

Financial Asset Returns, Direction-of-Change Forecasting, and Volatility Dynamics

Peter F. Christoffersen
Faculty of Management
McGill University
and CIRANO

peter.christoffersen@mcgill.ca

Francis X. Diebold
Wharton School
University of Pennsylvania
and NBER

fdiebold@wharton.upenn.edu

This Draft / Print: December 13, 2005

Send all editorial correspondence to:

F.X. Diebold
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297

fdiebold@wharton.upenn.edu

Management Insights: Forecasting the direction-of-change in asset prices is essential for financial analysts who make “buy”, “sell” and “hold” recommendations. We show that the direction-of-change in an asset price is forecastable from information on the volatility of the asset. Asset price volatility is known to be highly forecastable which in turn renders the direction-of-change forecastable. We argue that volatility-driven direction-of-change forecasting is likely to be most effective at the monthly and quarterly horizons.

Financial Asset Returns, Direction-of-Change Forecasting, and Volatility Dynamics

This Draft / Print: December 13, 2005

Abstract: We consider three sets of phenomena that feature prominently in the financial economics literature: conditional mean dependence (or lack thereof) in asset returns, dependence (and hence forecastability) in asset return signs, and dependence (and hence forecastability) in asset return volatilities. We show that they are very much interrelated, and we explore the relationships in detail. Among other things, we show that: (a) Volatility dependence produces sign dependence, so long as expected returns are nonzero, so that one should expect sign dependence, given the overwhelming evidence of volatility dependence; (b) It is statistically possible to have sign dependence without conditional mean dependence; (c) Sign dependence is not likely to be found via analysis of sign autocorrelations, runs tests, or traditional market timing tests, because of the special nonlinear nature of sign dependence, so that traditional market timing tests are best viewed as tests for sign dependence arising from variation in expected returns rather than from variation in volatility or higher moments; (d) Sign dependence is not likely to be found in very high-frequency (e.g., daily) or very low-frequency (e.g., annual) returns; instead, it is more likely to be found at intermediate return horizons; (e) The link between volatility dependence and sign dependence remains intact in conditionally non-Gaussian environments, as for example with time-varying conditional skewness and/or kurtosis.

Acknowledgments: This work was supported by the National Science Foundation, the Guggenheim Foundation, the Wharton Financial Institutions Center, FQRSC, IFM2 and SSHRC. For helpful comments we thank the Editor, David Hsieh, and two anonymous referees, as well as David Bates, Antulio Bomfim, Michael Brandt, Xiaohong Chen, Michel Dacorogna, Graham Elliott, Rene Garcia, Christian Gourieroux, Clive Granger, Anthony Hall, Wolfgang Härdle, Joanna Jasiak, Michael Johannes, Blake LeBaron, Bruce Lehman, Martin Lettau, Nour Meddahi, Theo Nijman, Jonathan Reeves, Sergei Sarkissian, Frank Schorfheide, Allan Timmermann, Harald Uhlig, Pietro Veronesi, Ken West, Hal White, Jonathan Wright, and seminar participants at the European Central Bank, the Federal Reserve Board, McGill University, UCSD, Penn, the Third Annual Conference on Financial Econometrics at the University of Waterloo, the Econometric Society Winter Meetings in Washington DC, the NBER/NSF Time Series Meeting in Philadelphia, the NYU Stern Conference in Honor of the 2003 Nobel Prize in Economics, the Northern Finance Association annual meetings, the European Finance Association annual meetings, Singapore Management University, and the Federal Reserve Bank of St. Louis. Sean Campbell, Chayawat Ornthanalai and Clara Vega provided outstanding research assistance. All inadequacies are ours alone.

1. Introduction

We consider three sets of phenomena that feature prominently in the financial economics literature: approximate conditional mean independence (and hence little or no forecastability) in asset returns, dependence (and hence forecastability) in asset return signs, and dependence (and hence forecastability) in asset return volatilities. We argue that they are very much interrelated, forming a tangled and intriguing web, a full understanding of which leads to a deeper understanding of the subtleties of financial market dynamics. Let us introduce them in turn.

First, consider conditional mean independence, by which we mean that an asset return's conditional mean does not vary with the conditioning information set. Approximate conditional mean independence then implies that out-of-sample return forecasting will be difficult, if not impossible. This view is widely held and stems from both introspection and observation. That is, financial economic theory suggests that asset returns should not be easily forecast using readily-available information and forecasting techniques, and a broad interpretation of four decades of empirical work suggests that the data support the theory (e.g., Fama, 1970, 1991). Consequently, we believe that conditional mean independence is reasonably viewed as a good working approximation to asset return dynamics.¹

Second, consider dependence and hence forecastability of market direction (the return *sign*). Profitable trading strategies result from successful forecasting of market direction, quite apart from whether one is successful at forecasting returns themselves. A well-known and classic example, discussed routinely even at the MBA textbook level (e.g., Levich, 2001, Chapter 8), involves trading in speculative markets. If, for example, the Yen/\$ exchange rate is expected to increase, reflecting expected depreciation of the Yen relative to the dollar and hence a negative expected return on the Yen, one would sell Yen for Dollars, whether in the spot or derivatives markets. Positive profits will be made when the sign forecast is correct. Generalizations to multiple asset classes, such as stock and bond markets, involve basing allocation strategies on forecasts of the sign of the return spread. Recent literature shows that asset

¹ We emphasize the word “approximation,” because weak conditional mean dependence may appear at both long horizons (e.g., Fama and French, 1988, 1989, and Campbell and Shiller, 1988) and short horizons (e.g., Lo and MacKinlay, 1999), for a variety of reasons ranging from low-frequency variation in risk premia to high-frequency variation in market microstructure noise. The evidence of long-horizon conditional mean dependence remains controversial, however, due to serious statistical complications including possibly spurious regressions (e.g., Kirby, 1997), data snooping biases (e.g., Foster, Smith and Whaley, 1997), and small-sample biases (Nelson and Kim, 1993), which may distort standard inference procedures when applied to long-horizon prediction regressions.

return sign forecasting can often be done with surprising success.²

Finally, consider dependence and forecastability of asset return volatility. A huge literature documents the notable dependence, and hence forecastability, of asset return volatility, with important implications not only for asset allocation, but also for asset pricing and risk management. Bollerslev, Chou and Kroner (1992) provide a fine review of evidence in the GARCH tradition, while Ghysels, Harvey and Renault (1996) survey results from stochastic volatility modeling, Franses and van Dijk (2000) survey results from regime-switching volatility models, and Andersen, Bollerslev and Diebold (2005) survey results from realized volatility modeling. Interesting extensions include models of time-variation in higher-ordered conditional moments, such as the conditional skewness models of Harvey and Siddique (2000) and the conditional density model of Hansen (1994). The recent literature also contains intriguing theoretical work explaining the empirical phenomenon of volatility forecastability, such as Brock and Hommes (1997) and de Fontnouvelle (2000).

In this paper we characterize in detail the relationships among the three phenomena and three literatures discussed briefly above: asset return conditional mean independence, sign dependence, and conditional variance dependence. It is well known that conditional mean independence and conditional variance dependence are statistically compatible. However, much less is known in general about sign dependence, and in particular about the relationship of sign dependence to conditional mean independence and volatility dependence. Hence we focus throughout on sign dependence. Among other things, we show that:

- (a) Volatility dependence produces sign dependence, so long as expected returns are nonzero.
Hence one should expect sign dependence, given the overwhelming evidence of volatility dependence.
- (b) It is statistically possible to have sign (and volatility) dependence, while at the same time, mean independence.
- (c) Sign dependence is not likely to be found via analysis of sign autocorrelations or other tests (such as runs tests or traditional tests of market timing), because the nature of sign dependence is highly nonlinear.
- (d) Sign dependence is not likely to be found in high-frequency (e.g., daily) or low-frequency (e.g., annual) returns. Instead, it is more likely to appear at intermediate return horizons

² See, among others, Breen, Glosten and Jagannathan (1989), Leitch and Tanner (1991), Wagner, Shellans and Paul (1992), Pesaran and Timmermann (1995), Kuan and Liu (1995), Larsen and Wozniak (1995), Womack (1996), Gencay (1998), Leung, Daouk and Chen (2000), Elliott and Ito (1999), White (2000), Pesaran and Timmermann (2004), and Cheung, Chinn and Pascual (2003).

of two or three months.

- (e) The link between volatility forecastability and sign forecastability remains intact in conditionally non-Gaussian environments, as for example with time-varying conditional skewness and/or kurtosis; in fact, it is significantly enriched.

We derive results (a), (b) and (e) theoretically in a general setting, we derive (c) and (d) via a complementary calibrated simulation experiment using a popular model of return dynamics, and we provide a preliminary empirical exploration using the S&P 500 in conjunction with RiskMetrics volatility forecasts.

Before moving on, we wish to emphasize that our basic contention, namely that sign dependence exists and has fundamental connections to volatility dependence, in no way *requires* conditional mean independence; instead, we simply invoke conditional mean independence to make our point as clearly as possible. This is fortunate, because although conditional mean independence is a reasonable approximation for many purposes, it is also the subject of some controversy on both the theoretical and empirical fronts, with both theory and empirics ultimately proving inconclusive. Classical dynamic asset pricing theory predicts a positive equilibrium relationship between conditional first and second moments (e.g., Merton, 1973; Ferson and Harvey, 1991), but it is now recognized that theory is also consistent with a negative relationship (e.g., Abel, 1988; Backus and Gregory, 1993; Genotte and Marsh, 1993). Similarly inconclusive results hold on the empirical side: positive but often insignificant dependence is found in Baillie and DeGennaro (1990), French, Schwert and Stambaugh (1987), Campbell and Hentschel (1992), and Ghysels, Santa-Clara and Valkanov (2005), whereas negative dependence is found in Campbell (1987), Nelson (1991), Lettau and Ludvigson (2003), and Brandt and Kang (2004), and *both* are found in Glosten, Jagannathan and Runkle (1993), Harvey (2001), and Turner, Startz and Nelson (1989).

We proceed as follows. In Section 2 we build intuition by sketching the main results in simple contexts, focusing primarily on the conditionally Gaussian case. We discuss the basic framework, we arrive at the basic result that volatility dynamics produce sign dynamics, and we draw the implications. In Sections 3 and 4 we focus in greater depth on sign dependence, and we provide basic results on sign realizations, sign forecasts, and the relation between the two, stressing both the measurement (Section 3) and detection (Section 4) of sign forecastability. In Section 5 we perform a detailed simulation experiment, which not only illustrates our basic results but also extends them significantly, by characterizing the nature of sign forecastability as a function of forecast horizon. We provide an illustrative empirical application in Section 6, and we conclude in Section 7.

2. Conditional Mean Dependence, Sign Dependence, and Volatility Dependence: Basic Results

Here we explore the links between conditional mean dependence, sign dependence, and volatility dependence. We have used the terms repeatedly but thus far not defined them precisely, relying instead on readers' intuition, so let us begin with some precise definitions. First, we will say that a return series R_{t+1} displays conditional mean dependence (conditional mean dynamics, conditional mean forecastability, conditional mean predictability) if $E(R_{t+1}|\Omega_t)$ varies with Ω_t .³ Second, we will say that R_{t+1} displays sign dependence (sign dynamics, sign forecastability, sign predictability) if the return sign indicator series $I(R_{t+1}>0)$ displays conditional mean dependence; that is, if $E(I(R_{t+1}>0)|\Omega_t)$ varies with Ω_t .⁴ Finally, we will say that R_{t+1} displays conditional variance dependence (conditional variance dynamics, conditional variance forecastability, conditional variance predictability, volatility dependence, volatility dynamics, volatility forecastability, volatility predictability) if $\sigma_{t+1|t}^2 \equiv \text{Var}(R_{t+1}|\Omega_t)$ varies with Ω_t .

We now proceed to characterize the relationships among sign, volatility dynamics, and conditional mean dynamics.

2.1 Sign Dynamics Follow from Volatility Dynamics

Consider the prevalence of volatility dynamics in high-frequency asset returns, and the positive expected returns earned on risky assets. To take the simplest possible example, which nevertheless conveys all of the basic points, assume the returns on a generic risky asset are distributed as

$$R_{t+1} | \Omega_t \sim N(\mu, \sigma_{t+1|t}^2), \quad \mu > 0, \quad (1)$$

and therefore display conditional variance dependence but no conditional mean dependence. The probability of a positive return is then

$$\Pr_t(R_{t+1} > 0) = 1 - \Pr_t(R_{t+1} < 0) = 1 - \Pr\left(\frac{R_{t+1} - \mu}{\sigma_{t+1|t}} < \frac{-\mu}{\sigma_{t+1|t}}\right) = \Phi\left(\frac{\mu}{\sigma_{t+1|t}}\right), \quad (2)$$

where $\Phi(\bullet)$ is the $N(0,1)$ c.d.f. Notice that although the distribution is symmetric around the conditional mean, and the conditional mean is constant by assumption, the sign of the return is nevertheless forecastable, because the probability of a positive return is time-varying (and above 0.5, if $\mu > 0$). As volatility moves, so too does the probability of a positive return: the higher the volatility, the lower the

³ Here and throughout, "returns" are excess returns relative to the risk-free rate.

⁴ Equivalently, R_{t+1} displays sign dependence if the conditional probability of a positive return, $\Pr(R_{t+1} > 0 | \Omega_t)$, varies with Ω_t , because $\Pr(R_{t+1} > 0 | \Omega_t) = E(I(R_{t+1}) > 0 | \Omega_t)$.

probability of a positive return, as illustrated in Figure 1.

The surprising result that the sign of the return is forecastable although the conditional mean is constant hinges interestingly on the interaction of a non-zero mean return and non-constant volatility. A zero mean would render the sign unforecastable, as would constant volatility; hence the tradition in financial econometrics of removing unconditional means and working with zero-mean series disguises sign forecastability. Notice also that a large volatility relative to the mean renders the sign nearly unpredictable. It is interesting to note that the key link between sign forecastability and volatility dynamics parallels the literature on optimal prediction under asymmetric loss. In sign forecasting, volatility dynamics interact with a nonzero mean to produce time variation in the probability of a positive return and hence sign forecastability. In forecasting under asymmetric loss, as in Christoffersen and Diebold (1996, 1997), volatility dynamics similarly produce time variation in the optimal point forecast of a series with a constant conditional mean.

Our setup above was intentionally simple, but it is easy to see that the results are maintained under a number of interesting variations. To take just one example (we discuss several in the online supplement), note that if returns are conditionally non-Gaussian (e.g., conditionally skewed), the result in equation (2) that volatility forecastability implies sign forecastability still holds.

2.2 Sign Dynamics do not Require Conditional Mean Dynamics

Forecasting market direction is of interest for active asset allocation, and a substantial body of evidence suggests that it *can* be done, as per the references given earlier. Successful directional forecasting implies that returns must be somehow dependent. When directional forecasting is found to be empirically successful, it is tempting to assert that it is driven by (perhaps subtle) nonlinear conditional mean dependence, which would be missed in standard analyses of (linear) dependence, such as those based on return autocorrelations.

The key insight is that although sign dynamics *could* be due to conditional mean dependence, they need not be. In particular, we have demonstrated that volatility dynamics produce sign dynamics, so that one should *expect* sign dynamics in asset returns, given the overwhelming evidence of volatility dynamics, even if returns display near conditional mean independence.

Moreover, as noted above, in a general equilibrium involving risk averse market participants observed prices will of course not in general evolve as martingales, which is to say that observed returns may in general display some conditional mean dynamics, due perhaps to time-varying risk-premia over the business cycle. Any such conditional mean dynamics may of course contribute to sign forecastability as well. In our analysis we intentionally assume the absence of conditional mean dynamics, for several reasons. First, as discussed earlier, the evidence suggests that conditional mean independence is a

reasonable empirical approximation to asset return dynamics, despite the fact that dynamic asset pricing theory predicts dependence between first and second moments. Second, we want to focus on the subtle and little-understood connection between sign dynamics and *volatility* dynamics. Third, standard expansions suggest that, even if conditional mean dynamics are operative, one may expect that conditional variance dynamics will dominate.⁵ To see this, consider a Taylor series expansion of $\mu_{t+1|t}/\sigma_{t+1|t}$ about the mean,

$$\mu_{t+1|t}/\sigma_{t+1|t} \approx \mu/\sigma + (\mu_{t+1|t} - \mu)\sigma^{-1} + (\sigma_{t+1|t}^{-1} - \sigma^{-1})\mu \quad (3)$$

Based upon this, consider the variance of $\mu_{t+1|t}/\sigma_{t+1|t}$. To first order, the variance of $\mu_{t+1|t}$ is multiplied by σ^{-2} , which is a large number, while the variance of $1/\sigma_{t+1|t}$ is multiplied by μ^2 , a smaller number. However, the variance of $1/\sigma_{t+1|t}$ is huge and should be the dominant term.

2.3 An Intriguing Decomposition

It is interesting to interpret the phenomena at hand through the decomposition,⁶

$$R_{t+1} = \text{sign}(R_{t+1}) \cdot |R_{t+1}|. \quad (4)$$

In the simple models just described (and, to a good approximation, in observed return data), both of the right-hand-side components of returns display persistent dynamics and hence are forecastable, yet the left-hand-side variable, returns themselves, are unforecastable. This is an example of a nonlinear “common feature,” in the terminology of Engle and Kozicki (1993): both signs of returns and absolute returns are conditional mean dependent and hence forecastable, yet their product can be conditional mean independent and hence unforecastable.

3. Measuring the Strength of Sign Forecastability

Here we examine a number of questions relevant to measuring sign forecastability. How, if at all, does the derivative of a sign forecast with respect to volatility vary as a function of volatility, and in what volatility region is the derivative largest? What is the correlation between sign forecasts and realizations, and how, if at all, is the correlation related to the volatility of sign forecasts?

3.1 The Responsiveness of Sign Forecasts to Volatility Changes

In the simple setup developed thus far, we achieve probability forecastability only from volatility dynamics. A key issue is how much the probability forecast changes when the volatility changes, and the obvious measure is the derivative,

⁵ We thank a referee for suggesting this, and for the Taylor series argument.

⁶ For an extensive and authoritative treatment of this and related decompositions in the context of ultra-high-frequency financial asset price movements, see Rydberg and Shephard (2003).

$$\mathfrak{R} = \frac{\partial \Pr_t(R_{t+1} > 0)}{\partial \sigma_{t+1|t}}, \quad (5)$$

where we choose the notation \mathfrak{R} for “responsiveness.”⁷ Immediately,

$$\mathfrak{R} = -f\left(\frac{-\mu}{\sigma_{t+1|t}}\right)\left(\frac{\mu}{\sigma_{t+1|t}^2}\right), \quad (6)$$

where $f(\bullet)$ is the pdf of standardized returns.

In Figure 2, we work in a Gaussian environment and plot \mathfrak{R} as a function of the information ratio, μ/σ . Notice that \mathfrak{R} is always negative (i.e., the probability of a positive return is always decreasing in the conditional standard deviation). Crucially, however, \mathfrak{R} is not monotone in μ/σ ; instead, it achieves a minimum at $\mu/\sigma = \sqrt{2} = 1.41$. This makes sense: for μ/σ close to zero, the conditional probability of a positive return can deviate little from $1/2$, and hence responsiveness is tiny. Similarly, for very large μ/σ , the conditional probability of a positive return can deviate little from 1, and hence responsiveness is again tiny. Intermediate values of μ/σ , however, can produce greater responsiveness. The maximal forecastability as measured by \mathfrak{R} that occurs when $\mu/\sigma \approx 1.41$ is rather high for an information ratio. The frequency with which we hit that “sweet spot” depends on the volatility of volatility, to which we shall return.

3.2 The Correlation Between Sign Forecasts and Realizations

To characterize the correlation between sign forecasts and realizations, first notice that

$$\text{Cov}(I_{t+1}, P_{t+1|t}) = E(I_{t+1} P_{t+1|t}) - E(I_{t+1}) E(P_{t+1|t}) = E(I_{t+1} P_{t+1|t}) - P^2, \quad (7)$$

where P is the unconditional probability of a positive return and I_{t+1} is the indicator variable of an ex-post realized positive return. Second, use the law of iterated expectations to get

$$E(I_{t+1} P_{t+1|t}) = E\left(E_t(I_{t+1} P_{t+1|t})\right) = E\left(E_t(I_{t+1}) P_{t+1|t}\right) = E(P_{t+1|t}^2). \quad (8)$$

Hence we have

$$\text{Cov}(I_{t+1}, P_{t+1|t}) = E(P_{t+1|t}^2) - P^2 = \text{Var}(P_{t+1|t}), \quad (9)$$

⁷ This is also known as the “marginal effect” in the binary response literature.

so the covariance between the forecast and the realization is equal to the variance of the forecast.⁸
 Converting to correlation, we can write

$$\text{Corr}(I_{t+1}, P_{t+1|t}) = \frac{\text{Var}(P_{t+1|t})}{\text{Std}(I_{t+1}) \text{Std}(P_{t+1|t})} = \frac{\text{Std}(P_{t+1|t})}{\sqrt{P(1-P)}}, \quad (10)$$

where $P = E(I_t)$. Notice that the correlation between sign forecasts and realizations depends only on the standard deviation of the forecast, which of course will depend on the particular return process at hand. Despite its generality, the correlation expression furnishes considerable insight. In particular, because the optimal probability forecast $P_{t+1|t}$ is driven entirely by the volatility $\sigma_{t+1|t}$, we have that $\text{Corr}(I_{t+1}, P_{t+1|t})$, which is proportional to $\text{Std}(P_{t+1|t})$, is therefore driven by the volatility of volatility.

4. Detecting Sign Forecastability

Now we examine questions relevant to the detection of sign forecastability. First, we examine the serial correlation structure of sign realizations, and we argue that is it not likely to be useful for identifying sign forecastability, because the optimal sign predictor turns out to be highly nonlinear. Hence linear diagnostic tools discard valuable information. Next, we study the efficacy of runs tests for detecting sign forecastability and show that they have parallel deficiencies. Finally, we examine the ability of traditional market timing tests to detect sign forecastability, and we argue that although they may have good power to detect sign forecastability arising from time-varying expected returns, they have little or no power to detect sign forecastability arising from variation in volatility (or higher-ordered conditional moments).

4.1 Serial Correlation of Signs

The basic insight is simply that

$$\text{Corr}(I_{t+1}, I_t) < \text{Corr}(P_{t+1|t}, I_{t+1}), \quad (11)$$

which follows immediately from the fact that the optimal time-t forecast of I_{t+1} , $P_{t+1|t}$, has a higher correlation with I_{t+1} than anything else observed at time t, including I_t .⁹ The inequality leads one to suspect – although of course it does not definitively prove – that the size of the sign autocorrelation is

⁸ Alternatively, and perhaps more intuitively, note that regression of I_{t+1} on $P_{t+1|t}$ yields a unit coefficient, by definition of I_{t+1} and $P_{t+1|t}$. That is, $\text{Cov}(I_{t+1}, P_{t+1|t}) / \text{Var}(P_{t+1|t}) = 1$, which implies that $\text{Cov}(I_{t+1}, P_{t+1|t}) = \text{Var}(P_{t+1|t})$.

⁹ The inequality can also be established by formally deriving the equality $\text{Corr}(I_{t+1}, I_t) = \text{Corr}(P_{t+1|t}, I_{t+1}) \text{Corr}(P_{t+1|t}, I_t)$.

likely to be much less than the maximum attainable, because I_t is not likely to be highly correlated with $P_{t+1|t}$ (which is driven directly by $\sigma_{t+1|t}$, not I_t). Hence we conjecture that the autocorrelations of the sign sequence may be small, even if sign predictability is high. In section 5, we shall provide a quantitative assessment of this effect in a realistically-calibrated simulation.

4.2 Runs Tests

Working with a 0-1 sign sequence naturally leads one to consider tests of sign forecastability based on the number of runs in the sign sequence, and there is a long tradition of doing so in empirical finance.¹⁰ We now consider whether runs tests are likely to be more useful for detecting sign dependence than the serial correlation coefficients discussed above.

A run is simply a string of consecutive zeros or ones. Hence the number of runs is the number of switches from 0 to 1, plus the number of switches from 1 to 0, plus 1,

$$N_{\text{runs}} = 1 + \sum_t (1 - I_t) I_{t+1} + \sum_t I_t (1 - I_{t+1}), \quad (12)$$

which can be written as

$$N_{\text{runs}} = 1 + 2\hat{P}T - 2\sum I_t I_{t+1}, \quad (13)$$

where $\hat{P} = \frac{1}{T} \sum I_t$. Solving the N_{runs} equation for $\sum I_t I_{t+1}$ and exploiting the fact that $I_t^2 = I_t$ yields the estimator,

$$\hat{\text{Corr}}(I_{t+1}, I_t) = \frac{\frac{1}{T} \sum I_t I_{t+1} - \hat{P}^2}{\hat{P}(1 - \hat{P})} = 1 - \frac{N_{\text{runs}} - 1}{2T\hat{P}(1 - \hat{P})}. \quad (14)$$

This expression makes clear that there is no information in the number of runs in a sign sequence that is not also in the first-order autocorrelation of the sequence, and conversely. Thus our earlier conjecture that sign predictability is unlikely to be found using autocorrelations of the sign sequence translates as well into a conjecture that it is unlikely to be found using runs tests. In general, tests that rely only on the sign sequence omit important information about volatility dynamics, which is potentially valuable for detecting sign predictability.

4.3 Market Timing Tests

Here we discuss some popular market timing tests and their relationship to return sign forecasts, and we argue that none of them are likely to be useful for capturing sign forecastability when it arises via volatility dynamics, as we have emphasized.

¹⁰ See, for example, Campbell, Lo and MacKinlay (1997, Chapter 2) for an overview.

The literature on market timing is intimately concerned with signs and sign forecasting. For example, Henriksson and Merton (1981) argue that $p_1 + p_2 - 1$ can be interpreted as the multiplier to apply to the value of perfect market timing, where p_1 is the probability of a correctly forecasted negative return and p_2 is the probability of a correctly forecasted positive return. Breen, Glosten, and Jagannathan (1989) show that in the regression,

$$I(R_{t+1} > 0) = a + b I(P_{t+1|t} > 0.5) + e_{t+1}, \quad (15)$$

we have that $b = p_1 + p_2 - 1$. Hence the absence of Henriksson-Merton (1981) market timing ability corresponds to $b=0$, which is easily tested using standard methods. Cumby and Modest (1987) examine the closely-related regression,

$$R_{t+1} = a + b I(P_{t+1|t} > 0.5) + e_{t+1}, \quad (16)$$

and similarly test the significance of b .¹¹

Note that all of the tests above process $P_{t+1|t}$ in a particular way: $P_{t+1|t}$ enters the them only through the indicator function $I(P_{t+1|t} > 0.5)$.¹² Hence, for example, a $P_{t+1|t}$ of 0.5001 is treated fundamentally differently from a $P_{t+1|t}$ of 0.4999, whereas a $P_{t+1|t}$ of 0.9999 is treated no differently from a $P_{t+1|t}$ of 0.5001. This is particularly unfortunate because, although volatility dynamics lead to sign forecastability (i.e., variation in $P_{t+1|t}$), the time-varying $P_{t+1|t}$ may well never drop below 0.5. Such is the case, for example, in the leading example of a fixed positive expected return with symmetric conditional density analyzed earlier, so that the tests discussed above would have no power to detect sign dependence. Hence the traditional market timing tests are best viewed as tests for sign dependence arising from variation in expected returns rather than from variation in volatility (or higher-ordered conditional moments).

5. Sign Forecasting for Various Data Frequencies and Forecast Horizons: A Simulation

Experiment

We have shown that return sign forecastability arises from the interaction of nonzero expected returns and volatility forecastability. As expected returns approach zero, or as volatility forecastability

¹¹ See also Merton (1981), Whitelaw (1997), and Busse (1999). Note, however, that this framework does not include the quadratic regression tests of Treynor and Mazuy (1966) and Ferson and Schadt (1996), which are used in mutual fund performance evaluation.

¹² The same is true for generalizations of the Henriksson-Merton (1981) test, such as Pesaran and Timmermann (1992).

approaches zero, sign forecastability approaches zero.¹³ Hence one does not expect strong sign forecastability for very high frequency returns such as daily, despite their high volatility forecastability, because expected daily returns are negligible. Similarly, one does not expect strong sign forecastability for very low frequency returns such as annual, despite the high expected returns, because annual return volatility forecastability is negligible. One might therefore conjecture that sign forecastability will be highest at some intermediate horizon between such very short and very long extremes. In this section, we evaluate this conjecture.

When analyzing sign dynamics at various horizons, one is quickly faced with the challenge that few discrete-time dynamic model with time-varying volatility are closed in distribution under increasing horizons.¹⁴ We therefore work with Heston's (1993) affine continuous-time stochastic volatility model, which is arguably the most widely applied modern univariate stochastic process for equity and equity index returns. In Heston's model, the conditional characteristic function is known for any horizon, so Fourier inversion techniques can be used to compute sign probability forecasts.

5.1 Simulation Design

The stochastic volatility model parsimoniously captures many of the stylized facts of asset returns, including skewness, leptokurtosis and volatility persistence, and its conditional density can be calculated easily at any forecast horizon. For all of these reasons – both substantive and methodological – it has become a standard benchmark in empirical asset pricing.¹⁵

The Heston (1993) stochastic volatility model is

¹³ Of course, higher-ordered conditional moment dynamics can also contribute to sign predictability, as we will discuss in detail in the online supplement.

¹⁴ For penetrating insight into the difficulties involved in the temporal aggregation of discrete-time volatility models, see Meddahi and Renault (2004), Meddahi (2001), Darolles, Gouriéroux, and Jasiak (2001), and Heston and Nandi (2000).

¹⁵ The stochastic volatility model has been estimated by Andersen, Benzoni and Lund (2002), Bakshi, Cao and Chen (1997), Benzoni (2002), Chernov, Gallant, Ghysels and Tauchen (2003), Chernov and Ghysels (2000), Eraker, Johannes, and Polson (2003), Jones (2003), and Pan (2002), among others. We intentionally work with a very simple version of the model, with a single volatility factor, no volatility jumps, and no volatility long memory. One could of course examine even richer models with multiple volatility factors and/or volatility jumps as in Alizadeh, Brandt and Diebold (2002), Bates (2000), Chernov, Gallant, Ghysels and Tauchen (2003), and Duffie, Pan and Singleton (2000), and long memory in volatility as in Andersen, Bollerslev, Diebold and Ebens (2001) and Andersen, Bollerslev, Diebold and Labys (2001, 2003).

$$\begin{aligned}
dS(t) &= \mu S dt + \sigma(t) S dz_1 \\
d\sigma^2(t) &= \kappa(\theta - \sigma^2(t)) dt + \eta \sigma(t) dz_2,
\end{aligned}
\tag{17}$$

where $S(t)$ is the asset price process and $\sigma^2(t)$ is the variance process, and where $\text{corr}(dz_1, dz_2) = \rho$. The expected instantaneous rate of return is μ , the long run variance is θ , the speed of variance adjustment is governed by κ , and the volatility of volatility is governed by η .

Using Ito's lemma, the stochastic volatility process can conveniently be written in terms of the log asset price, $x(t)$, as

$$\begin{aligned}
dx(t) &= (\mu - \sigma^2(t)/2) dt + \sigma(t) dz_1 \\
d\sigma^2(t) &= \kappa(\theta - \sigma^2(t)) dt + \eta \sigma(t) dz_2.
\end{aligned}
\tag{18}$$

Notice that although the instantaneous drift is simply a constant, the continuously compounded return has a slightly time-varying mean from the Ito transformation.

The probability of an increase in the asset price between time t and $t+\tau$, or equivalently, the probability of a positive return during $[t, t+\tau]$, can be calculated using the inverse characteristic function technique.¹⁶ In particular,

$$P_{t+\tau|t} \equiv \Pr(x(t+\tau) \geq x(t) \mid x(t)=x, \sigma^2(t)=\sigma^2) = \frac{1}{2} + \frac{1}{\pi} \int_0^{\infty} \text{Re} \left(\frac{\exp(-i\psi x) f(x, \sigma^2, \tau; \psi)}{i\psi} \right) d\psi, \tag{19}$$

where $f(x, \sigma^2, \tau; \psi)$ is the characteristic function for horizon τ , $i=\sqrt{-1}$, $\text{Re}(\bullet)$ takes the real part of a complex number, and the characteristic function is

$$f(x, \sigma^2, \tau; \psi) = \exp(C(\tau, \psi) + D(\tau, \psi) \sigma^2 + i\psi x), \tag{20}$$

where

¹⁶ As in the Gaussian case, computation of the sign probability requires numerical integration, but the well-behaved integrand renders the integration straightforward.

$$C(\tau, \psi) = \mu\psi i\tau + \frac{\kappa\theta}{\eta^2} \left((\kappa - \rho\eta\psi i + d)\tau - 2\log\left(\frac{1 - g \exp(d\tau)}{1 - g}\right) \right) \quad (21)$$

$$D(\tau, \psi) = \left(\frac{\kappa - \rho\eta\psi i + d}{\eta^2} \right) \left(\frac{1 - \exp(d\tau)}{1 - g \exp(d\tau)} \right) \quad (22)$$

$$g = \frac{\kappa - \rho\eta\psi i + d}{\kappa - \rho\eta\psi i - d} \quad (23)$$

and

$$d = \sqrt{(\rho\eta\psi i - \kappa)^2 + \eta^2(\psi i + \psi^2)}. \quad (24)$$

As we have shown, the structure of the Heston model makes certain key calculations tractable, which explains its popularity. We will now exploit that tractability to illustrate several important aspects of optimal sign forecasting, including the existence of a nontrivial optimal horizon for sign prediction. Before doing so, however, we note that the Heston model actually *limits* the possible amount of sign predictability. In particular, to ensure in discrete-time simulations that the variance process stays strictly positive almost surely in the Heston model, it is necessary to restrict the volatility of volatility (η), such that $\eta^2 \leq 2\theta\kappa$. This limits the amount of sign predictability that the Heston model can generate, because it requires the volatility of volatility to be small relative to the unconditional variance and the conditional variance persistence, whereas sign predictability is greatest when the volatility of volatility is high, when conditional variance persistence is high, and when the unconditional variance is low relative to the mean. Other stochastic volatility specifications, although less tractable mathematically, imply no such limits to sign predictability.¹⁷ In this sense, then, our simulation results below on sign predictability using the Heston model as data-generating process are, if anything, conservative relative to those that could be

¹⁷ The volatility of volatility restriction in the Heston model has been found to be restrictive in option valuation as well. Circumventing the restriction, Duffie, Pan and Singleton (2000) suggest a model with correlated jumps in returns and jumps in volatility. Bakshi and Cao (2003) find that the new model significantly improves on the Heston model when valuating individual equity options.

obtained using alternative models.

We simulate prices at 5-minute intervals and assume 24-hour trading with 250 trading days per year. For the purpose of sign prediction, we proceed by discarding the intra-day observations and take daily to be the highest frequency of interest. We calibrate the parameters to typical values estimated in the empirical literature. Our benchmark values are $\mu = 0.10$, $\kappa = 2$, $\theta = 0.015$, $\eta = 0.15$ and $\rho = -0.50$, which imply a daily mean of about 0.037%, a daily unconditional standard deviation of 0.77%, unconditional skewness of about -0.1, and unconditional excess kurtosis of about 1.¹⁸ The annualized mean reversion parameter $\kappa = 2$ implies a daily persistence of about $1 - 2/250 = .99$ in a standard GARCH(1,1) model. Notice also that the parameters satisfy the $\eta^2 \leq 2\theta\kappa$ condition.

5.2 Simulation Results

In Figure 3 we plot the sign forecasts from a typical sample path of the simulated process, computed using (19). We show daily, weekly, monthly, quarterly, semi-annual and annual conditional as well as unconditional sign probabilities. As we move from daily to annual returns, the volatility of the conditional sign probabilities first increases and then decreases. By (10), this supports our conjecture that sign predictability should increase and then decrease with horizon. In contrast, the unconditional probability of a positive return increases monotonically (and at a decreasing rate) with horizon.

In Figure 4 we focus more directly and thoroughly on our conjecture that sign dynamics will be most prevalent at intermediate frequencies; examining the correlation between sign forecasts and realizations as a function of horizon. Consider first the top line where $\mu = .10$ which is our benchmark parameter value. The correlation is quite low for the highest-frequency returns, then it increases, and then it tapers off again as we aggregate toward annual returns. The correlation is highest at horizons of approximately 2-3 months (corresponding to 40-60 trading days). Interestingly, then, despite the fact that sign predictability is driven by volatility predictability, which is highest at very high frequencies, the interaction between decreasing volatility predictability and increasing expected returns under temporal aggregation results in maximization of sign predictability at medium horizons.

Consider now the middle and lower lines in Figure 4 corresponding to $\mu = .05$ and $\mu = 0$ respectively, with all other parameters kept at their benchmark values. As expected, signs are less forecastable at all horizons for smaller μ ; the figure provides a precise quantitative characterization. Interestingly, some sign forecastability remains even when $\mu = 0$, due to the non-zero leverage effect, ρ ,

¹⁸ In the online supplement, we explore an alternative parameterization with no leverage, and the results are qualitatively similar.

interacting with the volatility dynamics.¹⁹

In Figure 5 we explore the effects of lower volatility persistence. The top line again corresponds to our benchmark parameter value, that is $\kappa=2$. The middle line shows the correlation between sign forecasts and realizations when $\kappa=5$ (corresponding to a daily volatility persistence of about .98) and the bottom line when $\kappa=10$ (corresponding to a daily volatility persistence of about .96), with all other parameters kept at their benchmark values. As expected, signs are less forecastable at all horizons for lower volatility persistence, and the figure again provides a precise quantitative characterization.

In Figure 6 we show another important result, also suggested but not conclusively established by our earlier analytic work: the simple autocorrelation of the sign realization is small compared to the correlation between the forecasted and realized signs, at all horizons, suggesting that attempts to detect, model, or forecast signs by simple linear autoregressive models are not likely to be fruitful. It appears that the nonlinear volatility dynamics that drive sign dynamics make the linear forecastability in signs quite small. In our setup, tomorrow's sign is linked to tomorrow's volatility, which is much more correlated with today's volatility than with today's sign.

Finally, in Figure 7, we take the numerical derivative of the conditional probability, $P_{t+\tau|t}$, in equation (19) with respect to the standard deviation $\sigma(t)$. We plot this derivative against the annualized information ratio defined as $\mu/\sigma_{t+\tau|t}$, where $\mu=0.1$ is the (annualized) mean of the instantaneous return drift and $\sigma_{t+\tau|t}$ is the annualized expected average standard deviation over the forecast horizon, τ .²⁰ We use the benchmark parameters from Figure 6, with a forecast horizon of forty days ($\tau = 40/250$). Figure 7 thus confirms in a richer environment the result found in the stylized Gaussian model in Figure 2. In particular, in a dynamic volatility setting when volatility is very low, the sign probability does not respond to changes in volatility – it is simply 1. Similarly, when the volatility is very high, it dominates the mean return, and the sign probability remains near 1/2. However, for intermediate levels of volatility, the sign probability responds to changes in volatility. Specifically, in this case, when $\mu/\sigma_{t+\tau|t}$ is close to 1.5, that is, when $\sigma_{t+\tau|t}$ is close to 6.67% in annual terms, the sign predictor reaches a maximum level of

¹⁹ It is interesting to note that, even when both drift and leverage are zero, there is a small degree of sign forecastability as the horizon increases, coming from the Ito term, $-\sigma^2(t)/2$, in the drift of the log price process. Meddahi and Renault (2004) find similar effects arising under aggregation. By defining $dS(t)/S(t)$ to have a symmetric distribution, an asymmetry is automatically introduced in the distribution of $S(t+\tau)$. Please see the online supplement for further detail.

²⁰ Note that the expected average variance in the Heston model can be calculated from

$$\sigma_{t+\tau|t}^2 = \theta + (1 - \exp(-\kappa\tau))/(\kappa\tau)(\sigma^2(t) - \theta).$$

responsiveness to volatility changes. Figure 7 thus shows that the “sweet spot” for sign forecast responsiveness to volatility changes is found at a low but not completely unrealistic level of volatility.

6. Empirics

We now present a simple empirical example, producing daily sign forecasts of returns on the S&P 500 index (SPINDEX, from CRSP) from January 1, 1963 through December 31, 2003, at horizons ranging from $h=1$ through $h=250$ days.²¹ The example illustrates our methods, provides preliminary evidence as to their applied relevance, and lays the groundwork for additional exploration.

Let $R_{t+1:t+h}$ be the h -day return, and define the “positive return” indicator as $I_{t+h} = 1$ if $R_{t+1:t+h} > 0$ and $I_{t+h} = 0$ otherwise. We want to forecast I_{t+h} , and our earlier theoretical analysis strongly suggests using a model of the form

$$I_{t+h} = F\left(\frac{\mu}{\sigma_t}\right) + e_{t+h}, \quad (25)$$

where $F(\bullet)$ is a monotone function with a left limit of zero and a right limit of one, μ is the h -day expected return, and σ_t is a forecast of h -day return volatility. An obvious choice of $F(\bullet)$ is the logistic,

$$F(x) = \frac{\exp(x)}{1 + \exp(x)}, \quad (26)$$

which produces the popular logistic regression (logit) model.

Hence we proceed via logistic regression of I_{t+h} on $1/\sigma_t$. The issue remains as to what to use for σ_t ; indeed, as we have shown and stressed, the key ingredient of a sign forecast is a volatility forecast. Volatility forecasts can be obtained using a variety of approaches, including GARCH volatilities, stochastic volatilities, implied volatilities, and realized volatilities, as surveyed in Andersen, Bollerslev, Christoffersen and Diebold (2005a, b). Here we proceed using the popular RiskMetrics approach, which has links not only to the traditional GARCH approach, but also to the more recent “realized volatility” models of Hsieh (1991) and Andersen, Bollerslev, Diebold and Labys (2003), measuring any day’s variance (and forecasting any future day’s variance) as an exponentially weighted moving average of past

²¹ We do not include dividends, because SPINDEX does not. In addition, we do not subtract a risk-free rate. Doing so would require complete daily term structure data every day since 1963. It would of course be of interest to move to excess returns, including dividends, in subsequent more elaborate investigations.

daily squared returns, with the customary smoothing parameter of 0.94.²²

Each day we compute out-of-sample 1-day through 250-day return sign probability forecasts using five-year rolling estimation windows and different logit models for each horizon, to allow expected returns to change over time and horizon.²³ In Figure 8 we show the volatility series. Its well-known pronounced fluctuations and high persistence are immediately apparent. The key issue, of course, is whether the pronounced and predictable volatility dynamics translate empirically into similarly pronounced and predictable sign dynamics, as in the earlier simulation exercise.

The results are encouraging. In Figure 9 we show the conditional sign probability forecasts for the same six horizons as in the simulations, together with unconditional probabilities superimposed for visual reference. The unconditional sign probabilities naturally increase with horizon, and the conditional sign probabilities fluctuate widely and persistently around them – indeed noticeably more so than in the simulation example – reflecting the wide and persistent empirical fluctuations in S&P 500 volatility.

In Figure 10 we show the correlation between the out-of-sample sign probability forecast and the sign realization, as well as the first-order autocorrelation of the sign realization, for various horizons. The humped patterns, with the correlations larger than the autocorrelations, conform roughly to the theory. It is interesting to note, however, that both the correlations and the autocorrelations are generally much larger than in our earlier simulation example, which uses an affine volatility model with low volatility of volatility.

7. Concluding Remarks and Directions for Future Research

Our contribution is two-fold. First, we show that, given the widely-accepted volatility dynamics in stock returns, one should not be surprised to find sign forecastability in returns, even if returns are conditional mean independent. Hence sign dependence is likely widespread, but it is not necessarily indicative of time-varying expected returns and should not be interpreted as such.

Second, we explore a variety of aspects of sign forecastability, analytically, numerically and empirically. For example, we show analytically that sign probability forecasts are most sensitive to changes in volatility when volatility is at an intermediate level, and we show in a realistically calibrated simulation exercise that sign forecastability appears strongest at intermediate horizons of two or three months. In addition, we show that the nonlinear nature of sign dependence makes it unlikely to be found

²² In future work it would be of interest to contrast the results reported below with those produced using other volatility models. In particular, certain of our empirical results may be driven by the fact that RiskMetrics volatility forecasts do not mean revert, whereas those from most other models do.

²³ All told, we estimate roughly 100,000 logit models.

via traditional approaches such as analysis of sign autocorrelations, runs tests, or market timing tests.

Importantly, our results relate not only to academic concerns, but also to those of practitioners, who commonly use market timing strategies linked to volatility movements, as documented for example in Rattray and Balasubramanian (2003) and Gross, Keh, Leventhal and Sarfati (2003). To the best of our knowledge, however, we are the first to provide a rigorous scientific investigation of the links between volatility dynamics and directional market movements.

As for future work, the obvious next step is to formulate trading strategies based on sign forecasts (e.g., by trading digital options) and to assess their efficacy in generating positive risk-adjusted excess returns. Interestingly, the analysis of this paper makes clear that such strategies should be related to “volatility timing” strategies, such as those of Fleming, Kirby and Ostdiek (2001, 2003) and Johannes, Polson and Stroud (2002), in which portfolio shares are dynamically adjusted based on forecasts of the variance-covariance matrix of the underlying assets. Sign/volatility dynamics could also be exploitable in safety-first problems, as in Roy (1952), Bawa (1978), and Hagigi and Kluger (1987). One could progress significantly, moreover, by incorporating skewness and kurtosis dynamics, using the methods of El Babsiri and Zakoian (2001). Such “moment timing” will require evaluation measures more widely applicable than simple Sharpe ratios, such as Stutzer’s (2001) portfolio performance index.

Other explorations may also prove interesting. One example is generation of probability forecasts for future returns exceeding any given value x or percentile α , and developing the links to related work such as Engle and Manganelli (2002) and Taylor (2005). Another example is estimating nonparametrically the maximally predictable nonlinear function $f(\bullet)$ of returns. The results of Chen, Hansen and Scheinkman (2000) indicate that the optimal function will be monotone, as is the sign function, but the issue remains open as to whether and how the maximally predictable nonlinear function of returns diverges from the sign function.

References

- Abel, Andrew, 1988, Stock prices under time-varying dividend risk, *Journal of Monetary Economics* 22, 375-393.
- Alizadeh, Sasan, Michael W. Brandt and Francis X. Diebold, 2002, Range-based estimation of stochastic volatility models, *Journal of Finance* 57, 1047-1092.
- Andersen, Torben G., Luca Benzoni, and Jesper Lund, 2002, Estimating jump-diffusions for equity returns, *Journal of Finance* 57, 1239-1284.
- Andersen, Torben G., Tim Bollerslev, Peter F. Christoffersen, and Francis X. Diebold, 2005a, Volatility and correlation forecasting, in G. Elliott, C.W.J. Granger, and Allan Timmermann (eds.), *Handbook of Economic Forecasting*. Amsterdam: North-Holland, forthcoming.
- Andersen, Torben G., Tim Bollerslev, Peter F. Christoffersen, and Francis X. Diebold, 2005b, Practical volatility and correlation modeling for financial market risk management, in M. Carey and R. Stulz (eds.), *Risks of Financial Institutions*, University of Chicago Press for NBER, forthcoming.
- Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold, 2005, Parametric and nonparametric volatility measurement, in L.P. Hansen and Y. Ait-Sahalia (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland, forthcoming.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens, 2001, The distribution of realized stock return volatility, *Journal of Financial Economics* 61, 43-76.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, 2001, The distribution of realized exchange rate volatility, *Journal of the American Statistical Association* 96, 42-55.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, 2003, Modeling and forecasting realized volatility, *Econometrica* 71, 579-626.
- Backus, David and Alan Gregory, 1993, The relations between risk premia and conditional variances, *Journal of Business and Economic Statistics*, 11, 177-185.
- Baillie, Richard T., and Ramon P. DeGennaro, 1990, Stock returns and volatility, *Journal of Financial and Quantitative Analysis* 25, 203-214.
- Bakshi, Gurdip, and Charles Cao, 2003, Risk-neutral kurtosis, jumps and option pricing: evidence from 100 most actively traded firms on the CBOE, Manuscript, University of Maryland.
- Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical performance of alternative option pricing models, *Journal of Finance* 52, 2003-2049.
- Bates, David, 2000, Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* 94, 181-238.
- Bawa, Vijay S., 1978, Safety-first, stochastic dominance, and optimal portfolio choice, *Journal of Financial and Quantitative Analysis* 13, 255-271.

- Benzoni, Luca, 2002, Pricing options under stochastic volatility: an empirical investigation, Manuscript, University of Minnesota.
- Bollerslev, Tim, Ray Y. Chou, and Kenneth F. Kroner, 1992, ARCH modeling in finance: a selective review of the theory and empirical evidence, *Journal of Econometrics* 52, 5-59.
- Brandt, Michael W., and Qiang Kang, 2004, On the relationship between the conditional mean and volatility of stock returns: a latent VAR approach, *Journal of Financial Economics* 72, 217-257.
- Breen, William, Lawrence R. Glosten, and Ravi Jagannathan, 1989, Economic significance of predictable variations in stock index returns, *Journal of Finance* 44, 1177-1189.
- Brock, William A., and Cars H. Hommes, 1997, A rational route to randomness, *Econometrica* 65, 1059-1095.
- Busse, Jeffrey A., 1999, Volatility timing in mutual funds: evidence from daily returns, *Review of Financial Studies* 12, 1009-1041.
- Campbell, John, 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373-399.
- Campbell, John, and Ludger Hentschel, 1992, No news is good news: An asymmetric model of changing volatility in stock returns, *Journal of Financial Economics* 31, 281-318.
- Campbell, John, Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets*. Princeton: Princeton University Press.
- Campbell, John, and Robert Shiller, 1988, Stock prices, earnings and expected dividends, *Journal of Finance* 46, 661-676.
- Chen, Xiaohong, Lars Peter Hansen, and José A. Scheinkman, 2000, Principal Components and the Long Run, Manuscript, New York University, University of Chicago, and Princeton University.
- Chernov, Mikhail, A. Ronald Gallant, Eric Ghysels, and George Tauchen, 2003, Alternative models for stock price dynamics, *Journal of Econometrics* 116, 225-257.
- Chernov, Mikhail, and Eric Ghysels, 2000, A Study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of option valuation, *Journal of Financial Economics* 56, 407-458.
- Cheung, Yin-Wong, Menzie D. Chinn, and Antonio G. Pascual, 2003, Empirical exchange rate models of the nineties: are any fit to survive?, NBER Working Paper 9393.
- Christoffersen, Peter F., and Francis X. Diebold, 1996, Further results on forecasting and model selection under asymmetric loss, *Journal of Applied Econometrics* 11, 561-572.
- Christoffersen, Peter F., and Francis X. Diebold, 1997, Optimal prediction under asymmetric loss, *Econometric Theory* 13, 808-817.
- Cumby, Robert E., and David M. Modest, 1987, Testing for market timing ability: a framework for forecast evaluation, *Journal of Financial Economics* 19, 169-89.


- Darolles, Serge, Christian Gouriéroux, and Joanna Jasiak, 2001, Compound autoregressive models, Manuscript, York University.
- Diebold, Francis X., and Jose A. Lopez, 1996, Forecast evaluation and combination, in G.S. Maddala and C.R. Rao (eds.), *Handbook of Statistics*. Amsterdam: North-Holland, 241-268.
- Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343-1376.
- El Babsiri, Mohamed, and Jean-Michel Zakoian, 2001, Contemporaneous asymmetry in GARCH processes, *Journal of Econometrics* 101, 257-294.
- Elliott, Graham, and Takatoshi Ito, 1999, Heterogeneous expectations and tests of efficiency in the Yen/Dollar forward foreign exchange market, *Journal of Monetary Economics* 43, 435-456.
- Engle, Robert F., and Sharon Kozicki, 1993, Testing for common features, *Journal of Business and Economic Statistics* 11, 369-380.
- Engle, Robert F. and Simone Manganelli, 2002, CAViaR: conditional autoregressive value at risk by regression quantiles, *Journal of Business and Economic Statistics* 11, 167-176.
- Eraker, Bjorn, Michael Johannes, and Nicholas Polson, 2003, The role of jumps in returns and volatility, *Journal of Finance* 58, 1269-1300.
- Fama, Eugene F., 1970, Efficient capital markets: a review of theory and empirical work, *Journal of Finance* 25, 383-417.
- Fama, Eugene F., 1991, Efficient capital markets II, *Journal of Finance* 46, 1575-1617.
- Fama, Eugene, and Kenneth French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 19, 3-29.
- Fama, Eugene, and Kenneth French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23-49.
- Ferson, Wayne E., and Campbell R. Harvey, 1991, The variation of economic risk premiums, *Journal of Political Economy* 99, 385-415.
- Ferson, Wayne E., and Rudi W. Schadt, 1996, Measuring fund strategy and performance in changing economic conditions, *Journal of Finance* 51, 425-462.
- Fleming, Jeff, Chris Kirby, and Barbara Ostdiek, 2001, The economic value of volatility timing, *Journal of Finance* 56, 329-352.
- Fleming, Jeff, Chris Kirby, and Barbara Ostdiek, 2003, The economic value of volatility timing using realized volatility, *Journal of Financial Economics* 67, 473-509.
- Fontnouvelle, Patrick de, 2000, Information dynamics in financial markets, *Macroeconomic Dynamics* 4, 139-169.

- Foster, Douglass, Tom Smith, and Robert Whaley, 1997, Assessing goodness-of-fit of asset pricing models: the distribution of the maximal R^2 , *Journal of Finance* 52, 591-607.
- Franses, Philip Hans, and Dick van Dijk, 2000, *Nonlinear Time Series Models in Empirical Finance*. Cambridge: Cambridge University Press.
- French, Kenneth R., William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.
- Gencay, Ramazan, 1998, Optimization of technical trading strategies and the profitability in security markets, *Economics Letters* 59, 249-254.
- Genotte, Gerard, and Terry A. Marsh, 1993, Variations in economic uncertainty and risk premiums on capital assets, *European Economic Review* 37, 1021-1044.
- Ghysels, Eric, Andrew Harvey, and Eric Renault, 1996, Stochastic volatility, in G.S. Maddala and C.R. Rao (eds.), *Statistical Methods in Finance* (Handbook of Statistics, Volume 14). Amsterdam: North-Holland.
- Ghysels, Eric, Pedro Santa-Clara and Rossen Valkanov, 2005, There is a risk-return tradeoff after all, *Journal of Financial Economics* 76, 509-548.
- Glosten, Lawrence R., Ravi Jagannathan and David E. Runkle, 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance* 48, 1779–1801.
- Gross, Leon, Joanna Keh, Sara Leventhal and Olivier Sarfati , 2003, Individual implied volatility predicts future stock movement, Research Report, Equity Derivative Sales, Citigroup.
- Hagigi, Moshe, and Brian Kluger, 1987, Safety first: an alternative performance measure, *Journal of Portfolio Management* 13, 34-40.
- Hansen, Lars Peter, 1982, Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029-1054.
- Hansen, Bruce E., 1994, Autoregressive conditional density estimation, *International Economic Review* 35, 705-730.
- Harvey, Campbell R., 2001, The specification of conditional expectations, *Journal of Empirical Finance* 8, 573–638.
- Harvey, Campbell R., and Akhtar Siddique, 2000, Conditional skewness in asset pricing tests, *Journal of Finance* 55, 1263-1296.
- Henriksson, Roy D, and Robert C. Merton, 1981, On market timing and investment performance II: statistical procedures for evaluating forecasting skills, *Journal of Business* 54, 513-33.
- Heston, Steven L., 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327-343.

- Heston, Steven L., and Saikat Nandi, 2000, A closed-form GARCH option valuation model, *Review of Financial Studies* 13, 585-625.
- Hsieh, David A., 1991, Chaos and nonlinear dynamics: application to financial markets, *Journal of Finance* 46, 1839-1877.
- Johannes, Michael, Nicholas Polson and Jon Stroud, 2002, Sequential optimal portfolio performance: market and volatility timing, Manuscript, Columbia University.
- Jones, Christopher, 2003, The dynamics of stochastic volatility: evidence from underlying and options markets, *Journal of Econometrics* 116, 181-224.
- Kirby, Chris, 1997, Measuring the predictable variation in stock and bond returns, *Review of Financial Studies* 10, 579-630.
- Kuan, Chung-Ming, and Tung Liu, 1995, Forecasting exchange rates using feed-forward and recurrent neural networks, *Journal of Applied Econometrics* 10, 347-64.
- Larsen, Glen A. Jr, and Gregory D. Wozniak, 1995, Market timing can work in the real world, *Journal of Portfolio Management* 21, 74-81.
- Leitch, Gordon, and J. Ernest Tanner, 1991, Economic forecast evaluation: profits versus the conventional error measures, *American Economic Review* 81, 580-590.
- Lettau, Martin, and Sydney Ludvigson, 2003, Measuring and modeling variation in the risk-return tradeoff, in Yacine Ait-Shalia and Lars Peter Hansen (eds.), *Handbook of Financial Econometrics*, forthcoming.
- Leung, Mark T., Hazem Daouk, and An-Sing Chen, 2000, Forecasting stock indices: a comparison of classification and level estimation models, *International Journal of Forecasting* 16, 173-190.
- Levich, Richard M., 2001, *International Financial Markets*, Second Edition. New York: McGraw-Hill.
- Lo, Andrew W., and A. Craig MacKinlay, 1999, *A Non-Random Walk Down Wall Street*. Princeton, NJ: Princeton University Press.
- Meddahi, Nour, 2001, An eigenfunction approach for volatility modeling, Manuscript, University of Montreal.
- Meddahi, Nour, and Eric Renault, 2004, Temporal aggregation of volatility models, *Journal of Econometrics* 119, 355-379.
- Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867-887.
- Merton, Robert C., 1981, On market timing and investment performance I: an equilibrium theory of value for market forecasts, *Journal of Business* 54, 363-406.
- Nelson, Charles and Myung Kim, 1993, Predictable stock returns: the role of small sample bias, *Journal of Finance* 48, 641-661.
- Nelson, Daniel B., 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* 59, 347-370.


- Pan, Jun, 2002, The jump-risk premia implicit in options: evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3-50.
- Pesaran, Hashem, and Allan G. Timmermann, 1992, A simple non-parametric test of predictive performance, *Journal of business and economic statistics* 10, 461-465.
- Pesaran, M. Hashem, and Allan G. Timmermann, 1995, Predictability of stock returns: robustness and economic significance, *Journal of Finance* 50, 1201-1228.
- Pesaran, M. Hashem, and Allan G. Timmermann, 2004, How costly is it to ignore breaks when forecasting the direction of a time series? *International Journal of Forecasting* 20, 411-425.
- Rattray, Sandy and Venkatesh Balasubramanian, 2003, The new VIX as a market signal – it still works, Research Report, Equity Derivatives Strategy, Goldman Sachs.
- Roy, A.D., 1952, Safety-first and the holding of assets, *Econometrica* 20, 431- 449.
- Rydberg, Tina H. and Neil Shephard, 2003, Dynamics of trade-by-trade price movements: decomposition and models, *Journal of Financial Econometrics* 1, 2-25.
- Stutzer, Michael, 2001, A portfolio performance index and its implications, Manuscript, University of Iowa.
- Taylor, J.W., 2005, Generating volatility forecasts from value at risk estimates, *Management Science* 51, 712-725.
- Treynor Jack, and Kay Mazuy, 1966, Can mutual funds outguess the market? *Harvard Business Review* 44, 131-136.
- Turner, Christopher M., Richard Startz, and Charles R. Nelson, 1989, A markov model of heteroskedasticity, risk, and learning in the stock market, *Journal of Financial Economics* 25, 3-22.
- Wagner, Jerry, Steve Shellans, and Richard Paul, 1992, Market timing works where it matters most: in the real world, *Journal of Portfolio Management* 18, 86-90.
- White, Halbert, 2000, A reality check for data snooping, *Econometrica* 68, 1097-1126.
- Whitelaw, Robert F., 1997, Time-varying Sharpe ratios and market timing, Manuscript, New York University.
- Womack, Kent L., 1996, Do brokerage analysts' recommendations have investment value?, *Journal of Finance* 51, 137-167.

Figure 1
The Dependence of Sign Probability on Volatility



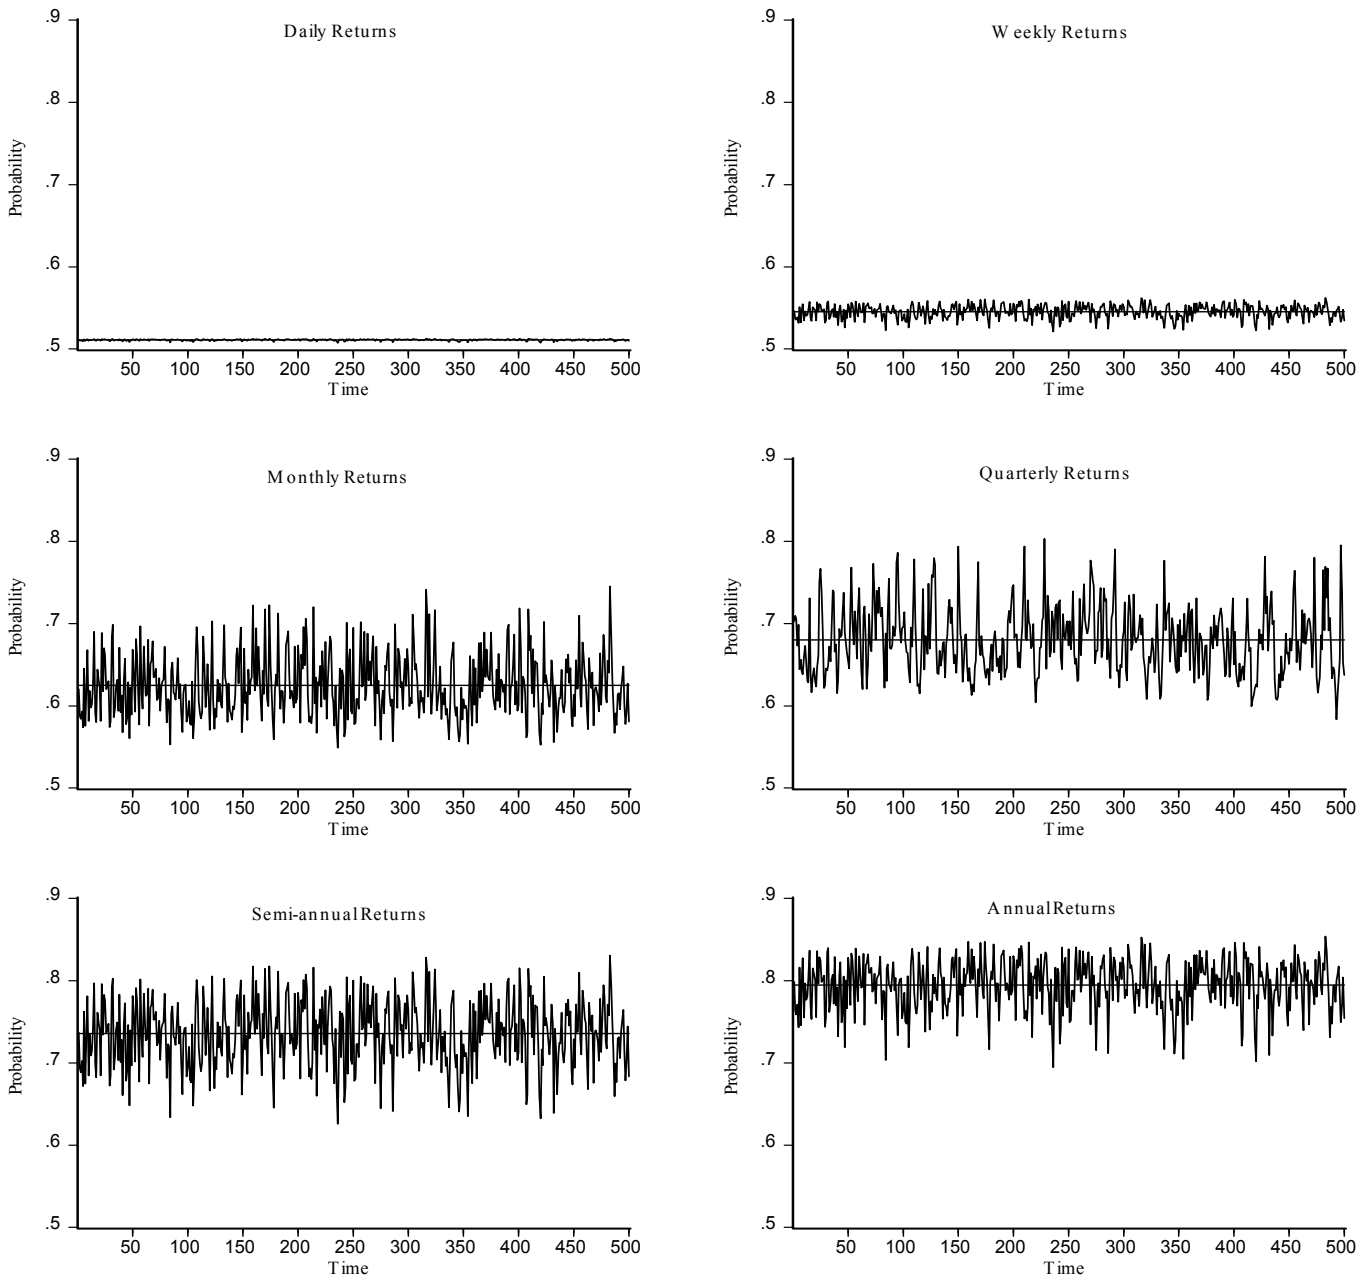
Notes: We show two Gaussian return densities, each with expected return of ten percent. The first return has a standard deviation of five percent and hence is positive with probability 0.98 (the area to the right of zero under the more peaked density function). The second return has a standard deviation of fifteen percent and hence is positive with smaller probability 0.75 (the area to the right of zero under the less peaked density function).

Figure 2
Responsiveness of Sign Probability to Volatility Movements
Plotted against the Information Ratio, μ / σ




Notes: We plot \mathfrak{R}_t , the derivative of the probability of a positive return with respect to return volatility, as a function of the information ratio, μ/σ . We assume Gaussian returns.

Figure 3
Time Series of Conditional Sign Probabilities, Various Return Horizons



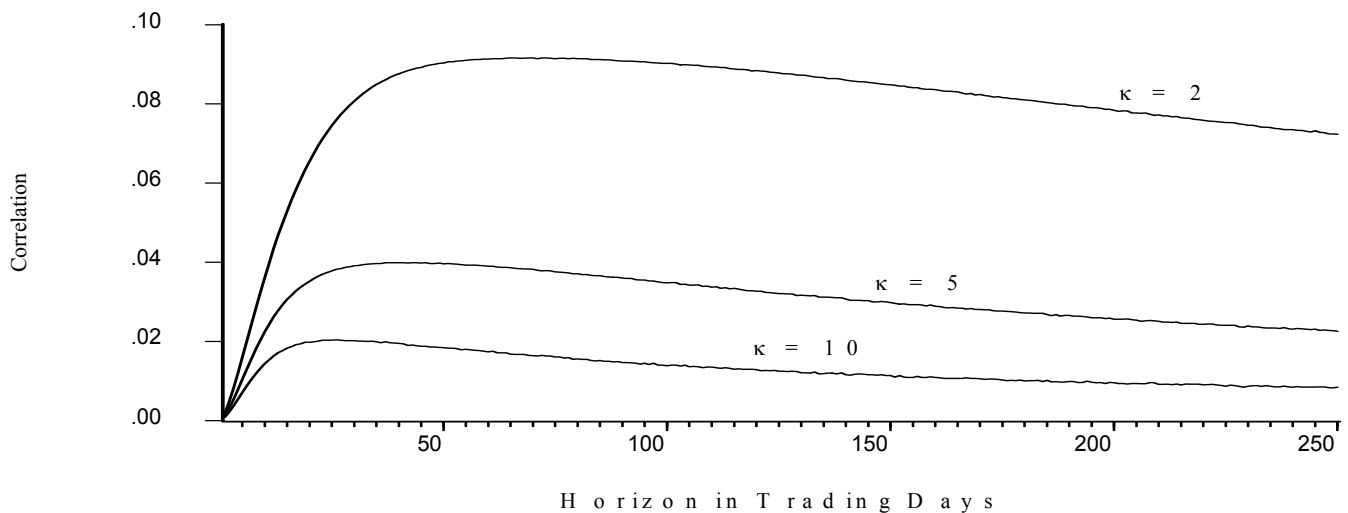
Notes: We simulate asset prices from a stochastic volatility model parameterized as $dS(t) = 0.10Sdt + \sigma(t)Sdz_1$ and $d\sigma^2(t) = 2(0.015 - \sigma^2(t))dt + 0.15\sigma(t)dz_2$, with $\text{corr}(dz_1, dz_2) = -0.5$. We then calculate, at each of 500 periods, the conditional probability of a positive return at daily, weekly, monthly and annual horizons. The horizontal line in each subplot denotes the unconditional probability of a positive return, which equals the average conditional probability.

Figure 4
Correlation Between Sign Forecasts and Realizations
Across Horizons, Various Expected Return Parameters



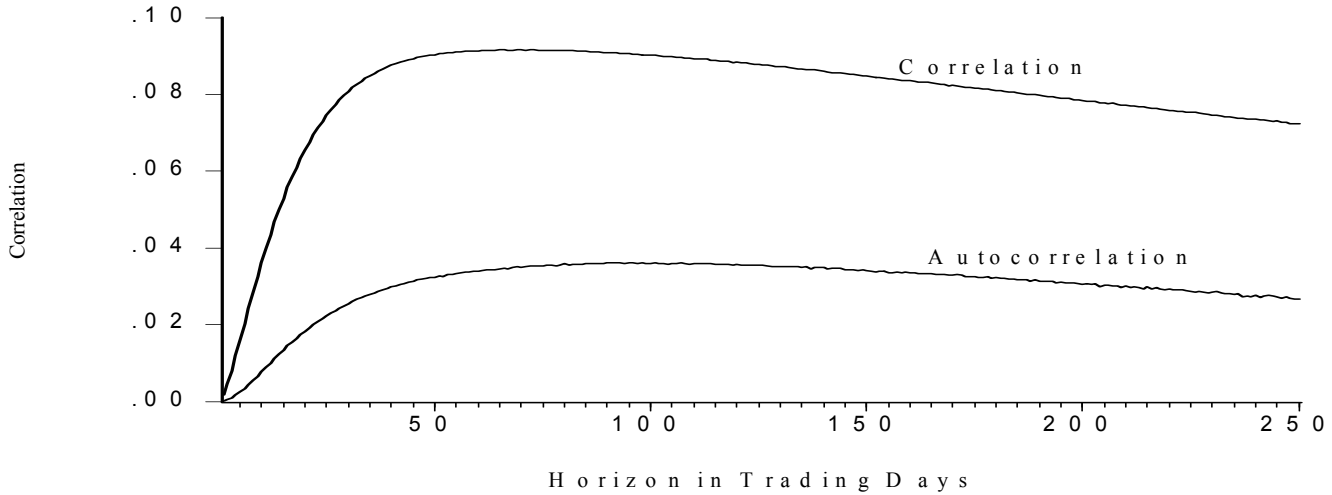
Notes: We simulate asset prices from a stochastic volatility model parameterized as $dS(t) = \mu S dt + \sigma(t) S dz_1$ and $d\sigma^2(t) = 2(0.015 - \sigma^2(t))dt + 0.15 \sigma(t) dz_2$, with $\text{corr}(dz_1, dz_2) = -0.5$ and $\mu = 0.10, 0.05$, and 0.00 respectively. We then calculate the ex ante conditional probability of a positive return as well as the ex post return sign realization at non-overlapping horizons ranging from one to 250 trading days (one year). We calculate the sample correlation between the forecast and the realization across a large number of realizations, making use of the quasi-analytic result (10) in the text.

Figure 5
Correlation between Sign Forecasts and Realizations
Across Horizons, Various Volatility Persistence Parameters




Notes: We simulate $dS(t) = 0.10S dt + \sigma(t) S dz_1$ and $d\sigma^2(t) = \kappa(0.015 - \sigma^2(t))dt + 0.15 \sigma(t) dz_2$, with $\text{corr}(dz_1, dz_2) = -0.5$, and with volatility persistence parameter $\kappa = 2, 5$ and 10 respectively. We then calculate the ex ante conditional probability of a positive return as well as the ex post return sign realization at non-overlapping horizons ranging from one to 250 trading days (one year). We calculate the sample correlation between the forecast and the realization using a large number of realizations, making use of the quasi-analytic result (10) in the text.

Figure 6
Correlation between Sign Forecasts and Realizations
and First Autocorrelation of Return Signs
Various Horizons, Benchmark Parameters




Notes: The correlation between sign forecasts and realizations is as in Figure 4. We compute the first autocorrelation of the sign sequence by simulating asset prices from a stochastic volatility model parameterized as $dS(t) = 0.10Sdt + \sigma(t)Sdz_1$ and $d\sigma^2(t) = 2(0.015 - \sigma^2(t))dt + 0.15\sigma(t)dz_2$, with $\text{corr}(dz_1, dz_2) = -0.5$. We then construct an indicator sequence of return signs for each horizon, and we calculate the sample autocorrelation from a long simulated sequence of returns, using the quasi-analytic result (10) in the text.

Figure 7
Responsiveness of Sign Probability to Volatility Movements in Heston's SV Model
Plotted against the Information Ratio, μ / σ



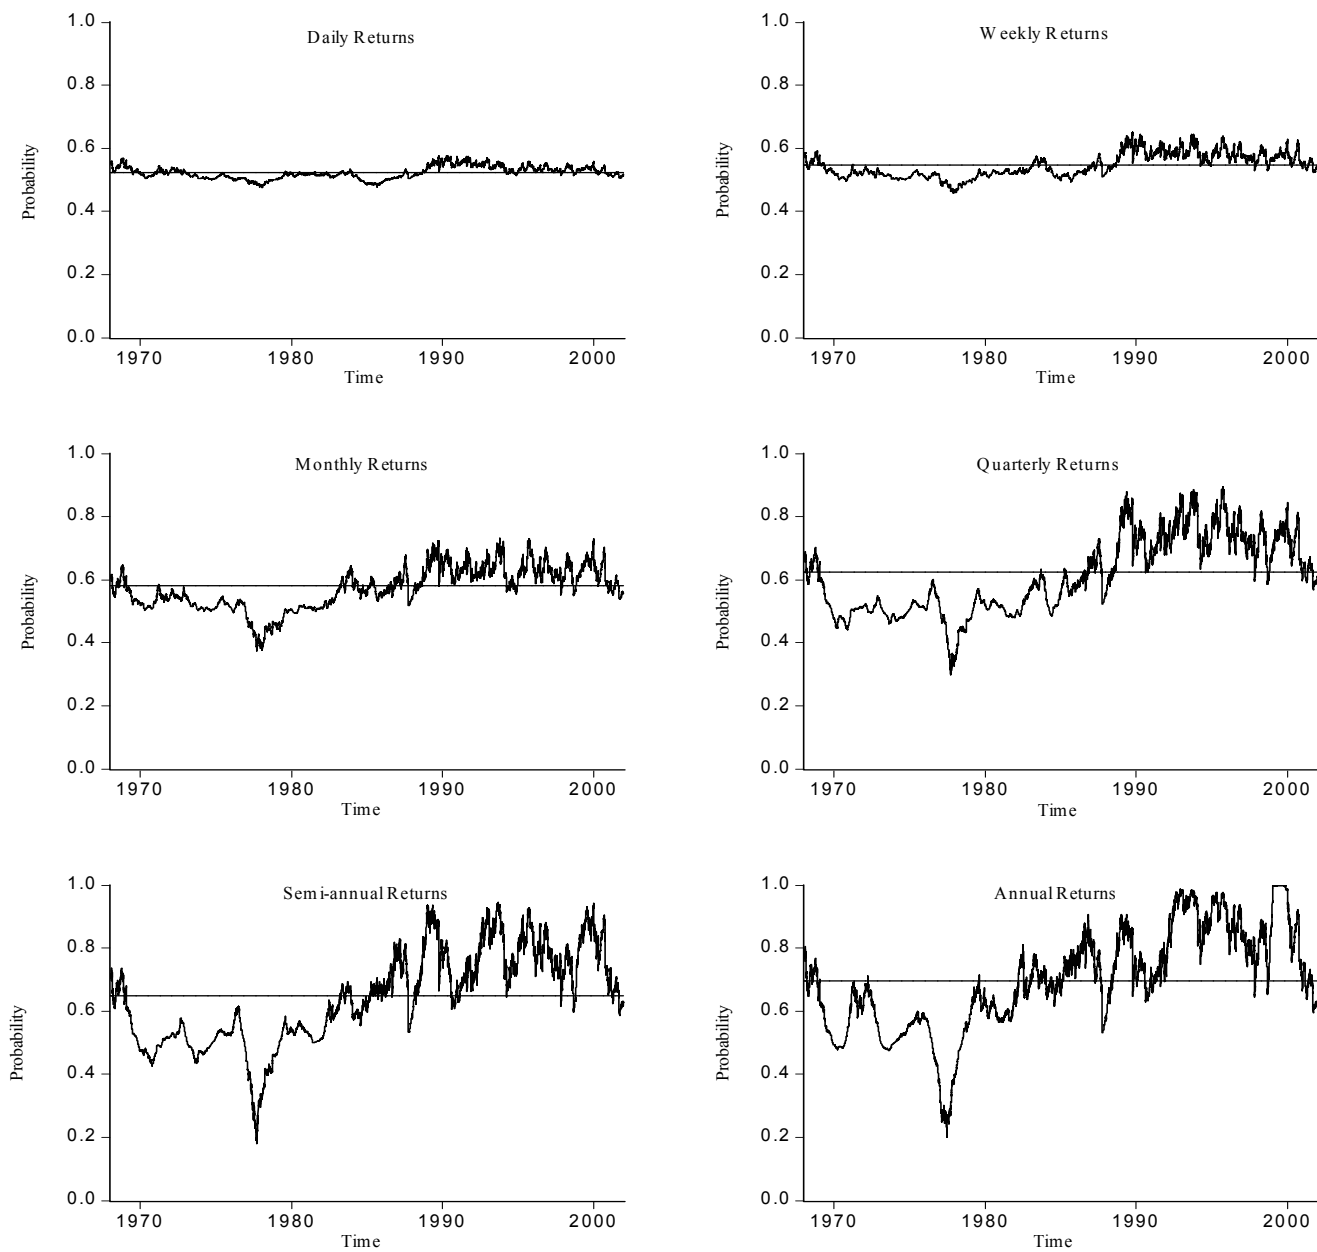
Notes: We take the numerical derivative of the conditional probability, $P_{t+\tau|t}$, from Heston's stochastic volatility model with respect to the standard deviation $\sigma(t)$. We plot this derivative against the annualized information ratio, $\mu/\sigma_{t+\tau|t}$, where μ is the annualized drift and $\sigma_{t+\tau|t}$ is the annualized expected average standard deviation over the forecast horizon, τ . We use the benchmark parameters from Figure 6, with a forecast horizon of forty trading days ($\tau = 40/250$).

Figure 8
Daily RiskMetrics Volatility (Annualized)



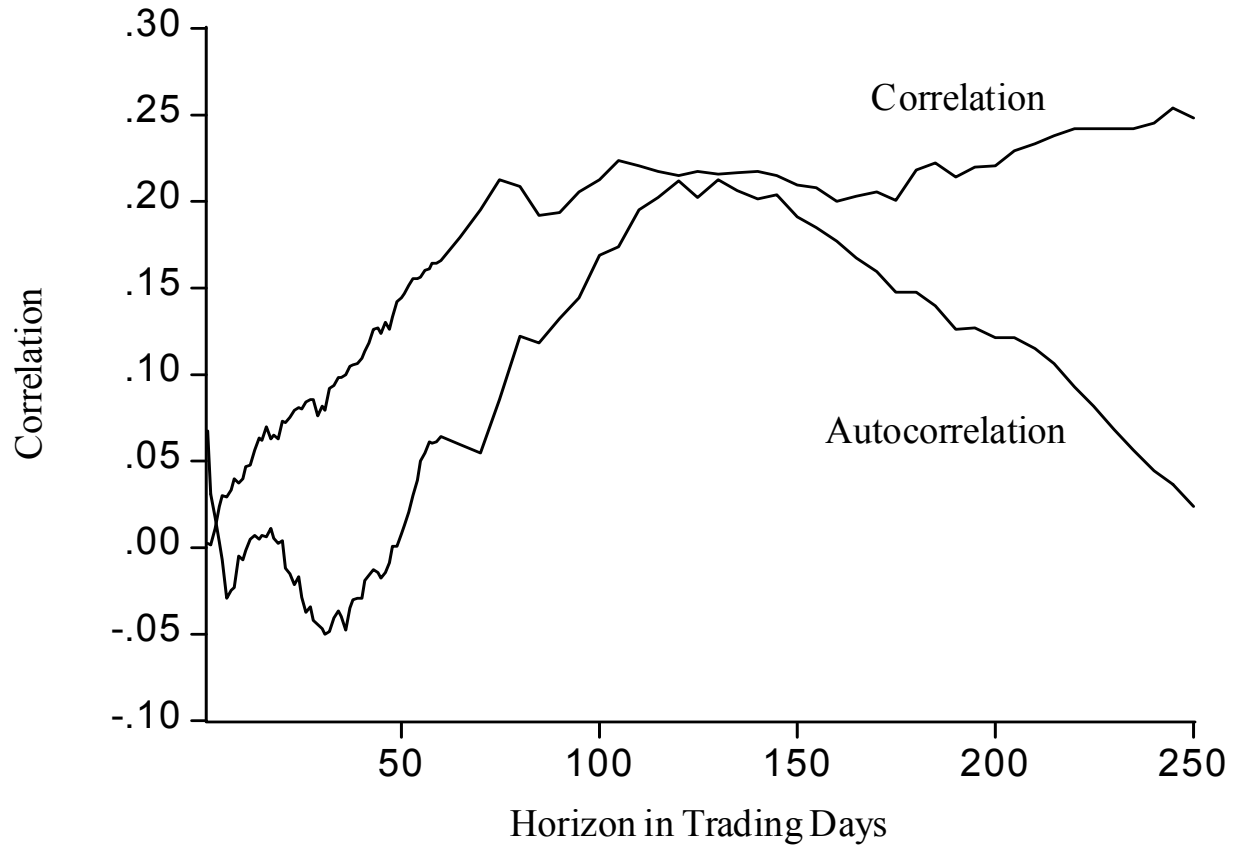
Notes: We show the time series of daily (annualized) RiskMetrics volatilities, from January 1, 1963 through December 31, 2003, obtained using the standard smoothing parameter of 0.94.

Figure 9
Conditional Probability Forecasts, Various Horizons



Notes: We show the time series of conditional probabilities of a positive S&P 500 returns at six horizons, assessed using a logit model in conjunction with a RiskMetrics volatility forecast.

Figure 10
Forecast Correlation and Return Sign Autocorrelation, Various Horizons



Notes: We show the empirical analog of Figure 6, based on probability forecasts produced using a logit model in conjunction with a RiskMetrics volatility forecast.

Financial Asset Returns, Direction-of-Change Forecasting, and Volatility Dynamics

Online Supplement

Peter F. Christoffersen
McGill University
and CIRANO

Francis X. Diebold
University of Pennsylvania
and NBER

In this online supplement we delve more deeply into both the theoretical and simulation results contained in the main text of the paper. On the theoretical side, we allow for potentially time-varying conditional skewness and kurtosis, establishing that, in general, sign predictability obtains even in the *absence* of drift and volatility predictability. On the simulation side, we present results for alternative drift and leverage specifications (i.e., no Ito drift term and/or no leverage). We begin with theory and then proceed to simulation.

1. Theoretical Results Allowing for Time-Varying Conditional Skewness and Kurtosis

In the main text we focused primarily on conditionally Gaussian asset return processes, for which sign prediction is driven by a time-varying conditional variance interacting with a constant, but nonzero, conditional mean. We also asserted, however, that the link between volatility forecastability and sign forecastability remains intact even in conditionally non-Gaussian environments. Here we defend that assertion by assessing the potential impact of time-varying higher-ordered conditional moments, the importance of which has been suggested for example by Hansen (1994) and Harvey and Siddique (2000).

Once the assumption of conditional normality is discarded, one is faced with choosing an acceptable alternative among countless possibilities. Instead of choosing a particular parametric distribution, we work with a Gram-Charlier expansion, which can be viewed as an approximation to any density with nontrivial higher-ordered moments. Define standardized returns as

$$z_{t+1} \equiv \frac{R_{t+1} - \mu}{\sigma_{t+1|t}}, \quad (\text{S1})$$

and assume they are approximately distributed according to the pdf

$$f_{t+1|t}(z) = \varphi(z) - \gamma_{3,t+1|t} \frac{1}{3!} D^3 \varphi(z) + \gamma_{4,t+1|t} \frac{1}{4!} D^4 \varphi(z), \quad (\text{S2})$$

where $\varphi(x)$ is the standard normal pdf, D^j is the j -th derivative, and the third and fourth conditional moments are defined in the usual way as

$$\gamma_{3,t+1|t} = \frac{E_t(R_{t+1} - \mu)^3}{\sigma_{t+1|t}^3} \quad (\text{S3})$$

and

$$\gamma_{4,t+1|t} = \frac{E_t(R_{t+1} - \mu)^4}{\sigma_{t+1|t}^4} - 3. \quad (\text{S4})$$

The conditional c.d.f. of standardized returns is then

$$F_{t+1|t}(z) = \Phi(z) - \gamma_{3,t+1|t} \frac{1}{3!} D^2 \varphi(z) + \gamma_{4,t+1|t} \frac{1}{4!} D^3 \varphi(z). \quad (\text{S5})$$

Notice that, because

$$\begin{aligned} D\varphi(z) &= -z\varphi(z) \\ D^2\varphi(z) &= (z^2-1)\varphi(z) \\ D^3\varphi(z) &= -(z^3-3z)\varphi(z), \end{aligned} \quad (\text{S6})$$

we have

$$F_{t+1|t}(z) = \Phi(z) - \varphi(z) \left(\frac{\gamma_{3,t+1|t}}{3!} (z^2-1) + \frac{\gamma_{4,t+1|t}}{4!} (z^3-3z) \right). \quad (\text{S7})$$

Because the probability of a positive return equals the probability of the standardized return being bigger than $-\mu/\sigma_{t+1|t}$, the probability of a positive return is

$$\begin{aligned} 1 - F_{t+1|t}(-\mu/\sigma_{t+1|t}) &= 1 - \Phi(-\mu/\sigma_{t+1|t}) \\ &+ \varphi(-\mu/\sigma_{t+1|t}) \left(\frac{\gamma_{3,t+1|t}}{3!} (\mu^2/\sigma_{t+1|t}^2 - 1) + \frac{\gamma_{4,t+1|t}}{4!} (-\mu^3/\sigma_{t+1|t}^3 + 3\mu/\sigma_{t+1|t}) \right), \end{aligned} \quad (\text{S8})$$

which is time-varying for nonzero μ , just as it was under conditional normality, so long as one or more of the second, third and fourth conditional moments are time-varying.

Importantly, sign forecastability arises even in the absence of conditional variance dynamics, in which case the probability of a positive return is

$$1 - F_{t+1|t}(-\mu/\sigma) = 1 - \Phi(-\mu/\sigma) + \varphi(-\mu/\sigma) \left(\frac{\gamma_{3,t+1|t}}{3!} (\mu^2/\sigma^2 - 1) + \frac{\gamma_{4,t+1|t}}{4!} (-\mu^3/\sigma^3 + 3\mu/\sigma) \right), \quad (\text{S9})$$

which is still time-varying for nonzero μ so long as either the third or fourth conditional moment is time-varying. Moreover, even if μ is zero, and regardless of whether volatility dynamics are present, sign forecastability arises so long as conditional skewness dynamics are present. In that case, the probability of a positive return becomes

$$1 - F_{t+1|t}(0) = 1/2 - \frac{\gamma_{3,t+1|t}}{\sqrt{2\pi} 3!} \approx 1/2 - 0.0665 \gamma_{3,t+1|t}, \quad (\text{S10})$$

which is time-varying as long as conditional skewness dynamics are present. Notice that a negative skewness implies that the probability of a positive return is greater than 0.5.


2. Simulation Results Under Alternative Drift and Leverage Specifications

In the simulation exercise of the main text we use an instantaneous log-price drift with a non-zero (indeed time-varying) Ito term, $\mu - \sigma^2(t)/2$, as well as a non-zero leverage correlation, ρ . We find that simulation to be a representative and realistic benchmark, which is why we adopted it. It is, however, of interest to examine variations on that benchmark, to assess robustness and gain deeper understanding.

In particular, we want to be sure that the simulation results of the main text are not driven purely by the time-varying Ito term and/or leverage effect. Hence we now explore the effects of dropping the Ito term and leverage from the log price diffusion, both jointly and in isolation. In each case we derive the differential equation for the new drift and solve it to get the corresponding conditional characteristic function, which yields the conditional sign probability via Fourier inversion. (For assistance we thank Chayawat Ornthanalai, an industrious Ph.D. student at McGill University.) Throughout, we fix the stochastic volatility persistence parameter at $\kappa = 2$.


The results appear graphically in Figures S1-S3, in which we plot the correlation between sign forecasts and realizations across horizons, for various specifications of diffusion drift and leverage. Although the correlations are generally smaller than those reported in the text, in each case the results are qualitatively identical: the interplay of non-zero drift and stochastic volatility produces sign predictability that varies with horizon, and maximal sign predictability occurs at moderate horizons of roughly one quarter. This basic qualitative result holds even when both the Ito effect and leverage are absent.

Figure S1
Correlation Between Sign Forecasts and Realizations
Across Horizons, Various Expected Return Parameters
Benchmark Drift Specification, $\rho=0$




Notes: We simulate asset prices from a stochastic volatility model parameterized as $dS(t) = \mu S dt + \sigma(t) S dz_1$ and $d\sigma^2(t) = 2(0.015 - \sigma^2(t))dt + 0.15\sigma(t)dz_2$, with $\text{corr}(dz_1, dz_2) = 0$ and $\mu = 0.10, 0.05$, and 0.00 respectively. We then calculate the ex ante conditional probability of a positive return as well as the ex post return sign realization at non-overlapping horizons ranging from one to 250 trading days (one year). We calculate the sample correlation between the forecast and the realization across a large number of realizations, making use of the quasi-analytic result (10) in the text.

Figure S2
Correlation Between Sign Forecasts and Realizations
Across Horizons, Various Expected Return Parameters
New Drift Specification, $\rho = -0.5$



Notes: We simulate asset prices from a stochastic volatility model parameterized as $dS(t) = (\mu + \sigma^2(t)/2)S dt + \sigma(t)S dz_1$ and $d\sigma^2(t) = 2(0.015 - \sigma^2(t))dt + 0.15\sigma(t)dz_2$, with $\text{corr}(dz_1, dz_2) = -0.5$ and $\mu = 0.10, 0.05,$ and 0.00 respectively. We then calculate the ex ante conditional probability of a positive return as well as the ex post return sign realization at non-overlapping horizons ranging from one to 250 trading days (one year). We calculate the sample correlation between the forecast and the realization across a large number of realizations, making use of the quasi-analytic result (10) in the text.

Figure S3
Correlation Between Sign Forecasts and Realizations
Across Horizons, Various Expected Return Parameters
New Drift Specification, $\rho=0$



Notes: We simulate asset prices from a stochastic volatility model parameterized as $dS(t) = (\mu + \sigma^2(t)/2)S dt + \sigma(t)S dz_1$ and $d\sigma^2(t) = 2(0.015 - \sigma^2(t))dt + 0.15\sigma(t)dz_2$, with $\text{corr}(dz_1, dz_2) = 0$ and $\mu = 0.10$ and 0.05 . We then calculate the ex ante conditional probability of a positive return as well as the ex post return sign realization at non-overlapping horizons ranging from one to 250 trading days (one year). We calculate the sample correlation between the forecast and the realization across a large number of realizations, making use of the quasi-analytic result (10) in the text.