{t=1}^{T} \log \widehat{f}\,(\Delta s_t \mid \mathcal{F}_{t-1}, \theta) = \sum_{t=1}^{T} \log \widehat{f}_t\,(\Delta s_t \mid h_t, \theta) \tag{3.17}$$

where θ is taken as the posterior mean estimate from the MCMC simulations. The key to this calculation is simulating the one-step ahead predictive log-variance $h_t \mid \mathcal{F}_{t-1}, \theta$, which is a non-trivial task as it is sampled using the particle filter of Pitt and Shephard (1999). The particle filter is summarised in the appendix B. For more details see also Chib, Nardari and Shephard (2002) and Han (2006).

3.6.2 Model Risk and Posterior Probability

Model risk arises from the uncertainty over selecting a model specification. Consistent with our Bayesian approach, a natural statistical criterion for resolving this uncertainty is the posterior probability of each model. Hence, we rank the competing models using the posterior probability, which has three important advantages relative to the log-likelihood: (i) it is based on the marginal likelihood and therefore accounts for parameter uncertainty, (ii) it imposes a penalty for lack of parsimony (higher dimension), and (iii) it forms the basis of the Bayesian Model Averaging strategy discussed below. Ranking the models using the highest posterior probability is equivalent to choosing the best model in terms of density forecasts and is a robust model selection criterion in the presence of misspecification and non-nested models (e.g. Fernandez-Villaverde and Rubio-Ramirez, 2004).[41]

Consider a set of N models $M_1, ..., M_N$. We form a prior belief $\pi(M_i)$ on the probability that the ith model is the true model, observe the FX returns data Δs, and then update our belief that the ith model is true by computing the posterior probability of each model defined as follows:

$$p(M_i \mid \Delta s) = \frac{p(\Delta s \mid M_i)\, \pi(M_i)}{\sum_{j=1}^{N} p(\Delta s \mid M_j)\, \pi(M_j)} \tag{3.18}$$

where $p(\Delta s \mid M_i)$ is the marginal likelihood of the ith model defined as follows:

[41] The information one can extract from the posterior probability of a model is similar to using the Kullback-Leibler Information Criterion ($KLIC$). Specifically, Fernandez-Villaverde and Rubio-Ramirez (2004) show that choosing the model with the highest posterior probability is equivalent to selecting the best model under the $KLIC$. This is an attractive feature of our Bayesian approach because there is a complete axiomatic foundation that justifies why $KLIC$ is the best criterion a rational agent should use in choosing between models (e.g. Csiszar, 1991). See Burnham and Anderson (2002) for a review of $KLIC$.

$$p(\Delta s \mid M_i) = \int_\theta p(\Delta s, \theta \mid M_i)\, d\theta = \int_\theta p(\Delta s \mid \theta, M_i)\, \pi(\theta \mid M_i)\, d\theta \qquad (3.19)$$

In Equation 3.18 above we set our prior belief to be that all models are equally likely, i.e. $\pi(M_i) = \frac{1}{N}$.

Note that the marginal likelihood is an averaged (not a maximised) likelihood. This implies that the posterior probability is an automatic "Occam's Razor" in that it integrates out parameter uncertainty.[42] Furthermore, the marginal likelihood is simply the normalizing constant of the posterior density and (suppressing the model index for simplicity) it can be written as:

$$p(\Delta s) = \frac{f(\Delta s \mid \theta)\, \pi(\theta)}{\pi(\theta \mid \Delta s)} \qquad (3.20)$$

where $f(\Delta s \mid \theta)$ is the likelihood, $\pi(\theta)$ the prior density of the parameter vector θ, $\pi(\theta \mid \Delta s)$ the posterior density, and θ is evaluated at the posterior mean. Since θ is drawn in the context of MCMC sampling, the posterior density $\pi(\theta \mid \Delta s)$ is computed using the technique of reduced conditional MCMC runs of Chib (1995). For the θ_2 parameters in GARCH and SV, which are sampled in the MCMC chain by implementing a Metropolis-Hastings algorithm, the posterior density is computed as in Chib and Jeliazkov (2001).

3.6.3 Combined Forecasts

Assessing the predictive ability of empirical exchange rate models primarily involves a pairwise comparison of the competing models. However, given that we do not know

[42] *Occam's Razor* is the principle of parsimony, which states that among two competing theories that make exactly the same prediction, the simpler one is best.

which one of the models is true, it is important that we assess the performance of combined forecasts proposed by the seminal work of Bates and Granger (1969). Specifically, we design three strategies based on a combination of forecasts for both the conditional mean and volatility of exchange rate returns: the Deterministic Model Average (DMA) strategy, the Bayesian Model Average (BMA) strategy, and the Bayesian Winner (BW) strategy.[43]

We assess the economic value of combined forecasts by treating the DMA, BMA and BW strategies the same way as any of the 15 individual empirical models. For instance, we compute the performance fee, Φ, for the BMA one-month ahead forecasts of the conditional mean and volatility and compare them to the random walk benchmark. In particular, we focus on two distinct universes of models: the restricted universe of the five SV models (because the five conditional mean specifications with SV innovations have the highest marginal likelihood), and the unrestricted universe of all 15 empirical exchange rate models.

Consequently, our empirical analysis of exchange rate predictability and volatility timing further contributes to the literature by incorporating both a statistical view of Bayesian parameter uncertainty and an economic view of the effect of model uncertainty on asset allocation decisions and performance. In contrast to Avramov (2002), however, our approach does not attempt to separate the effects of parameter and model uncertainty. Finally, we only consider model uncertainty within the universe of the 15 model specifications implied by economic fundamentals and dynamic volatility.

[43]See Diebold and Pauly (1990), Diebold (1998, 2004), and Timmermann (2006) for a review of forecast combinations.

3.6.4 The DMA Strategy

Quite simply, the DMA strategy involves taking an equally weighted average of the conditional mean and volatility forecasts from a given universe of available models. Hence, for a set of N models the DMA strategy is referred to as the $1/N$ strategy. Since this is a strategy that does not require period-by-period updating of the weights in the forecast combination, it can be readily evaluated in-sample and out-of-sample on the basis of conditioning information available at the time of the forecast.

3.6.5 The BMA Strategy

In the context of our Bayesian approach, it is natural to implement the BMA method originally discussed in Leamer (1978), and surveyed in Hoeting, Madigan, Raftery and Volinsky (1999). The BMA strategy accounts directly for uncertainty in model selection, and is in fact easy to implement once we have the output from the MCMC simulations. Define $f_{i,t}$ as the forecast density of each of the N competing models at time t. Then, the BMA forecast density is given by:

$$f_t^{BMA} = \sum_{i=1}^{N} p_t\left(M_i \mid \Delta s_t\right) f_{i,t} \tag{3.21}$$

where $p_t\left(M_i \mid \Delta s_t\right)$ is the posterior probability of model M_i given the data Δs_t.

It is important to note that the BMA weights vary not only across models but also across time periods as does the marginal predictive density (and hence marginal likelihood) of each model. In particular, at each time period we estimate the one-step ahead predictive density $f_t\left(\Delta s_t \mid \mathcal{F}_{t-1}, \theta\right)$ and the posterior density $\pi_t\left(\theta \mid \Delta s_t\right)$. We can then compute the time-varying marginal predictive density using Equation

3.20, and insert it into Equation 3.18 to finally calculate the posterior probability of each model at each time period. It is crucial to emphasise that we evaluate the BMA strategy ex-ante. We do this by lagging the posterior probability of each model for the following reason. Suppose that we need to compute the period t BMA forecasts of the conditional mean and volatility for the four bonds we include in the portfolio. Knowing the mean and volatility forecasts implied by each model for the three exchange rates is not sufficient. We also need the realised data point Δs_t in order to evaluate the predictive density $f_t\left(\Delta s_t \mid \mathcal{F}_{t-1}, \theta\right)$. Since the realised data point Δs_t is only observed ex post, the only way to form the BMA weights ex ante is to lag the predictive density and thus use $f_{t-1}\left(\Delta s_{t-1} \mid \mathcal{F}_{t-2}, \theta\right)$. The same method is applied both in-sample and out-of-sample.

3.6.6 The BW Strategy

Under the BW strategy, in each time period we select the set of one-step ahead conditional mean and volatility from the empirical model that has the highest marginal predictive density in that period. In other words, the BW strategy only uses the forecasts of the "winner" model in terms of marginal predictive density, and hence discards the forecasts of the rest of the models. Clearly, there is no model averaging in the BW strategy. Similar to the BMA, the BW strategy is evaluated ex ante using the lagged predictive marginal densities.

3.7 Empirical Results

3.7.1 Estimation of Exchange Rate Models

We begin our statistical and economic evaluation of short-horizon exchange rate predictability by performing Bayesian estimation of the parameters of our 15 candidate models: the five conditional mean specifications (RW, MF_1, MF_2, MF_3, FP) under the three volatility frameworks (LR, $GARCH$, SV). The posterior mean estimates for the parameters of each empirical model are presented in Tables 3.2, 3.3 and 3.4. We particularly focus on the size, sign and statistical significance of the β estimate because it captures the effect of either monetary fundamentals or the forward premium in the conditional mean of exchange rate returns. In our Bayesian MCMC framework we assess statistical significance using two diagnostics. First, we report the highest posterior density (HPD) region for each parameter estimate. For example, the 95% HPD region is the shortest interval that contains 95% of the posterior distribution. We check whether the 90%, 95% and 99% HPD regions contain zero, which is equivalent to two-sided hypothesis testing at the 10%, 5% and 1% level respectively. Second, we compute the Numerical Standard Error (NSE) as defined in Section 3.6.1.

Tables 3.2 through 3.4 illustrate that for the three monetary fundamentals specifications (MF_1, MF_2, and MF_3) the in-sample β estimate tends to be a low positive number, which increases in size as we move from MF_1 to MF_3. This suggests that when s_t is below its fundamental value z_t, it is expected to slowly rise over time. In contrast, the in-sample β estimate for the FP model has a large negative value. For example, in the case of the pound sterling, β rises from 0.0028 for MF_1^{SV}, to 0.0211 for MF_2^{SV}, and then to 0.0226 for MF_3^{SV}, whereas for the FP^{SV} model $\beta = -0.653$.

The tables also report the estimates of the conditional variance parameters. For the Linear Regression model, the monthly variance of FX returns remains largely unchanged across the five conditional mean specifications and is around 10 (i.e.

$\upsilon \approx 3\%$) for all three currencies. For the $GARCH(1,1)$ models, the conditional monthly variance is highly persistent since the sum $\gamma_1 + \gamma_2$ revolves around 0.96 for all specifications. The SV models exhibit (i) high persistence (ϕ) in the conditional monthly log-variance, ranging from $\phi = 0.75$ for the Deutsch mark/euro, $\phi = 0.82$ for the yen, to $\phi = 0.89$ for the pound sterling, and (ii) a sizeable stochastic component in the conditional monthly log-variance, which ranges from $\sigma^2 = 0.070$ for the Deutsch mark/euro, $\sigma^2 = 0.090$ for the pound sterling, to $\sigma^2 = 0.150$ for the yen. Finally, all parameters in both the conditional mean and volatility exhibit very low NSE values and therefore a high degree of statistical significance.

3.7.2 Evaluating Forecasts Using Statistical Criteria

We assess the statistical evidence on short-horizon exchange rate predictability by ranking our set of 15 candidate models according to their log-likelihood and posterior probability. The conditional performance of the models is evaluated in-sample as well as out-of-sample. The in-sample period for the three monthly exchange rates covers 29 years ranging from January 1976 to December 2004. The out-of-sample exercise involves two steps: (i) initial parameter estimation for the 14-year period of January 1976 to December 1989, and (ii) sequential monthly updating of the parameter estimates for the out-of-sample 15-year period of January 1990 to December 2004. In other words, the forecasts at any given month are constructed according to a recursive procedure that is conditional only upon information up to the date of the forecast. The model is then successively re-estimated as the date on which forecasts are conditioned moves through the data set. Hence the design of the out-of-sample exercise is computationally intensive.

Our analysis of the statistical evidence begins with Table 3.5, which presents the log-likelihood values and demonstrates that across volatility models, the SV model always has higher log-likelihood than both LR and $GARCH$. This result is very robust as it holds for all three currencies both in-sample and out-of-sample. Similarly, the $GARCH(1,1)$ models always beat the constant variance LR models in terms of log-likelihood. Furthermore, across conditional mean specifications, the RW model is always worse in-sample than the model specifications which condition on either monetary fundamentals or the forward premium. Specifically, in-sample the MF models are best for the pound sterling and the Deutsch mark/euro, whereas the FP model is best for the yen. Finally, the out-of-sample log-likelihood values lead to the following conclusions: FP is still the best model for the yen, but now the RW model is best for the pound sterling and the Deutsch mark/euro.

In Table 3.6 we rank the in-sample and out-of-sample performance of our set of candidate models according to their posterior probability. The key input to this statistical criterion is the calculation of the marginal likelihood. Therefore, Table 3.6 gives us a distinct statistical perspective on performance because the marginal likelihood is computed in a way that integrates out parameter uncertainty and imposes a penalty for lack of parsimony (higher dimension). The results in Table 3.6 indicate two clear patterns in ranking the models. The first pattern confirms one of our most robust results: the best models for all three currencies both in-sample and out-of-sample have SV innovations. The second pattern provides a result that is slightly different from the log-likelihood findings: for all three exchange rates, both in-sample and out-of-sample, the best model is FP^{SV}, the second best is RW^{SV}, and third best is one of the three MF^{SV} specifications. The single exception is the pound sterling for which RW^{SV} is the best out-of-sample model. Hence, in contrast to the likelihood evidence, the MF specifications lose to RW even in-sample. In other words, the penalty the posterior probability imposes on the three monetary fundamentals models for lack of parsimony offsets their log-likelihood advantage.

3.7.3 Evaluating Forecasts Using Economic Criteria

We assess the economic value of short-horizon exchange rate predictability by analyzing the performance of the dynamically rebalanced portfolios constructed using our set of 15 candidate models. Our analysis focuses on the performance fee, Φ, a US investor is willing to pay for switching from one FX strategy to another. The fees are reported in Table 3.7, which displays the economic value of each mean and volatility specification relative to the benchmark random walk model with constant variance (RW^{LR}). We present the fees for the degrees of RRA $\delta = 2$ and $\delta = 6$.

Panel A of Table 3.7 presents the in-sample performance fees and demonstrates that the three monetary fundamentals specifications generally have no economic value as indicated by the negative Φ values. Only under stochastic volatility does the canonical MF_1 model beat the random walk benchmark with constant variance. On the other hand, the forward premium model (FP) exhibits high economic value, especially under stochastic volatility. For example, at the target portfolio volatility of $\sigma_p^* = 10\%$ and for $\delta = 2$, a US investor is willing to pay a substantial 248 annual basis points (bps) for switching from the RW^{LR} model to FP^{SV}. Consistent with our statistical evidence, for all conditional mean specifications there tends to be high economic value associated with stochastic volatility. However, in contrast to our statistical evidence, the performance of the $GARCH(1,1)$ model is surprisingly poor relative to the constant variance Linear Regression model. Specifically, at $\sigma_p^* = 10\%$ and $\delta = 2$, the in-sample fee for switching from RW^{LR} to RW^{GARCH} is -24 *bps*, whereas the fee for switching from RW^{LR} to RW^{SV} is 42 *bps*.[44] Finally, as investors become less risk averse, the fees tend to increase in absolute value, strengthening the evidence against the random walk benchmark and in favour of the FP^{SV} specification.

The out-of-sample performance fees are displayed in Panel B of Table 3.7 and suggest that even out-of-sample there is still high economic value in both the forward premium and stochastic volatility. This is a new and important result, which adds to the existing literature that is anchored around the seminal contribution of Meese and Rogoff (1983). Specifically, at $\sigma_p^* = 10\%$ and $\delta = 2$, the annual performance fees for switching from RW^{LR} to another model are: 127 *bps* for RW^{SV} and 266 *bps* for FP^{SV}. We can therefore conclude that there is substantial economic value both in-sample and out-of-sample against the naive random walk model and in favour of conditioning on the forward premium with stochastic volatility. This finding is in fact consistent with the large profits made by financial institutions that engage in sophisticated multi-currency forward bias strategies. For example, Galati and Melvin (2004) show that simple carry trades aiming at exploiting the forward bias constitute a significant source of the surge in FX trading observed in recent years.

In addition to the results associated with individual models, even stronger eco-

[44] At first sight, the poor performance of the GARCH model in terms of economic value appears rather surprising. For instance, Fleming and Kirby (2003) find that SV models only marginally outperform GARCH models. However, there is no study to date which assesses the economic value of GARCH and SV models, especially when applied to exchange rates. Furthermore, the negative in-sample and out-of-sample performance fees of RW^{GARCH} are not far from zero.

nomic evidence is found for the combined forecasts reported in Table 3.8. In particular, we compare the three methods of forecast combination described in Section 3.6.3 to the RW^{LR} benchmark for two cases: (i) the restricted universe of the five SV models (because the SV models generally perform the best), and (ii) the unrestricted universe of all 15 models. A purely agnostic approach to forecast combination would use the full set of 15 models (case ii). At first glance, therefore, restricting the universe of models appears to be conceptually inconsistent with the adopted model uncertainty framework. However, we still consider the restricted universe of SV models (case i) as an additional exercise, which represents the case where an investor first examines the in-sample results at the end of 1989, realises the superior in-sample performance of SV models, and consequently decides to use the restricted universe of only SV models in the out-of-sample forecast combination.

The results in Table 3.8 provide robust evidence against the naive random walk model as all performance fees based on combined forecasts are positive and high, both in-sample and out-of-sample. The BMA and BW perform similarly well and, in turn, both perform better than DMA. For example, when selecting among the SV models and setting $\sigma_p^* = 10\%$ and $\delta = 2$, the annual in-sample performance fee for switching away from the benchmark random walk model (RW^{LR}) is 169 *bps* for DMA, 255 *bps* for BMA, and 235 *bps* for BW. The out-of-sample fees are even higher: 219 *bps* for DMA, 317 *bps* for BMA, and 340 *bps* for BW. In short, therefore, there is clear in-sample and out-of-sample economic evidence on the superiority of combined forecasts relative to the naive random walk benchmark. In fact, the simple DMA $(1/N)$ strategy comfortably beats the RW^{LR} model and indeed its performance is not drastically lower than the more sophisticated BMA and BW strategies.

In conclusion, Figure 3.1 offers a visual description of the time variation in the weights investing in the three risky assets: the UK, German and Japanese bonds. The figure displays the weights for four cases: the benchmark RW^{LR} model, the best performing individual model FP^{SV}, the DMA combined forecast strategy, and finally the more sophisticated BMA strategy. As expected, the weights are very smooth over time for RW^{LR} and DMA, and remain reasonably smooth for the FP^{SV} model and the BMA strategy.[45]

3.7.4 Transaction Costs

If transaction costs are sufficiently high, the period-by-period fluctuations in the dynamic weights of an optimal strategy will render the strategy too costly to implement relative to the static random walk model. We address this concern by computing the break-even transaction cost, τ^{BE}, as the minimum monthly proportional cost which cancels out the utility advantage (and hence positive performance fee) of a given strategy. In comparing a dynamic strategy with the static random walk strategy, an investor who pays a transaction cost lower than τ^{BE} will prefer the dynamic strategy. The τ^{BE} values are expressed in monthly basis points and are reported only when Φ is positive.

The in-sample break-even transaction costs are reported in Panel A of Table

[45]However, the dynamic weights appear to be quite volatile in the beginning of the sample before they stabilize. We have experimented with alternative initial values and starting dates, and found that the results are robust to different initializations. Therefore, we believe that the initial instability in the weights is due to the high exchange rate volatility at the start of the sample, especially during the 1992 crisis of the Exchange Rate Mechanism that forced the UK to abandon the target zone system following a speculative attack.

3.7, which demonstrates that for the forward premium and stochastic volatility the values of τ^{BE} are positive and high; they tend to be higher than 100 bps and can be as high as 556 bps. For instance, at $\sigma_p^* = 10\%$ and $\delta = 2$, a US investor will switch back to the RW^{LR} model if he is subject to a proportional transaction cost of at least 120 bps for FP^{LR}, 101 bps for FP^{GARCH}, 132 bps for FP^{SV}, and 471 bps for RW^{SV}. In other words, at the reasonably high transaction cost of 50 bps (e.g. Marquering and Verbeek, 2004), there is still significant in-sample economic value in empirical models which condition on the forward premium, especially under stochastic volatility.

Determining the out-of-sample robustness to transaction costs is one of the most important considerations in assessing the forecasting performance of empirical exchange rates models. Panel B of Table 3.7 shows that conditioning on the forward premium and stochastic volatility leads to reasonably high τ^{BE} values. Specifically, at $\sigma_p^* = 10\%$ and $\delta = 2$, the break-even transaction cost which would eliminate the performance fee of 266 bps of the FP^{SV} model relative to the RW^{LR} benchmark is 90 bps. Furthermore, the τ^{BE} for RW^{SV} versus RW^{LR} is a very large 321 bps.

The evidence on the τ^{BE} of combined forecasts displayed in Table 3.8 is even stronger. Compared to the benchmark RW^{LR} at $\sigma_p^* = 10\%$ and $\delta = 2$, a combined forecast of all 15 model specifications exhibits an in-sample τ^{BE} of: 240 bps for DMA, 141 bps for BMA, and 114 bps for BW. Additionally, Panel B of Table 3.8 shows that the out-of-sample τ^{BE} values for combined forecasts are generally as high as the in-sample values. It is particularly interesting to note that for the simple DMA $(1/N)$ strategy we find a positive Φ over the RW^{LR} benchmark and a high τ^{BE}. In short, as the τ^{BE} values are generally positive and reasonably high, we conclude that the in-sample and out-of-sample economic value we have reported is robust to reasonably high transaction costs for empirical exchange rate models conditioning on the forward premium, for models with SV innovations, and for combined forecasts.

3.7.5 Summary of Results

The statistical and economic evidence on short-horizon exchange rate predictability supports the following four results: (i) the forward premium model unequivocally beats the random walk both in-sample and out-of-sample; (ii) conditioning on monetary fundamentals has no economic value either in-sample or out-of-sample; (iii) the stochastic volatility process always leads to superior portfolio performance both in-sample and out-of-sample; and (iv) the combined forecasts, including the simple $1/N$ strategy, consistently outperform the constant variance random walk benchmark both in-sample and out-of-sample. All these results are robust to reasonably high transaction costs.

3.7.6 Robustness and Extensions

This section discusses directions in which one can possibly extend the analysis of the chapter. First, we perform an additional robustness test by evaluating the out-of-sample performance of the empirical models in three 5-year subsamples. Recall that the full sample period at our disposal covers 29 years ranging from January 1976 to December 2004. We use data from January 1976 to December 1989 for in-sample estimation, whereas the out-of-sample period contains 15 years ranging from January 1990 to December 2004. The out-of-sample results we report in Tables 3.5 through 3.8 are for the entire 15-year out-of-sample period. In addition, Panel A of Table 3.9 presents the performance fees for selected models for three subsamples: 1990-1994, 1995-1999 and 2000-2004. We find that the economic value in conditioning on the forward premium and stochastic volatility is positive in all periods but is substantially higher in the last two subsamples. This is consistent with the well-known fact in the literature that the forward bias is very small in the early 1990s (e.g. Flood and Rose, 2002).[46] For all models, the best subsample period is 1995-1999. Furthermore, it is important to note that the combined forecast strategies, including the simple DMA, substantially outperform the random walk benchmark in all three subsamples. Finally, the best performing combined forecast strategies, BMA and BW, display similar performance fees to FP^{SV} for the last two subsamples. However, for the first subsample when the forward bias is small, the BMA and BW strategies significantly outperform FP^{SV} by optimally using predictive information from the entire universe of models, including monetary fundamentals.

Second, our analysis of the conditional variance of exchange rate returns includes the $GARCH(1,1)$ specification because this is the benchmark model in the seminal study of West, Edison and Cho (1993). As a further robustness check, we also examine the out-of-sample performance of the $tGARCH(1,1)$ model of Bollerslev (1987) in order to determine whether departing from the assumption of conditional

[46]In a separate experiment we start the out-of-sample exercise in 1985 and find significant economic value in the forward premium and stochastic volatility for the 1985-1989 period. However, starting the out-of-sample period in 1985 leaves too few in-sample observations for initial parameter estimation. Therefore, the tables present the out-of-sample results for the period starting in 1990.

normality can improve the performance of the GARCH model.[47] The results from this exercise – a subset of which is reported in Panel B of Table 3.9 – reveal that using a Student-t distribution leads to substantial performance gains in the GARCH framework.[48] In particular, the results in Panel B of Table 3.9 show that the out-of-sample performance fees of the $tGARCH(1,1)$ model are much higher than for $GARCH(1,1)$, especially for the forward premium and random walk conditional mean specifications. For instance, setting $\sigma_p^* = 10\%$ and $\delta = 2$ and comparing the results in Tables 3.7 and 3.9 indicates that the out-of-sample fees for switching from the RW^{LR} model to the forward premium models are as follows: 76 *bps* for FP^{LR}, 70 *bps* for FP^{GARCH}, 140 *bps* for FP^{tGARCH}, and 266 *bps* for FP^{SV}. Similarly, when switching from the random walk with constant variance, RW^{LR}, to a random walk with time-varying volatility the fees are: -32 *bps* for RW^{GARCH}, 28 *bps* for RW^{tGARCH}, and 127 *bps* for RW^{SV}. Therefore, we can conclude that in terms of economic value the $tGARCH$ model performs better than $GARCH$, although the SV model outperforms both normal and Student-t GARCH specifications. Hence, our main conclusions remain qualitatively the same.

Third, Table 3.10 presents the in-sample and out-of-sample annualised Sharpe ratios for selected models. The Sharpe ratio values are generally in agreement with the performance fees and hence confirm our conclusions. Specifically, FP^{SV} and all combined forecast strategies consistently outperform the random walk model both

[47] In estimating the $tGARCH$ model, we implement an algorithm similar to the $GARCH$ case as described in Appendix B.3, with an additional Metropolis-Hastings step for sampling the degrees of freedom parameter ν.

[48] Note that the degrees of freedom parameter estimate revolves around $\nu = 10$ for the UK pound, $\nu = 25$ for the Deutsch mark/euro, and $\nu = 7$ for the Japanese yen (not reported). This indicates that the unconditional distribution of exchange rate returns is not normal, especially for the UK pound and the Japanese yen.

in-sample and out-of-sample. Indeed, the simple DMA strategy also performs better than the random walk, but not as well as FP^{SV}. The best performing strategies are BMA and BW. For example, at $\sigma_p^* = 10\%$, the out-of-sample Sharpe ratios are: 0.76 for RW^{LR}, 0.98 for FP^{SV}, 0.86 for DMA, 1.06 for BMA, and 1.12 for BW.

Fourth, this chapter explores the predictability in exchange rates by focusing on the frequency and horizon of one month. On the one hand, adopting the monthly frequency is a natural choice because this is the highest frequency at which monetary fundamentals are observed. On the other hand, our motivation for investigating predictability at the one-month horizon is founded on the prevailing view in this literature that exchange rates are not predictable at short horizons. It is clear, therefore, that one possible direction in extending the analysis of this chapter is to study the predictability of the forward premium, stochastic volatility and combined forecasts for higher frequencies and longer horizons. We leave this for future research.

Finally, we study short-horizon exchange rate predictability by estimating a set of univariate conditional mean and volatility models. However, in assessing the economic value of exchange rate predictability we build multivariate dynamic asset allocation strategies. Specifically, the optimal weights of the dynamically rebalanced portfolios are computed using the conditional mean forecasts, the conditional volatility forecasts and the dynamic covariances implied by the constant conditional correlation (CCC) model of Bollerslev (1990). In the CCC model, the dynamics of covariances are driven by the time-variation in the conditional volatilities. By design, therefore, the advantage of this setting is that the optimal weights will vary across models only to the extent that forecasts of the conditional mean and volatility will vary, which is precisely what the empirical models provide. Indeed, introducing multivariate stochastic volatility models for capturing the dynamic heteroskedasticity of the covariances of exchange rate returns remains an important extension to this line of research. Multivariate stochastic volatility models are high dimensional and their estimation is computationally challenging (e.g. Chib, Nardari and Shephard, 2006). Additionally the dynamic conditional correlation (DCC) model of Engle (2002) has yet to be examined in a Bayesian SV framework. Hence, we will revisit this issue in future research.

3.8 Conclusion

This chapter draws from three separate, yet related strands of international finance literature. A large body of empirical research finds that models which condition on monetary fundamentals cannot outperform the naive random walk model in out-of-sample forecasting of exchange rates. Despite the increasing sophistication of the econometric techniques implemented and the improving quality of the data sets utilised, evidence of exchange rate predictability remains elusive. A second and related research strand indicates that the rejection of the risk-neutral FX efficient market hypothesis implies that exchange rate movements can be predicted using information contained in forward premia. Finally, financial economists agree that exchange rate volatility is predictable by specifying either GARCH or stochastic volatility innovations.

Prior research in this area has largely relied on standard statistical measures of forecast accuracy. In this chapter, we complement this approach in two critical aspects. First, in assessing the predictive performance of the set of empirical exchange rate models, we implement a Bayesian methodology which explicitly accounts for parameter and model uncertainty. Second, we provide a comprehensive economic evaluation of the models in the context of dynamic asset allocation strategies. In doing so, our study contributes to the growing empirical literature on exchange rate predictability in the following manner. We assess the economic value of exchange rate forecasts derived from empirical models which condition on information contained in either monetary fundamentals or forward premia. This is done in a framework that allows for time-varying volatility. The empirical exchange rate models are set against the naive random walk benchmark. Finally, we evaluate the performance of combined forecasts based on Deterministic and Bayesian Model Averaging.

Our results provide robust evidence against the random walk (no predictability) benchmark, and therefore our empirical findings reinforce the notion that exchange rates are predictable. Specifically, we find that the predictive ability of the forward premium has substantial economic value in a dynamic portfolio allocation context and that stochastic volatility significantly outperforms the constant variance and GARCH(1,1) models irrespective of the conditional mean specification. Combined forecasts which are formed using Deterministic and Bayesian Model Averaging also substantially outperform the random walk benchmark. These results are robust to reasonably high transaction costs and they hold for all currencies both in-sample and out-of-sample. In short, these findings suggest that the random walk hypothesis as applied to exchange rates might have been overstated, while at the same time they justify the widespread use of forward bias and volatility timing strategies in the practice of currency management.

Table 3.1

Descriptive Statistics for Monthly FX Returns and Fundamentals

	UK (USD/GBP)					Germany (USD/DEM − EURO)					Japan (USD/JPY)				
	Δs_t	MF_1	MF_2	MF_3	FP	Δs_t	MF_1	MF_2	MF_3	FP	Δs_t	MF_1	MF_2	MF_3	FP
Panel A: Percent Returns															
Mean	−0.012	0.000	0.000	0.000	−0.159	0.165	0.000	0.000	0.000	0.131	0.309	0.000	0.000	0.000	0.267
Std Dev	2.95	21.52	13.50	11.57	0.266	3.05	33.43	17.56	16.82	0.300	3.32	42.00	29.61	16.88	0.300
Min	−10.53	−50.02	−37.41	−24.56	−0.736	−9.71	−56.90	−34.87	−29.02	−0.779	−10.08	−72.24	−45.68	−43.25	−2.12
Max	11.69	52.89	41.86	44.00	1.35	9.21	75.85	45.75	46.78	1.59	11.89	72.29	70.28	30.85	1.31
Skewness	−0.101	−0.137	−0.372	0.474	1.49	0.156	0.164	0.240	0.473	−0.162	0.619	0.234	0.575	−0.247	−1.05
Kurtosis	4.04	2.41	3.60	4.24	9.19	3.12	2.05	2.17	2.39	4.46	4.06	1.63	2.20	2.00	13.94
$Corr\,(\Delta s_t, \Delta s_{t-1})$	0.106	0.988	0.976	0.965	0.786	0.115	0.994	0.983	0.983	0.641	0.147	0.996	0.987	0.980	0.558
$Corr\,(\Delta s_t, \Delta s_{t-3})$	0.013	0.960	0.920	0.888	0.480	0.040	0.982	0.945	0.943	0.600	0.108	0.989	0.969	0.933	0.510
$Corr\,(\Delta s_t, \Delta s_{t-6})$	0.037	0.911	0.838	0.780	0.208	0.025	0.959	0.886	0.880	0.593	−0.057	0.973	0.938	0.856	0.403
$Corr\,(\Delta s_t, \Delta s_{t-12})$	0.025	0.793	0.674	0.569	0.334	0.030	0.897	0.734	0.719	0.538	0.078	0.935	0.882	0.704	0.267
Panel B: Absolute Percent Returns															
Mean	2.25	17.92	10.46	8.57	0.237	2.43	28.14	15.07	14.18	0.271	2.48	37.17	25.65	14.60	0.315
Std Dev	1.90	11.87	8.51	7.76	0.200	1.84	17.99	8.98	9.01	0.184	2.22	19.46	14.71	8.43	0.248
Min	0.000	0.023	0.041	0.050	0.000	0.006	0.320	0.007	0.039	0.000	0.000	0.028	0.477	0.009	0.000
Max	11.69	52.89	41.86	44.00	1.35	9.71	75.85	45.75	46.78	1.59	11.89	72.29	70.28	43.25	2.12
Skewness	1.48	0.687	1.37	1.46	1.55	1.12	0.238	0.547	0.711	1.64	1.52	−0.281	0.729	0.312	1.95
Kurtosis	6.04	2.89	4.64	5.67	6.93	4.38	2.21	3.11	3.40	10.40	5.76	2.11	3.46	2.82	11.33
$Corr\,(\|\Delta s_t\|, \|\Delta s_{t-1}\|)$	0.065	0.960	0.944	0.933	0.878	0.135	0.981	0.939	0.944	0.647	0.129	0.983	0.950	0.927	0.746
$Corr\,(\|\Delta s_t\|, \|\Delta s_{t-3}\|)$	0.149	0.876	0.832	0.797	0.731	0.085	0.943	0.807	0.818	0.466	0.051	0.950	0.884	0.761	0.590
$Corr\,(\|\Delta s_t\|, \|\Delta s_{t-6}\|)$	0.093	0.724	0.690	0.606	0.577	0.051	0.876	0.633	0.641	0.410	0.010	0.882	0.786	0.541	0.500
$Corr\,(\|\Delta s_t\|, \|\Delta s_{t-12}\|)$	0.030	0.431	0.414	0.248	0.385	−0.046	0.685	0.305	0.271	0.212	−0.035	0.703	0.659	0.280	0.330

The table summarizes the descriptive statistics for the spot exchange rate percent returns (Δs_t), the three demeaned percent monetary fundamentals specifications (MF_1, MF_2, MF_3), and the percent forward premium (FP). The data sample ranges from January 1976 through December 2004 for a sample size of 348 monthly observations. The exchange rates are defined as US dollars per unit of foreign currency. For a detailed definition of the three monetary fundamentals specifications see Section 3.1.

Table 3.2

Posterior Means for the UK Pound Sterling (*USD/GBP*)

Parameter	*RW*	MF_1	MF_2	MF_3	*FP*
Panel A: Bayesian Linear Regression					
α	−0.012 (0.0005)	−0.012 (0.0005)	−0.012 (0.0005)	−0.012 (0.0005)	−0.110 (0.0005)
β	–	0.0079 (2.4e−05)	0.0254** (3.7e−05)	0.0254* (4.4e−05)	−0.629 (0.0016)
v^2	8.72*** (0.0021)	10.12*** (0.0021)	8.63*** (0.0021)	10.06*** (0.0021)	8.69*** (0.0021)
Panel B: Bayesian GARCH(1,1)					
α	0.018 (0.0022)	0.027 (0.0021)	0.005 (0.0021)	0.017 (0.0020)	−0.101 (0.0022)
β	–	0.0042 (0.0001)	0.0215 (0.0002)	0.0193 (0.0002)	−0.816 (0.0078)
ω	0.331*** (0.0015)	0.346*** (0.0026)	0.324*** (0.0017)	0.329*** (0.0017)	0.387*** (0.0064)
γ_1	0.905*** (0.0003)	0.902*** (0.0004)	0.903*** (0.0003)	0.902*** (0.0003)	0.897*** (0.0009)
γ_2	0.055*** (0.0002)	0.056*** (0.0002)	0.0572*** (0.0002)	0.058*** (0.0002)	0.056*** (0.0004)
Panel C: Bayesian Stochastic Volatility					
α	0.048 (0.0026)	0.046 (0.0028)	0.022 (0.0032)	0.022 (0.0027)	−0.045 (0.0034)
β	–	0.0028 (0.0001)	0.0211 (0.0002)	0.0226 (0.0004)	−0.653 (0.0107)
μ	2.01*** (0.0037)	2.02*** (0.0033)	2.01*** (0.0035)	2.01*** (0.0036)	2.00*** (0.0033)
ϕ	0.882*** (0.0012)	0.878*** (0.0014)	0.885*** (0.0014)	0.884*** (0.0013)	0.871*** (0.0015)
σ^2	0.093*** (0.0010)	0.092*** (0.0010)	0.086*** (0.0009)	0.090*** (0.0010)	0.097*** (0.0011)

The table presents the Bayesian MCMC estimates of the posterior means of the Linear Regression, GARCH(1,1) and SV model parameters for the *USD/GBP* monthly percent FX returns. The MCMC chain run for 5,000 iterations after an initial burn-in of 1,000 iterations. The numbers in parenthesis indicate the Numerical Standard Error (NSE). The superscripts *, ** and *** indicate that the 90%, 95% and 99% highest posterior density (HPD) regions, respectively, do not contain zero. The HPD region for each MCMC parameter estimate is the shortest interval that contains 95% of the posterior distribution.

Table 3.3

Posterior Means for the Deutsch Mark/Euro (*USD/DEM-EURO*)

Parameter	*RW*	MF_1	MF_2	MF_3	*FP*
Panel A: Bayesian Linear Regression					
α	0.160 (0.0005)	0.160 (0.0005)	0.160 (0.0005)	0.160 (0.0005)	0.206 (0.0005)
β	–	0.0077 (1.6e−0.5)	0.0104 (3.0e−05)	0.0148 (3.1e−05)	−0.355 (0.0015)
v^2	9.30*** (0.0023)	9.26*** (0.0022)	9.30*** (0.0022)	9.27*** (0.0022)	10.81*** (0.0022)
Panel B: Bayesian GARCH(1,1)					
α	0.153 (0.0024)	0.164 (0.0022)	0.154 (0.0023)	0.159 (0.0023)	0.216 (0.0023)
β	–	0.0068 (7.6e−05)	0.0097 (0.0001)	0.0134 (0.0001)	−0.463 (0.0072)
ω	0.405*** (0.0019)	0.404*** (0.0026)	0.400*** (0.0028)	0.399*** (0.0027)	0.409*** (0.0023)
γ_1	0.930*** (0.0003)	0.929*** (0.0003)	0.928*** (0.0003)	0.928*** (0.0003)	0.929*** (0.0003)
γ_2	0.028*** (0.0002)	0.029*** (0.0002)	0.030*** (0.0002)	0.030*** (0.002)	0.028*** (0.0002)
Panel C: Bayesian Stochastic Volatility					
α	0.163 (0.0031)	0.176 (0.0010)	0.165 (0.0028)	0.173 (0.0028)	0.219 (0.0030)
β	–	0.0074 (9.4e−05)	0.0091 (0.0002)	0.0136 (0.0002)	−0.440 (0.0084)
μ	2.17*** (0.0030)	2.16*** (0.0028)	2.17*** (0.0028)	2.16*** (0.0030)	2.17*** (0.0029)
ϕ	0.746*** (0.0032)	0.751*** (0.0036)	0.757*** (0.0033)	0.755*** (0.0036)	0.751*** (0.0033)
σ^2	0.068*** (0.0005)	0.069*** (0.0006)	0.068*** (0.0005)	0.067*** (0.0005)	0.069*** (0.0005)

The table presents the Bayesian MCMC estimates of the posterior means of the Linear Regression, GARCH(1,1) and SV model parameters for the *USD/DEM-EURO* monthly percent FX returns. The MCMC chain run for 5,000 iterations after an initial burn-in of 1,000 iterations. The numbers in parenthesis indicate the Numerical Standard Error (NSE). The superscripts *, ** and *** indicate that the 90%, 95% and 99% highest posterior density (HPD) regions, respectively, do not contain zero. The HPD region for each MCMC parameter estimate is the shortest interval that contains 95% of the posterior distribution.

Table 3.4

Posterior Means for the Japanese Yen (*USD/JPY*)

Parameter	RW	MF_1	MF_2	MF_3	FP
Panel A: Bayesian Linear Regression					
α	0.299^{*} (0.0006)	0.299^{*} (0.0005)	0.299^{*} (0.0005)	0.299^{*} (0.0005)	0.615^{***} (0.0007)
β	–	0.0070 (1.3e−0.5)	0.0075 (1.9e−05)	0.0189^{*} (3.4e−05)	-1.224^{**} (0.0016)
v^2	11.0^{***} (0.0027)	10.96^{***} (0.0026)	10.99^{***} (0.0026)	10.94^{***} (0.0026)	10.79^{***} (0.0026)
Panel B: Bayesian GARCH(1,1)					
α	0.343^{*} (0.0027)	0.344^{*} (0.0023)	0.342^{*} (0.0024)	0.333^{*} (0.0024)	0.623^{***} (0.0027)
β	–	0.0065 (6.4e−05)	0.0077 (8.6e−05)	0.0164 (0.0002)	-1.170^{**} (0.0069)
ω	0.595^{***} (0.0028)	0.593^{***} (0.0034)	0.599^{***} (0.0043)	0.594^{***} (0.0042)	0.676^{***} (0.0076)
γ_1	0.911^{***} (0.0004)	0.911^{***} (0.0004)	0.909^{***} (0.0005)	0.911^{***} (0.0005)	0.898^{***} (0.0011)
γ_2	0.037^{***} (0.0003)	0.036^{***} (0.0003)	0.038^{***} (0.0003)	0.036^{***} (0.0003)	0.041^{***} (0.0007)
Panel C: Bayesian Stochastic Volatility					
α	0.166 (0.0039)	0.166 (0.0039)	0.161 (0.0041)	0.138 (0.0040)	0.532^{***} (0.0038)
β	–	0.0055 (8.1e−05)	0.0059 (0.0001)	0.0155 (0.0002)	-1.763^{***} (0.0109)
μ	2.16^{***} (0.0040)	2.16^{***} (0.0042)	2.16^{***} (0.0037)	2.15^{***} (0.0036)	2.06^{***} (0.0040)
ϕ	0.814^{***} (0.0017)	0.818^{***} (0.0020)	0.818^{***} (0.0018)	0.816^{***} (0.0019)	0.801^{***} (0.0019)
σ^2	0.149^{***} (0.0025)	0.147^{***} (0.0021)	0.146^{***} (0.0022)	0.150^{***} (0.0020)	0.230^{***} (0.0043)

The table presents the Bayesian MCMC estimates of the posterior means of the Linear Regression, GARCH(1,1) and SV model parameters for the *USD/JPY* monthly percent FX returns. The MCMC chain run for 5,000 iterations after an initial burn-in of 1,000 iterations. The numbers in parenthesis indicate the Numerical Standard Error (NSE). The superscripts *, ** and *** indicate that the 90%, 95% and 99% highest posterior density (HPD) regions, respectively, do not contain zero. The HPD region for each MCMC parameter estimate is the shortest interval that contains 95% of the posterior distribution.

Table 3.5
The Likelihood of the Models

Panel A: In-Sample log-Likelihood					
		USD/GBP			
Model	*RW*	MF_1	MF_2	MF_3	*FP*
LR	−867.47	−866.90	−865.12	−865.76	−866.50
GARCH(1,1)	−860.94	−860.51	−858.00	−858.67	−860.28
SV	−801.80	−804.07	−802.25	−801.15	−801.33
		USD/DEM-EURO			
LR	−878.71	−877.47	−878.09	−877.55	−878.49
GARCH(1,1)	−878.26	−876.96	−877.45	−876.91	−877.97
SV	−847.84	−846.23	−846.52	−846.72	−846.64
		USD/JPY			
LR	−907.95	−906.60	−907.17	−906.34	−904.08
GARCH(1,1)	−906.97	−905.60	−905.93	−905.37	−903.05
SV	−828.49	−827.76	−828.66	−827.61	−791.47

Panel B: Out-of-Sample log-Likelihood					
		USD/GBP			
Model	*RW*	MF_1	MF_2	MF_3	*FP*
LR	−439.52	−440.173	−438.97	−439.25	−439.91
GARCH(1,1)	−427.58	−427.69	−426.27	−426.31	−428.25
SV	−412.08	−412.91	−412.24	−412.63	−412.18
		USD/DEM-EURO			
LR	−451.91	−452.37	−452.50	−453.06	−452.76
GARCH(1,1)	−448.47	−448.68	−449.00	−449.36	−449.03
SV	−427.120	−433.68	−433.91	−434.69	−434.58
		USD/JPY			
LR	−465.04	−465.58	−465.47	−466.99	−464.84
GARCH(1,1)	−458.33	−458.68	−458.68	−460.14	−458.05
SV	−433.814	−426.40	−425.50	−425.68	−413.52

The table reports the in-sample and out-of-sample log-likelihood values for the three FX rates (USD/GBP, USD/DEM-EURO and USD/JPY), five conditional mean specifications (*RW*, MF_1, MF_2, MF_3 and *FP*) and three volatility frameworks (Linear Regression, GARCH and Stochastic Volatility). The out-of-sample data runs from January 1990 through December 2004.

Table 3.6
The Models with the Highest Posterior Probability

Panel A: The Best In-Sample Models			
	Best Model	*Second Best Model*	*Third Best Model*
USD/GBP	FP^{SV}	RW^{SV}	MF_3^{SV}
$USD/DEM-EURO$	FP^{SV}	RW^{SV}	MF_2^{SV}
USD/JPY	FP^{SV}	RW^{SV}	MF_3^{SV}

Panel B: The Best Out-of-Sample Models			
	Best Model	*Second Best Model*	*Third Best Model*
USD/GBP	RW^{SV}	FP^{SV}	MF_2^{SV}
$USD/DEM-EURO$	FP^{SV}	RW^{SV}	MF_2^{SV}
USD/JPY	FP^{SV}	RW^{SV}	MF_3^{SV}

The table shows the three best models according to the highest in-sample and out-of-sample posterior probability for the three FX rates (USD/GBP, USD/DEM-EURO, and USD/JPY). The out-of-sample data runs from January 1990 through December 2004. Ranking the models using the highest posterior probability is equivalent to choosing the best model in terms of density forecasts and is a robust model selection criterion in the presence of misspecification and non-nested models.

Table 3.7

The Economic Value of the Empirical Exchange Rate Models

Panel A: In-Sample Performance for Models *vs.* RW^{LR}

	MF_1^{LR}				MF_2^{LR}				MF_3^{LR}				FP^{LR}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−26	–	−58	–	−129	–	−144	–	−127	–	−167	–	144	120	145	118
10%	−37	–	−90	–	−164	–	−190	–	−165	–	−228	–	180	120	181	117
12%	−51	–	−129	–	−200	–	−239	–	−205	–	−299	–	217	120	218	117

	MF_1^{GARCH}				MF_2^{GARCH}				MF_3^{GARCH}				FP^{GARCH}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	3	14	−13	–	−120	–	−128	–	−128	–	−151	–	132	101	131	98
10%	1	9	−25	–	−152	–	−165	–	−164	–	−202	–	165	101	163	96
12%	−3	–	−41	–	−184	–	−205	–	−201	–	−257	–	198	101	195	96

	MF_1^{SV}				MF_2^{SV}				MF_3^{SV}				FP^{SV}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	102	166	38	59	−55	–	−91	–	−79	–	−144	–	203	136	166	108
10%	118	151	15	18	−74	–	−132	–	−108	–	−213	–	248	132	190	97
12%	130	137	−21	–	−95	–	−180	–	−142	–	−296	–	291	129	206	87

	RW^{GARCH}				RW^{SV}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−19	–	−19	–	36	556	9	51
10%	−24	–	−24	–	42	471	16	150
12%	−28	–	−29	–	48	404	19	267

continued

Table 3.7 *(continued)*

Panel B: Out-of-Sample Performance for Models *vs.* RW^{LR}

	MF_1^{LR}				MF_2^{LR}				MF_3^{LR}				FP^{LR}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−173	–	−178	–	23	26	11	9	−154	–	−149	–	61	34	56	31
10%	−217	–	−224	–	27	23	9	2	−192	–	−184	–	76	34	68	30
12%	−261	–	−271	–	30	21	4	1	−229	–	−218	–	90	34	79	29

	MF_1^{GARCH}				MF_2^{GARCH}				MF_3^{GARCH}				FP^{GARCH}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−124	–	−120	–	−1	–	−2	–	−195	–	−185	–	55	28	63	32
10%	−154	–	−147	–	−1	–	−2	–	−242	–	−227	–	70	29	82	34
12%	−184	–	−172	–	−1	–	−3	–	−289	–	−267	–	86	29	103	36

	MF_1^{SV}				MF_2^{SV}				MF_3^{SV}				FP^{SV}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−18		−68		14	−3	−33		55	40	26	19	218	92	174	73
10%	−30		−108		10	−11	−64		65	37	19	10	266	90	197	65
12%	−45		−158		4	−19	−105		73	34	6	2	311	87	212	58

	RW^{GARCH}				RW^{SV}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−27	–	−17	–	105	342	76	246
10%	−32	–	−16	–	127	321	82	205
12%	−37	–	−13	–	147	304	82	166

The table presents the in-sample and out-of-sample performance fees (Φ) and break-even transaction costs (τ^{BE}) for selected models against the RW^{LR} benchmark for three target portfolio volatilities (8%, 10% and 12%). Each maximum return strategy builds an efficient portfolio by investing in the monthly return of four bonds from the US, UK, Germany and Japan and using the three exchange rates to convert the portfolio return in US dollars. The fees denote the amount an investor with quadratic utility and a degree of relative risk aversion equal to either 2 or 6 is willing to pay for switching from RW^{LR} to another model (such as FP^{SV}). The performance fee Φ is expressed in annual basis points. The transaction cost τ^{BE} is defined as the minimum monthly proportional cost which cancels out the utility advantage (and hence positive performance fee) of a given strategy. The τ^{BE} values are expressed in monthly basis points and are reported only when Φ is positive. The in-sample period starts in January 1979 and the out-of-sample data runs from January 1990 through December 2004.

Table 3.8
The Economic Value of Combined Forecasts

Panel A: In-Sample Performance												
All Models (*vs.* RW^{LR})												
	DMA				BMA				BW			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	109	244	93	204	207	145	156	120	192	117	156	93
10%	134	240	109	189	254	141	172	109	235	114	177	83
12%	158	239	121	177	299	138	178	100	276	111	191	74
Stochastic Volatility Models (*vs.* RW^{LR})												
	DMA				BMA				BW			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	140	262	102	147	208	146	157	120	192	117	156	93
10%	169	253	109	158	255	142	173	110	235	114	178	83
12%	197	244	110	131	300	139	179	100	276	111	191	74
Panel B: Out-of-Sample Performance												
All Models (*vs.* RW^{LR})												
	DMA				BMA				BW			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	83	126	75	116	250	130	206	108	268	128	222	107
10%	103	124	91	111	306	127	237	100	329	125	255	99
12%	122	121	105	106	360	124	259	91	386	122	279	90
Stochastic Volatility Models (*vs.* RW^{LR})												
	DMA				BMA				BW			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	179	174	147	142	259	134	215	112	277	131	231	111
10%	219	169	168	129	317	131	249	104	340	129	267	103
12%	258	164	183	116	373	128	273	95	400	126	294	94

The table reports the in-sample and out-of-sample performance fees (Φ) and break-even transaction costs (τ^{BE}) for all maximum return strategies based on combined forecasts for three target portfolio volatilities (8%, 10% and 12%). *DMA* denotes *Deterministic Model Average* (*1/N* strategy), *BMA* denotes *Bayesian Model Average*, and *BW* is *Bayesian Winner.* The combined forecasts are shown for two cases: (i) the unrestricted universe of all 15 models, and (ii) the restricted universe of only the five stochastic volatility models. The fees denote the amount an investor with quadratic utility and a degree of relative risk aversion equal to either 2 or 6 is willing to pay for switching from the RW^{LR} benchmark to (say) the *BMA* strategy. τ^{BE} is defined as the minimum monthly proportional cost which cancels out the utility advantage (and hence positive performance fee) of a given strategy. The transaction costs are only reported when Φ is positive. The performance fees are expressed in annual basis points, and the transaction costs in monthly basis points. The in-sample period starts in January 1979 and the out-of-sample data runs from January 1990 through December 2004.

Table 3.9

Out-of-Sample Robustness

Panel A: Subsample Analysis for Selected Models *vs.* RW^{LR}

($\sigma_p^* = 10\%$, $\delta = 2$)

	FP^{SV}		DMA		BMA		BW	
Subsample	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}
1990 – 1994	40	12	90	46	196	50	260	56
1995 – 1999	539	347	185	446	519	346	523	357
2000 – 2004	229	83	39	79	227	93	224	96
1995 – 2004	381	193	109	279	363	208	364	281
1990 – 2004	266	90	103	124	306	127	329	125

Panel B: The Performance of $tGARCH$ Models *vs.* RW^{LR}

($\delta = 2$)

	MF_1^{tGARCH}		MF_2^{tGARCH}		MF_3^{tGARCH}		FP^{tGARCH}		RW^{tGARCH}	
σ_p^*	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}
8%	−34	–	−32	–	−29	–	110	49	21	78
10%	−43	–	−40	–	−36	–	140	50	28	82
12%	−50	–	−48	–	−42	–	169	51	35	88

The table provides an analysis of out-of-sample robustness for the performance fees (Φ) and break-even transaction costs (τ^{BE}) of selected models against the RW^{LR} benchmark. Panel A conducts a subsample analysis and Panel B examines the performance of the tGARCH(1,1) model with Student-t innovations. DMA denotes *Deterministic Model Average (1/N* strategy), BMA denotes *Bayesian Model Average*, and BW is *Bayesian Winner*. All maximum return strategies build an efficient portfolio by investing in the monthly return of four bonds from the US, UK, Germany and Japan and using the three exchange rates to convert the portfolio return in US dollars. The fees denote the amount an investor with quadratic utility and a degree of relative risk aversion equal to 2 is willing to pay for switching from RW^{LR} to (say) FP^{SV}. The target portfolio volatility in Panel A is set at 10%. τ^{BE} is defined as the minimum monthly proportional cost which cancels out the utility advantage (and hence positive performance fee) of a given strategy. The transaction costs are only reported when Φ is positive. The performance fees are expressed in annual basis points, and the transaction costs in monthly basis points. The combined forecasts are for the universe of all 15 models. The out-of-sample period runs from January 1990 through December 2004.

Table 3.10

Sharpe Ratios for Selected Models

Panel A: In-Sample					
σ_p^*	RW^{LR}	FP^{SV}	DMA	BMA	BW
8%	0.88	1.09	1.00	1.11	1.13
10%	0.91	1.14	1.05	1.15	1.17
12%	0.94	1.17	1.07	1.19	1.21
Panel B: Out-of-Sample					
σ_p^*	RW^{LR}	FP^{SV}	DMA	BMA	BW
8%	0.76	0.98	0.86	1.06	1.11
10%	0.76	0.98	0.86	1.06	1.12
12%	0.76	0.98	0.86	1.06	1.12

The table presents the in-sample and out-of-sample annualized Sharpe ratios for selected models. *DMA* denotes *Deterministic Model Average (1/N* strategy), *BMA* denotes *Bayesian Model Average*, and *BW* is *Bayesian Winner.* The Sharpe ratios are adjusted for the serial correlation in the monthly portfolio returns generated by the dynamic strategies (e.g. Lo, 2002). All maximum return strategies build an efficient portfolio by investing in the monthly return of four bonds from the US, UK, Germany and Japan and using the three exchange rates to convert the portfolio return in US dollars. The maximum return strategies are evaluated at three target portfolio return volatilities: 8%, 10%, and 12%. The in-sample period starts in January 1979 and the out-of-sample data runs from January 1990 through December 2004.

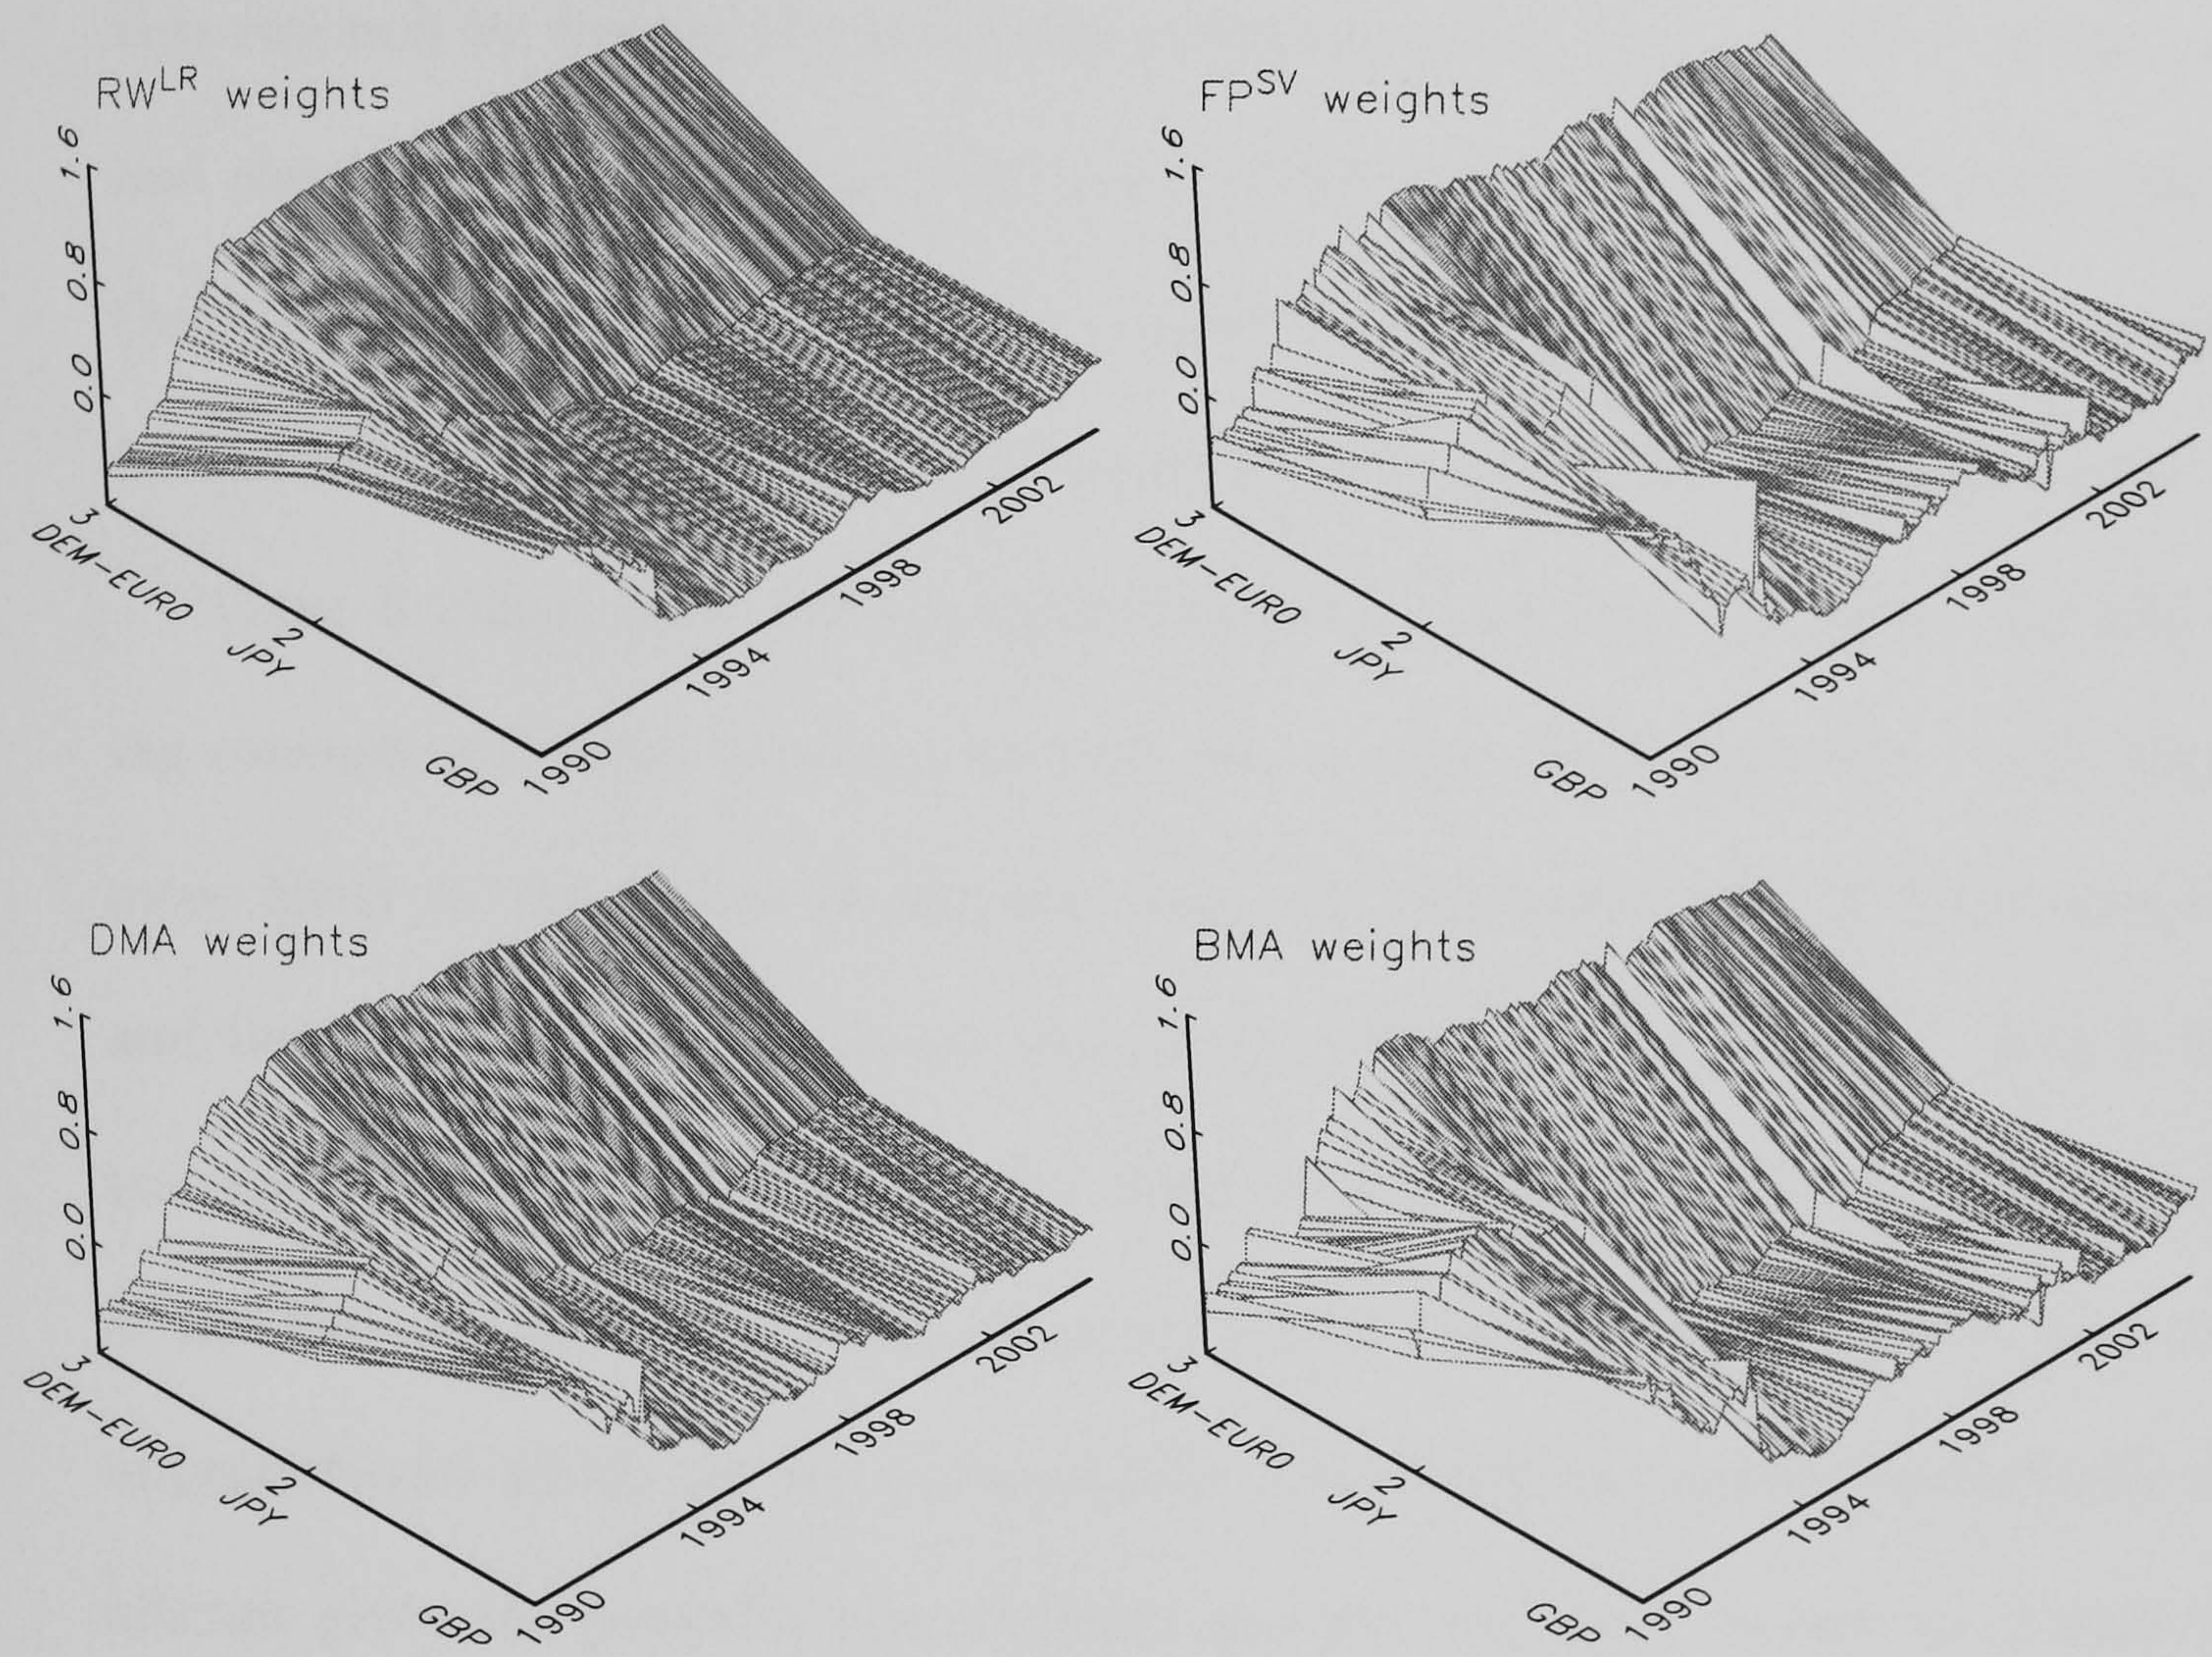

Figure 1: The Out-of-Sample Dynamic Weights for Selected Models: This is the out-of-sample time variation in the weights investing in the three risky assets (the UK, Japanese and German bonds) at a target portfolio volatility of 10% and a degree of relative risk aversion of 2. The figure presents four cases: the benchmark random walk model with constant variance (upper left), the forward premium model with stochastic volatility (upper right), the Deterministic Model Average strategy (lower left), and the Bayesian Model Average strategy (lower right).

4 Concluding Remarks

In the first chapter, we re-examine the result that the EH appears to fit the behaviour of US repo rates at the shortest end of the term structure, measured at daily frequency from overnight to the 3-month maturity (Longstaff, 2000a). We extend this research by testing the restrictions implied by the EH on a VAR of the long- and short-term repo rate using the test proposed by Bekaert and Hodrick (2001). Our empirical investigation, in contrast with Longstaff (2000a), is not encouraging for the EH, which is statistically rejected across the term structure considered.

These findings differ from Longstaff (2000a), who does not reject the EH using conventional tests, because the VAR test is particularly powerful – and, hence, more likely to detect fine departures from the null hypothesis in finite sample – and because our sample is larger than Longstaff's (2000a). However, despite this statistical evidence, a legitimate and unanswered concern is whether the rejection of the EH may be due to small departures from the null hypothesis (or tiny data imperfections) which are not economically meaningful but appear statistically significant given the powerful test statistics and the very large sample size employed. Moreover, the VAR tests are not designed to incorporate the fact that if one wanted to trade on departures from the EH – rather than assuming that the EH holds in a simple buy-and-hold allocation strategy – transactions costs create a wedge between returns from an active strategy exploiting departures from the EH and a simple buy-and-hold strategy. Finally, while the VAR tests rely on the ability of the VAR to capture the time-series properties of the term structure of repo rates, we are aware that the simple VAR tests, inspired by the literature on affine term structure models, is in fact unable to satisfactorily explain conditional means and volatility of interest rates. Hence, potential model misspecification and model uncertainty could play an important role in determining the rejection of the EH recorded in Table 1.8.

In the second chapter, we shed light on the economic significance of the statistical rejections of the EH recorded in previous section, and proceed to an economic evaluation of the EH departures. We innovate in this context by moving beyond statistical tests and providing complementary evidence on the validity of the EH using some economic value calculations. We assess the economic value of exploiting departures from the EH – i.e. using empirical models which condition on information contained in EH deviations – relative to the economic value of using a model that assumes the EH holds. The empirical results indicate that the economic value of departures from the EH is modest and generally smaller than the costs that an investor would incur to exploit the mispricing implied by EH violations. These findings are consistent with the thrust of Longstaff's (2000a) original conclusion.

The results from economic value calculations are in contrast with the results from VAR tests reported earlier. This difference confirms that statistical rejections of a hypothesis do not always imply economic rejections and raises doubts about the ability of the simple linear VAR framework to capture the relationship between repo rates at different maturities. Activities in the repo market at maturities of days or weeks are largely driven by liquidity considerations and by the attempts of banks to manage the quantity of reserves and to hedge interest rate risk on a short-term basis, rather than to speculate in search of excess returns. Hence, it seems unlikely that investors would be actively exploiting EH departures on a very short-term basis. Our main conclusion is that, even though the EH may be rejected statistically, it still provides a very reasonable approximation to the term structure of repo rates and constitutes a useful theory for practitioners in the repo market.

Finally, third chapter draws from three separate, yet related strands of international finance literature. A large body of empirical research finds that models which condition on monetary fundamentals cannot outperform the naive random walk model in out-of-sample forecasting of exchange rates. Despite the increasing sophistication of the econometric techniques implemented and the improving quality of the data sets utilised, evidence of exchange rate predictability remains elusive. A second and related research strand indicates that the rejection of the risk-neutral FX efficient market hypothesis implies that exchange rate movements can be predicted using information contained in forward premia. Finally, financial economists agree that exchange rate volatility is predictable by specifying either GARCH or stochastic volatility innovations.

Prior research in this area has largely relied on standard statistical measures of forecast accuracy. In this chapter, we complement this approach in two critical aspects. First, in assessing the predictive performance of the set of empirical exchange rate models, we implement a Bayesian methodology which explicitly accounts for parameter and model uncertainty. Second, we provide a comprehensive economic evaluation of the models in the context of dynamic asset allocation strategies. In doing so, our study contributes to the growing empirical literature on exchange rate predictability in the following manner. We assess the economic value of exchange rate forecasts derived from empirical models which condition on information contained in either monetary fundamentals or forward premia. This is done in a framework that allows for time-varying volatility. The empirical exchange rate models are set against the naive random walk benchmark. Finally, we evaluate the performance of combined forecasts based on Deterministic and Bayesian Model Averaging.

Our results provide robust evidence against the random walk (no predictability) benchmark, and therefore our empirical findings reinforce the notion that exchange rates are predictable. Specifically, we find that the predictive ability of the forward premium has substantial economic value in a dynamic portfolio allocation context and that stochastic volatility significantly outperforms the constant variance and GARCH(1,1) models irrespective of the conditional mean specification. Combined forecasts which are formed using Deterministic and Bayesian Model Averaging also substantially outperform the random walk benchmark. These results are robust to reasonably high transaction costs and they hold for all currencies both in-sample and out-of-sample. In short, these findings suggest that the random walk hypothesis as applied to exchange rates might have been overstated, while at the same time they justify the widespread use of forward bias and volatility timing strategies in the practice of currency management.

A Appendix: The Expectation Hypothesis

A.1 The EH Restrictions in the VAR Framework

In this appendix we derive the restrictions implied by the EH in the VAR framework. Define the indicator vectors $e_1 = (1, 0, \dots, 0)'$ and $e_2 = (0, 1, 0, \dots, 0)'$ with dimension $2p$ and select from the companion VAR the long-term rate and expected future short-term rates as $i_t^{(n)} = e_1' Y_t$ and $E_t[i_{t+i}^{(m)}] = e_1' \Gamma^i Y_t$, respectively.[49] Hence, the general statement in equation (1.1)

$$i_t^{(n)} = k^{-1} \left\{ i_t^{(m)} + E_t[i_{t+m}^{(m)}] + E_t[i_{t+2m}^{(m)}] + \dots + E_t[i_{t+m(k-1)}^{(m)}] \right\} \tag{A.1.1}$$

can be rewritten, under the maintained assumption that the joint process of the short- and long-term interest rates is accurately described by a linear VAR, as

$$e_2' Y_t = e_1' k^{-1} \left[I + \Gamma^m + \Gamma^{2m} + \dots + \Gamma^{m(k-1)} \right] Y_t \tag{A.1.2}$$

which converges, if the eigenvalues λ_i of Γ are such that $|\lambda_i| < 1$, to the following compact form

$$e_2' Y_t = e_1' k^{-1} \left(I - \Gamma^m \right)^{-1} \left(I - \Gamma^n \right) Y_t. \tag{A.1.3}$$

Notice that right-hand-side of equation (A.1.3) gives the sum of the current and expected short-term rates implied by the predictions of the VAR representation, while the left-hand-side of equation (A.1.3) gives the current long-term rate. In order to satisfy this equality and, hence, makes equation (A.1.3) consistent with equation (A.1.1), equation (A.1.3) implies the following system of nonlinear equations

$$e_2' = e_1' k^{-1} \left(I - \Gamma^m \right)^{-1} \left(I - \Gamma^n \right) \tag{A.1.4}$$

[49] The expectation is with respect to the information set of the VAR.

whose solution implies a $2p$ dimensional vector of highly nonlinear restrictions in the underlying parameters of the VAR. In the case where $m = 1$, the system of equation in (A.1.4) has a simple analytical solution (see Campbell and Shiller, 1987), but in the general case analysed in this chapter and in Bekaert and Hodrick (2001) we have to rely on the numerical outcome of the GMM maximization.

A.2 GMM Iterative Procedure

In this appendix we present the iterative procedure used for the constrained GMM maximization. The first-order conditions for the Lagrangian problem in equation (1.10) can be written as

$$\begin{bmatrix} 0 \\ 0 \end{bmatrix} = \begin{bmatrix} -G_T'\Omega_T^{-1}\sqrt{T}g_T(\overline{\theta}) - A_T'\sqrt{T}\overline{\gamma} \\ -\sqrt{T}a_T(\overline{\theta}) \end{bmatrix} \tag{A.2.1}$$

where $A_T \equiv \nabla_\theta a_T(\theta)$ and $G_T \equiv \nabla_\theta g_T(\theta)$. By using the Taylor's expansion of $g_T(\theta)$ and $a_T(\theta)$ around the true parameter value, θ_0, and substituting into the first-order conditions, Newey and McFadden (1994) derive an approximate asymptotic solution under the null hypothesis $a_T(\theta_0) = 0$ as

$$\begin{bmatrix} 0 \\ 0 \end{bmatrix} = \begin{bmatrix} -G_T'\Omega_T^{-1}\sqrt{T}g_T(\theta_0) \\ 0 \end{bmatrix} - \begin{bmatrix} B_T & A_T' \\ A_T & 0 \end{bmatrix} \begin{bmatrix} \sqrt{T}\left(\overline{\theta} - \theta_0\right) \\ \sqrt{T}\overline{\gamma} \end{bmatrix}. \tag{A.2.2}$$

Next, the formula for a partitioned inverse implies that

$$\begin{bmatrix} B_T & A_T' \\ A_T & 0 \end{bmatrix}^{-1} = \begin{bmatrix} B_T^{-1/2}M_TB_T^{-1/2} & B_T^{-1}A_T'\left(A_TB_T^{-1}A_T'\right)^{-1} \\ \left(A_TB_T^{-1}A_T'\right)^{-1}A_TB_T^{-1} & -\left(A_TB_T^{-1}A_T'\right)^{-1} \end{bmatrix} \tag{A.2.3}$$

where $M_T = I - B_T^{-1/2}A_T'\left(A_TB_T^{-1}A_T'\right)^{-1}A_TB_T^{-1/2}$ is an idempotent matrix, and $B_T \equiv G_T'\Omega_T^{-1}G_T$. Hence, the asymptotic distribution for the constrained estimator and the Lagrange multiplier turns out to be $\sqrt{T}[\overline{\theta} - \theta_0] \rightarrow N[0, B_T^{-1/2}M_TB_T^{-1/2}]$ and $\sqrt{T}\overline{\gamma} \rightarrow N[0, \left(A_TB_T^{-1}A_T'\right)^{-1}]$, respectively. Then, given an initial consistent unconstrained estimate $\widehat{\theta}$, by deriving $g_T(\overline{\theta}) \approx g_T(\widehat{\theta}) + G_T(\overline{\theta} - \widehat{\theta})$ and $a_T(\overline{\theta}) \approx a_T(\widehat{\theta}) + A_T(\overline{\theta} - \widehat{\theta})$, and substituting into the first-order conditions, Bekaert and Hodrick (2001) define the following iterative scheme

$$\overline{\theta} \approx \widehat{\theta} - B_T^{-1/2} M_T B_T^{-1/2} G_T' \Omega_T^{-1} g_T(\widehat{\theta}) - B_T^{-1} A_T' \left(A_T B_T^{-1} A_T'\right)^{-1} a_T(\widehat{\theta}) \quad (A.2.4)$$

$$\overline{\gamma} \approx -\left(A_T B_T^{-1} A_T'\right)^{-1} A_T B_T^{-1} G_T' \Omega_T^{-1} g_T(\widehat{\theta}) + \left(A_T B_T^{-1} A_T'\right)^{-1} a_T(\widehat{\theta}) \quad (A.2.5)$$

To obtain the constrained parameters $\overline{\theta}$, we iterate on equations (A.8) and (A.9), substituting the first constrained estimate for the initial consistent unconstrained estimate to derive a second constrained estimate and so forth. The iterative process continues until the constrained estimate satisfies the constraints, that is $a_T(\overline{\theta}) = 0$.

A.3 Small Sample bias correction

Let $Z_t = [i_t, S_t^{(1w)}, S_t^{(2w)}, S_t^{(3w)}, S_t^{(1m)}, S_t^{(2m)}, S_t^{(3m)}]'$, where $S_t^{(j)}$ denotes the spread between repo rate $i_t^{(j)}$ and the overnight repo rate i_t, and assume a $VAR(p)$ dynamics

$$Z_t = \varphi + \sum_{j=1}^{p} \Phi_j Z_{t-j} + \varepsilon_t \quad (A.3.1)$$

where φ is a vector of constant and Φ_j is a square matrix. Under the assumption of homoskedastic innovations, we proceed as follows. Estimate equation (A.10) on the original data set and simulate 100,000 artificial data sets of 3,625 by using an i.i.d. bootstrap of ε_t. Next, reestimate equation (A.10) for each replication and determine bias as the difference between the parameter estimates of the initial data set and the average of the parameter estimates of the artificial data sets. Then, correct the original parameters, simulate 70,000 observations, and add the simulated i_t to each simulated spread $S_t^{(j)}$. This bias corrected data set is, hence, subjected for each pairwise combination of short-term and long-term rate to the analysis described in Section 1.3.

In the second DGP, reparameterise $\varepsilon_t = F\eta_t$, to capture the effects of temporal heteroskedasticity, where η_t is a vector of idiosyncratic innovations and F is a 7×7 factor loadings matrix defined as

$$F = \begin{bmatrix} 1 & & & 1 \\ f_{21} & \ddots & & f_{27} \\ \vdots & & \ddots & \vdots \\ f_{71} & & & 1 \end{bmatrix} \tag{A.3.2}$$

where the blank elements are zero. Define $E_{t-1}\left[\eta_t'\eta_t\right] = V_t$, and $E_{t-1}\left[\varepsilon_t'\varepsilon_t\right] = FV_tF'$, where V_t is a diagonal matrix and each element is assumed to follow an GARCH(1,1) process augmented with square root of overnight rate, $h_{jt} = \omega_j\sqrt{i_{t-1}} + \beta_j h_{jt-1} + \alpha_j\eta_{jt-1}^2$ with $j \in \{1,\ldots,7\}$, as in Gray (1996), Bekaert and Hodrick (2001), Longstaff (2000a), and Ang and Bekaert (2002), in order to accommodate shifts in the short-rate volatility. Hence, estimate equation (A.3.1) and proceed with bias correction as in the previous experiment, Next, compute the residual vector ε_t, estimate the factor GARCH parameters via quasi-maximum likelihood, and simulate a second bias corrected data set as in the previous experiment. Finally, we always generate additional $1,000$ discarding values to avoid any dependence on the starting values.

B Appendix: Bayesian MCMC Estimation

B.1 Prior Specification

We perform Bayesian MCMC estimation of the parameters of the empirical exchange rate models by constructing a Markov chain whose limiting distribution is the target posterior density. This Markov chain is a Gibbs sampler in which all parameters are drawn sequentially from their full conditional posterior distribution. The chain is then iterated and the sampled draws, beyond a burn-in period, are treated as variates from the target posterior distribution.

For the conditional mean parameters, $\theta_1 = \{\alpha, \beta\}$, we assume a Normal prior $N(\underline{\theta}_1,\ \underline{V})$, where $\underline{\theta}_1 = 0_2$ and $\underline{V} = I_2$. In the Linear Regression model, we define $\theta_2 = \{v^{-2}\}$ as the inverse of the variance and assume a prior $Gamma\left(\frac{\underline{\nu}}{2}, \frac{2\underline{s}^{-2}}{\underline{\nu}}\right)$ with mean $\underline{s}^{-2} = 1$, and degrees of freedom $\underline{\nu} = 2$.

In the GARCH(1,1) model, $\theta_2 = \{\omega, \gamma_1, \gamma_2\}$ are the conditional variance parameters. We ensure that the conditional variance is covariance stationary by specifying ω as a logNormal prior: $\omega \sim LogN\left(\underline{w}, \underline{W}\right)$, with $\underline{w} = -1$ and $\underline{W} = 2$. The prior specification is completed by assuming $\gamma_1 \sim Beta\left(\underline{g}_1, \underline{G}_1\right)$, and $\gamma_2 \sim Beta\left(\underline{g}_2, \underline{G}_2\right)$, where $\underline{g}_1 = 40$, $\underline{G}_1 = 5$, $\underline{g}_2 = 2$, and $\underline{G}_2 = 40$. These hyperparameters imply a mean of 0.89 and 0.05 for γ_1 and γ_2, respectively.

In the SV model, $\theta_2 = \{\mu, \phi, \sigma\}$ are the conditional log-variance parameters. Our prior for μ is $N\left(\underline{m}, \underline{M}\right)$ with $\underline{m} = 1$ and $\underline{M} = 25$. Following Kim, Shephard, and Chib (1998), we formulate the prior for ϕ in terms of $\phi = 2\phi^* - 1$, where ϕ^* is distributed as $Beta(\underline{f}, \underline{F})$. This implies that the prior on $\phi \in (-1, 1)$ is $p(\phi) = \kappa\{0.5(1+\phi\}^{\underline{f}-1}\{0.5(1-\phi\}^{\underline{F}-1}$, $\underline{f}, \underline{F} > 0.5$, where $\kappa = 0.5\frac{\Gamma(\underline{f}+\underline{F})}{\Gamma(\underline{f})+\Gamma(\underline{F})}$. Specifying

$\underline{f} = 20$ and $\underline{F} = 1.5$ yields a mean of 0.86 with variance of 0.01. For σ, the prior is inverse gamma $IG(\underline{s}, \underline{S})$ with $\underline{s} = 3$ and $\underline{S} = 2.5$ so that the distribution has a mean of 0.20 with variance 0.006.

For all models, the hyperparameters are set to reasonable values, but the algorithms described below are robust to the prior specification and initial values.

B.2 The Linear Regression Algorithm

In the Bayesian Linear Regression (LR) model, we need to estimate $\theta = \{\theta_1, \theta_2\}$, where $\theta_1 = \{\alpha, \beta\}$ is the set of the conditional mean parameters, and $\theta_2 = \{v^{-2}\}$ is the constant precision defined as the inverse of the variance. The simple Gibbs algorithm is summarised below (for more details see Koop, 2003):

1. Initialise θ_2.

2. Sample θ_1 from $\theta_1 \mid \Delta s, \theta_2 \sim N\left(\overline{\theta}_1, \overline{V}\right)$, where $\overline{V} = \left(\underline{V}^{-1} + \theta_2 X'X\right)^{-1}$, and $\overline{\theta}_1 = \overline{V}\left(\underline{V}^{-1}\underline{\theta}_1 + \theta_2 X'\Delta s\right)$.

3. Sample θ_2 from $\theta_2 \mid \Delta s, \theta_1 \sim Gamma\left(\frac{\overline{\nu}}{2}, \frac{2\overline{s}^{-2}}{\overline{\nu}}\right)$, where $\overline{\nu} = T + \underline{\nu}$, and $\overline{s}^2 = \frac{(\Delta s - X\theta_1)'(\Delta s - X\theta_1) + \underline{\nu s}^2}{\overline{\nu}}$.

4. Go to step 2 and iterate 100000 times beyond a burn-in of 20000 iterations.

B.3 The GARCH(1,1) Algorithm

In the Bayesian GARCH(1,1) model, we need to estimate $\theta = \{\theta_1, \theta_2\}$, where $\theta_1 = \{\alpha, \beta\}$ is the set of the conditional mean parameters, and $\theta_2 = \{\omega, \gamma_1, \gamma_2\}$ are the conditional variance parameters. The algorithm is summarised below:[50]

[50]We have performed a simulation study for comparing the mean square error (MSE) of the GARCH(1,1) parameter estimates resulting from maximum likelihood and Bayesian MCMC es-

1. Initialise θ_1 and transform the data into $\Delta s_t^* = (\Delta s_t - \alpha - \beta x_{t-1})$.

2. Sample the variance parameters θ_2 from their full conditional posterior density: $\theta_2 \mid \Delta s^*, \theta_1$. This posterior density is not available analytically. We compute the log-likelihood of the transformed data Δs_t^* as function of θ_2 (conditional on θ_1) and then we optimise the conditional log-posterior. We generate a proposal from a t-distribution $t(m, V, \xi)$, where m is the mode, V is the inverse of the negative Hessian, and ξ a tuning parameter. The proposal is then accepted as according to the independence chain Metropolis-Hasting algorithm (e.g. Chib and Greenberg, 1995).

3. Sample all the conditional mean coefficients $\theta_1 \mid \Delta s, \theta_2$ using a precision-weighted average of a set of normal priors and the normal likelihood conditional on θ_2.

4. Update the data $\Delta s_t^* = (\Delta s_t - \alpha - \beta x_{t-1})$.

5. Go to step 2 and iterate 5000 times beyond a burn-in of 1000 iterations.

B.4 The Stochastic Volatility Algorithm

In the Bayesian SV model, we need to estimate $\theta = \{\theta_1, \theta_2\}$, where $\theta_1 = \{\alpha, \beta\}$ is the set of the conditional mean parameters, and $\theta_2 = \{\mu, \phi, \sigma^2\}$ are the conditional log-variance parameters. The parameters of the SV model are estimated using the Bayesian MCMC algorithm of Chib, Nardari, and Shephard (2002), which builds on

timation methods. We set the true model parameters: $\omega = 0.0005$, $\gamma_1 = 0.70$, $\gamma_2 = 0.25$, and $v_0 = 0.003$ as in Vrontos, Dellaportas and Politis (2000). For small ($T = 300$), medium ($T = 1000$), and large ($T = 5000$) sample sizes, we generated 10000 artificial samples for which we then estimated the GARCH parameters using the two estimation methods. We find that the Bayesian MCMC estimates have lower MSE values than the maximum likelihood estimates.

the procedures developed by Kim, Shephard, and Chib (1998), and is summarised below:

1. Initialise θ, mx, and transform the data into

$$\Delta s_t^* = \ln\left((\Delta s_t - \alpha - \beta x_{t-1})^2 + c\right),$$

$c = 0.001$ to put the model in state-space form. The "offset" constant c eliminates the inlier problem.

2. Sample the log-variance parameters θ_2 from their full conditional posterior density: $\theta_2 \mid \Delta s^*, mx$. This posterior density is not available analytically. We use the Kalman filter to compute the log-likelihood of the transformed data Δs_t^* as a function of θ_2 (conditional on mx_t) and then optimise the conditional log-posterior. We generate a proposal from a t-distribution $t(m, V, \xi)$, where m is the mode, V is the inverse of the negative Hessian, and ξ a tuning parameter. The proposal is then accepted according to the independence chain Metropolis-Hastings algorithm (e.g. Chib and Greenberg, 1995).

3. Sample the log-variance vector $\{h_t\}$ in one block from the posterior distribution: $h \mid \Delta s^*, mx, \theta_2$. This step uses the de Jong and Shephard (1995) simulation smoother, which is an algorithm designed for efficient sampling of the state vector in a state-space model.

4. Sample all the conditional mean coefficients θ_1 from $\theta_1 \mid \Delta s, h$ using a precision-weighted average of a set of normal priors and the normal likelihood conditional on h. Then update the transformed data

$$\Delta s_t^* = \ln\left((\Delta s_t - \alpha - \beta x_{t-1})^2 + c\right),$$

with $c = 0.001$.

5. Finally, sample the mixture indicator variable $mx \mid \Delta s^*, h, \theta$ directly from its posterior:

$$\Pr(mx_t \mid \Delta s_t^*, h_t) \propto \Pr(mx_t) f_N\left(\Delta s_t^* \mid h_t + m_{mx_t}, \upsilon_{mx_t}^2\right), \qquad t \leq T$$

where $\left\{m_{mx_t}, \upsilon_{mx_t}^2\right\}$ are the means and variances of the seven-component mixture of normal densities which are used to approximate the $\log \chi^2(1)$ distribution (see Kim, Shephard, and Chib, 1998).

6. Go to step 2 and iterate 5000 times beyond a burn-in of 1000 iterations.

B.5 The Particle Filter

The particle filter of Pitt and Shephard (1999) generates a sample from the density $h_t \mid F_t, \theta$. This is a non-trivial task performed by an Auxiliary Sampling-Importance Resampling algorithm. The SV application of the algorithm is detailed in Chib, Nardari and Shephard (2002) and sketched below:

1. Given a sample $\left\{h_{t-1}^1, ..., h_{t-1}^M\right\}$ from $(h_{t-1} \mid F_{t-1}, \theta)$ calculate: $\widehat{h}_t^{*j} = \mu + \phi\left(h_{t-1}^j - \mu\right)$, $\omega_j = N\left(\Delta s_t \mid \alpha + \beta x_{t-1}, \exp(\widehat{h}_t^{*j})\right)$, for $j = 1, ...M$. Sample $R = 10000$ times the integers $1, 2, ..., M = 2000$ with probability proportional to $\{\omega_j\}$. Let the sampled indices be $k_1, ..., k_R$ and associate these with $h_t^{*k_1}, ..., h_t^{*k_R}$.

2. For each value of k_j from Step 1, simulate the values $\left\{h_t^1, ..., h_t^R\right\}$ from the volatility process as: $h_t^{*j} = \mu + \phi\left(h_{t-1}^{k_j} - \mu\right) + \sigma\eta_t^j$, $j = 1, ..., R$, where $\eta_t^j \sim N(0, 1)$.

3. Resample the values $\{h_t^{*1}, ..., h_t^{*R}\}$ M times with replacement using probabilities proportional to: $\frac{N\left(\Delta s_t|\alpha+\beta_1 x_{t-1}, \exp\left(h_t^{*j}\right)\right)}{N\left(\Delta s_t|\alpha+\beta_1 x_{t-1}, \exp\left(h_t^{*k_j}\right)\right)}$, for $j = 1, ..., R$, to produce the desired filtered sample $\{h_t^1, ..., h_t^M\}$ from $(h_t \mid F_t, \theta)$.

References

Abhyankar, A., L. Sarno, and G. Valente (2005). "Exchange Rates and Fundamentals: Evidence on the Economic Value of Predictability," *Journal of International Economics*, **66**, 325-348.

Ahn, D.H., R.F. Dittmar, and A.R. Gallant (2002). "Quadratic Term Structure Models: Theory and Evidence," *Review of Financial Studies*, **15**, 243-288.

Andersen, T.G. (1996). "Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility," *Journal of Finance*, **51**, 169-204.

Ang, A., and G. Bekaert (2002). "Regime Switches in Interest Rates," *Journal of Business and Economic Statistics*, **20**, 163-182.

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

Backus, D.K., A.W. Gregory, and C.I. Telmer (1993). "Accounting for Forward Rates in Markets for Foreign Currency," *Journal of Finance*, **48**, 1887-1908.

Backus, D.K., S. Foresi, and C.I. Telmer (2001). "Affine Term Structure Models and the Forward Premium Anomaly," *Journal of Finance*, **56**, 279-304.

Baks, K., A. Metrick, and J.A. Wachter (2001). "Should Investors Avoid All Actively Managed Mutual Funds? A Study in Bayesian Performance Evaluation," *Journal of Finance*, **56**, 45-85.

Bansal, R. (1997). "An Exploration of the Forward Premium Puzzle in Currency Markets," *Review of Financial Studies*, **10**, 369-403.

Bansal, R., and M. Dahlquist (2000). "The Forward Premium Puzzle: Different Tales from Developed and Emerging Economies," *Journal of International Economics*, **51**, 115-144

Bansal, R., and H. Zhou (2002).. "Term Structure of Interest Rates with Regime Shifts," *Journal of Finance*, **57**, 1997–2043.

Bansal, R., A.R. Gallant, R. Hussey, and G. Tauchen (1995). "Nonparametric Estimation of Structural Models for High Frequency Currency Market Data," *Journal of Econometrics*, **66**, 251-287.

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

Bates, J.M., and C.W.J. Granger (1969). "The Combination of Forecasts," *Operations Research Quarterly*, **20**, 451-468.

Bauer, G.H. (2001). "Conditional Currency Hedging and Asset Market Shocks", working paper, University of Rochester.

Bekaert, G. (1996). "The Time-Variation of Risk and Return in Foreign Exchange Markets: A General Equilibrium Perspective," *Review of Financial Studies*, **9**, 427-470.

Bekaert, G., and R.J. Hodrick (1993). "On Biases in the Measurement of Foreign Exchange Risk Premiums," *Journal of International Money and Finance*, **12**, 115-138.

Bekaert, G., and R.J. Hodrick (2001). "Expectations Hypothesis Tests," *Journal of Finance*, **56**, 115-138.

Bekaert, G., R.J. Hodrick, and D. Marshall (1997). "The Implications of First-Order Risk Aversion for Asset Market Risk Premiums," *Journal of Monetary Economics*, **40**, 3-39.

Berkowitz, J., and L. Giorgianni (2001). "Long-Horizon Exchange Rate Predictability?," *Review of Economics and Statistics*, **83**, 81-91.

Bilson, J.F.O. (1981). "The 'Speculative Efficiency' Hypothesis," *Journal of Business*, **54**, 435-451.

Bollerslev, T. (1986). "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, **31**, 307-327.

Bollerslev, T. (1987). "A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return," *Review of Economics and Statistics*, **69**, 542-547.

Bollerslev, T. (1990). "Modeling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH Model," *Review of Economics and Statistics*, **72**, 498-505.

Boudoukh, J., M. Richardson, and R.F. Whitelaw (2006). "The Information in Long-Maturity Forward Rates: Implications for Exchange Rates and the Forward Premium Anomaly," working paper, New York University.

Burnham, K.P., and D.P. Anderson (2002). *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*, 2nd edition, New York: Springer-Verlag.

Campbell, J.Y., and R.J. Shiller (1987). "Cointegration and Tests of Present Value Models," *Journal of Political Economy*, **95**, 1062-1088.

Campbell, J.Y., and R.J. Shiller (1991). "Yield Spreads and Interest Rate Movements: A Bird's Eye View," *Review of Economic Studies*, **58**, 495-514.

Chib, S. (1995). "Marginal Likelihood from the Gibbs Output," *Journal of the American Statistical Association*, **90**, 1313-1321.

Chib, S., and E. Greenberg (1995). "Understanding the Metropolis-Hastings Algorithm," *American Statistician*, **49**, 327-335.

Chib, S., and I. Jeliazkov (2001). "Marginal Likelihood from the Metropolis-Hastings Output," *Journal of the American Statistical Association*, **96**, 270-281.

Chib, S., F. Nardari, and N. Shephard (2002). "Markov Chain Monte Carlo Methods for Stochastic Volatility Models," *Journal of Econometrics*, **108**, 281-316.

Chib, S., F. Nardari, and N. Shephard (2006). "Analysis of High Dimensional Multivariate Stochastic Volatility Models," *Journal of Econometrics*, **134**, 341-371.

Clarida, R.H., and M.P. Taylor (1997). "The Term Structure of Forward Exchange Premiums and the Forecastability of Spot Exchange Rates: Correcting the Errors," *Review of Economics and Statistics*, **79**, 353-361.

Clarida, R.H., L. Sarno, M.P. Taylor, and G. Valente (2003). "The Out-of-Sample Success of Term Structure Models as Exchange Rate Predictors: A Step Beyond," *Journal of International Economics*, **60**, 61-83.

Clarida, R.H., L. Sarno, M.P. Taylor, and G. Valente (2006). "The Role of Asymmetries and Regime Shifts in the Term Structure of Interest Rates," *Journal of Business*, **79**, 1193-1225.

Collin-Dufresne, P., and R.S. Goldstein (2002). "Do Bonds Span the Fixed Income Markets? Theory and Evidence for Unspanned Stochastic Volatility," *Journal of Finance*, **57**, 1685-1730.

Collin-Dufresne, P., R.S. Goldstein, and C.S. Jones (2007). "Identification of Maximal Affine Term Structure Models," *Journal of Finance*, forthcoming.

Cremers, M. (2002). "Stock Return Predictability: A Bayesian Model Selection Perspective," *Review of Financial Studies*, **15**, 1223-1249.

Csiszar, I. (1991). "Why Least Squares and Maximum Entropy? An Axiomatic Approach to Inference for Linear Inverse Problems," *Annals of Statistics*, **19**, 2032-2066.

Dai, Q., and K.J. Singleton (2000). "Specification Analysis of Affine Term Structure Models," *Journal of Finance*, **55**, 1943-1978.

de Jong, P., and N. Shephard (1995). "The Simulation Smoother for Time Series Models," *Biometrika*, **82**, 339-350.

Diebold, F.X. (1998). "The Past, Present and Future of Macroeconomic Forecasting," *Journal of Economic Perspectives*, **12**, 175-192.

Diebold, F.X. (2004). *Elements of Forecasting*, Third Edition, South-Western College Publishing.

Diebold, F.X., and P. Pauly, (1990). "The Use of Prior Information in Forecast Combination," *International Journal of Forecasting*, **6**, 503-508.

Duffie, D., and K.J. Singleton (1999). "Modeling Term Structures of Defaultable Bonds," *Review of Financial Studies*, **12**, 687-720.

Engel, C. (1996). "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence," *Journal of Empirical Finance*, **3**, 123-192.

Engel, C., and K.D. West (2005). "Exchange Rates and Fundamentals," *Journal of Political Economy*, **113**, 485-517.

Engle, R. (2002). "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models," *Journal of Business and Economic Statistics*, **20**, 339-350.

Fama, E.F. (1984a) "Forward and Spot Exchange Rates," *Journal of Monetary Economics*, **14**, 319-338.

Fama, E.J. (1984b). "The Information in the Term Structure," *Journal of Financial Economics*, **13**, 509-528.

Fama, E.J., and R.R. Bliss (1987). "The information in Long-Maturity Forward Rates," *American Economic Review*, **77**, 680-692.

Faust, J., J.H. Rogers, and J. Wright (2003). "Exchange Rate Forecasting: The Errors We've Really Made," *Journal of International Economics*, **60**, 35-59.

Fernandez-Villaverde, J., and J.F. Rubio-Ramirez (2004). "Comparing Dynamic Equilibrium Models to Data: A Bayesian Approach," *Journal of Econometrics*, **123**, 153-187.

Fisher, I. (1896). "Appreciation and Interests," *Publications of the American Economic Association*, **11**, 21-29.

Fleming, J., and C. Kirby (2003). "A Closer Look at the Relation between GARCH and Stochastic Autoregressive Volatility," *Journal of Financial Econometrics*, **1**, 365-419.

Fleming, J., C. Kirby, and B. Ostdiek (2001). "The Economic Value of Volatility Timing," *Journal of Finance*, **56**, 329-352.

Flood, R.B., and A.K. Rose (2002). "Uncovered Interest Parity in Crisis," *International Monetary Fund Staff Papers*, **49**, 252-266.

Frankel, J.A., and K.A. Froot (1987). "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations," *American Economic Review*, **77**, 133-153.

Froot, K.A. (1989). "New Hope for the Expectations Hypothesis of the Term Structure of Interest Rates," *Journal of Finance*, **44**, 283-305.

Froot, K.A., and R.H. Thaler (1990). "Foreign Exchange," *Journal of Economic Perspectives*, **4**, 179-192.

Galati, G., and M. Melvin (2004). "Why Has FX Trading Surged? Explaining the 2004 Triennial Survey," *BIS Quarterly Review*, December, 67-74.

Geweke, J. (1989). "Bayesian Inference in Econometric Models Using Monte Carlo Integration," *Econometrica*, **57**, 1317-1339.

Gomez, V., and A. Maravall (2000). "Seasonal Adjustment and Signal Extraction in Economic Time Series," in Pena D., G.C. Tiao, and T.R. Tsay (eds.), *A Course in Time Series Analysis*, New York: Wiley.

Gray, S. F. (1996). "Modeling the Conditional Distribution of Interest Rates as a Regime Switching Process," *Journal of Financial Economics*, **42**, 27-62.

Groen, J.J.J. (2000). "The Monetary Exchange Rate Model as a Long-Run Phenomenon," *Journal of International Economics,* **52**, 299-319.

Hai, W., N.C. Mark, and Y. Wu (1997). "Understanding Spot and Forward Exchange Rate Regressions," *Journal of Applied Econometrics*, **12**, 715-34.

Han, Y. (2006). "Asset Allocation with a High Dimensional Latent Factor Stochastic Volatility Model," *Review of Financial Studies*, **19**, 237-271.

Hansen, L.P. (1982). "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, **50**, 1029-1054.

Hansen, L.P., and T.J. Sargent (1980). "Formulating and Estimating Dynamic Linear Rational Expectations Models," *Journal of Economic Dynamics and Control,* **2**, 7-46.

Harvey A.C., E. Ruiz, and N. Shephard (1994). "Multivariate Stochastic Variance Models," *Review of Economics Studies*, **61**, 247-264.

Hlawitschka, W. (1994). "The Empirical Nature of Taylor-Series Approximations to Expected Utility," *American Economic Review*, **84**, 713-719.

Hodrick, R.J. (1987). *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*, London: Harwood.

Hoeting, J.A., D. Madigan, A.E. Raftery, and C.T. Volinsky (1999). "Bayesian Model Averaging: A Tutorial," *Statistical Science*, **14**, 382-417.

Jagannathan, R., A. Kaplan, and S.G. Sun (2003). "An Evaluation of Multi-Factor CIR Models Using LIBOR, Swap Rates, and Cap and Swaption Prices," *Journal of Econometrics*, **116**, 113-146.

Kandell, S., and R.F. Stambaugh (1996). "On the Predictability of Stock Returns: An Asset-Allocation Perspective," *Journal of Finance*, **51**, 385-424.

Karolyi, G.A., and R.M. Stulz (2003). "Are Financial Assets Priced Locally or Globally?," in: G. Constantinides, M. Harris, and R.M. Stulz (eds.), *Handbook of Economics of Finance*, Amsterdam: Elsevier.

Kim, S., N. Shephard, and S. Chib (1998). "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models," *Review of Economic Studies*, **65**, 361-393.

Koop, G. (2003). *Bayesian Econometrics*, New York: Wiley.

Leamer, E.E. (1978). *Specification Searches*, New York: Wiley.

Levich, R.M. (1985). "Empirical Studies of Exchange Rates: Price Behavior, Rate Determination, and Market Efficiency," in Jones, R.W., and P.B. Kenen (eds.), *Handbook of International Money and Finance*, Amsterdam: Elsevier.

Lewis, K.K. (1995). "Puzzles in International Financial Markets," in Grossman, G.M., and K. Rogoff (eds.), *Handbook of International Economics*, Amsterdam: Elsevier.

Lo, A.W. (2002). "The Statistics of Sharpe Ratios," *Financial Analysts Journal*, **58**, 36-52.

Longstaff, F. (2000a). "The Term Structure of Very Short Term Rates: New Evidence for the Expectations Hypothesis," *Journal of Financial Economics*, **58**, 397-415.

Longstaff, F. (2000b). "Arbitrage and the Expectations Hypothesis," *Journal of Finance*, **55**, 989-994.

Lucas, R.E. (1982). "Interest Rates and Currency Prices in a Two-Country World," *Journal of Monetary Economics*, **10**, 335-359.

Lustig, H., and A. Verdelhan (2007). "The Cross-Section of Currency Risk Premia and US Consumption Growth Risk," *American Economic Review*, forthcoming.

Lyons, R.K. (2001). *The Microstructure Approach to Exchange Rates*, Cambridge: MIT Press.

Mark, N.C. (1995). "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability," *American Economic Review*, **85**, 201-218.

Mark, N.C., and D. Sul. (2001). "Nominal Exchange Rates and Monetary Fundamentals: Evidence from a Small Post-Bretton Woods Panel," *Journal of International Economics*, **53**, 29-52.

Marquering, W., and M. Verbeek (2004). "The Economic Value of Predicting Stock Index Returns and Volatility," *Journal of Financial and Quantitative Analysis*, **39**, 407-429.

Meese, R.A., and K. Rogoff (1983). "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?" *Journal of International Economics*, **14**, 3-24.

Melino, A. (2001). "Estimation of a Rational Expectations Model of the Term Structure," *Journal of Empirical Finance*, **8**, 639-668.

Mishkin, F.S. (1982). "Does Anticipated Monetary Policy Matter? An Econometric Investigation," *Journal of Political Economy*, **90**, 22-51.

Modigliani, F., and L. Modigliani (1997). "Risk-Adjusted Performance," *Journal of Portfolio Management,* **23**, 45-54.

Newey, W. K., and D.L. McFadden (1994). "Large Sample Estimation and Hypothesis Testing," in Engle, R.F., and D.L. McFadden (eds.), *Handbook of Econometrics,* **4**, Amsterdam: Elsevier.

Obstfeld, M., and K. Rogoff (1995). "Exchange Rate Dynamics Redux," *Journal of Political Economy*, **103**, 624-660.

Pagan, A. (1984). "Econometric Issues in the Analysis of the Regressions With Generated Regressors," *International Economic Review,* **25**, 221-247.

Pitt, M.K., and N. Shephard (1999). "Filtering via Simulation Based on Auxiliary Particle Filters," *Journal of the American Statistical Association*, **94**, 590-599.

Rapach, D.E., and M.E. Wohar (2002). "Testing the Monetary Model of Exchange Rate Determination: New Evidence from a Century of Data," *Journal of International Economics*, **58**, 359-385.

Roll, R. (1970). *The Behavior of Interest Rates: An Application of the Efficient Market Model to U.S. Treasury Bills.* New York, NY: Basic Books.

Sargent, T.J. (1977). "The Demand for Money during Hyperinflation under Rational Expectations," *International Economic Review,* **18**, 59-82.

Sarno, L., and M.P. Taylor (2003). *The Economics of Exchange Rates*, Cambridge: Cambridge University Press.

Sarno, L., D.L. Thornton, and G. Valente (2007). "The Empirical Failure of the Expectations Hypothesis of the Term Structure of Bond Yields," *Journal of Financial and Quantitative Analysis,* **42**, 81-100.

Sarno, L., G. Valente, and H. Leon (2006). "Nonlinearity in Deviations from Uncovered Interest Parity: An Explanation of the Forward Bias Puzzle," *Review of Finance*, **10**, 443-482.

Shanken, J., and A. Tamayo (2001). "Risk, Mispricing, and Asset Allocation: Conditioning on Dividend Yield," NBER working paper 8666.

Shiller, R.J. (1979). "The Volatility of Long-Term Interest Rates and Expectation Models of the Term Structure," *Journal of Political Economy*, **87**, 1190-1219.

Shiller, R., J.Y. Campbell, and K.L. Schoenholtz (1983). "Forward Rates and Future Policy: Interpreting the Term Structure of Interest Rates," *Brookings Papers on Economic Activity*, **1**, 173-217.

Sims, C.A. (1980). "Macroeconomics and Reality," *Econometrica*, **48**, 1-48

Sims, C.A. (1988). "Bayesian Skepticism on Unit Root Econometrics," *Journal of Economic Dynamics and Control*, **12**, 463-474.

Stambaugh, R.F. (1988). "The Information in Forward Rates: Implications for Models of the Term Structure," *Journal of Financial Economics*, **21**, 41–70.

Tauchen, G.E., and M. Pitts (1983). "The Price Variability-Volume Relationship on Speculative Markets," *Econometrica*, **51**, 485-505.

Thornton, D.L. (2005). "Tests of the Expectations Hypothesis: Resolving Some Anomalies When the Short-Term Rate is the Federal Funds Rate," *Journal of Banking and Finance*, **29**, 2541-2556.

Thornton, D.L. (2006). "Tests of the Expectations Hypothesis: Resolving the Campbell-Shiller Paradox," *Journal of Money, Credit, and Banking*, **38**, 511-542.

Timmermann, A. (2006). "Forecast Combinations," in Elliott, G., C.W.J. Granger, and A. Timmermann (eds.), *Handbook of Economic Forecasting*, Amsterdam: Elsevier.

Vrontos, I.D., P. Dellaportas, and D.N. Politis (2000). "Full Bayesian Inference for GARCH and EGARCH Models," *Journal of Business and Economic Statistics*, **18**, 187-198.

Wallis, K.F. (1980). "Econometric Implications of the Rational Expectations Hypothesis," *Econometrica*, **48**, 49-72.

West, K.D., H.J. Edison, and D. Cho (1993). "A Utility-Based Comparison of Some Models of Exchange Rate Volatility," *Journal of International Economics*, **35**, 23-45.