

NBER WORKING PAPER SERIES

FINANCIAL ASSET RETURNS, DIRECTION-OF-CHANGE FORECASTING,
AND VOLATILITY DYNAMICS

Peter F. Christoffersen
Francis X. Diebold

Working Paper 10009
<http://www.nber.org/papers/w10009>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2003

This work is supported by the National Science Foundation, the Guggenheim Foundation, the Wharton Financial Institutions Center, FCAR, IFM2 and SSHRC. For helpful comments we thank David Bates, Antulio Bomfim, Michael Brandt, Michel Dacorogna, Graham Elliott, Rene Garcia, Christian Gourieroux, Clive Granger, Joanna Jasiak, Michael Johannes, Blake LeBaron, Bruce Lehman, Martin Lettau, Nour Meddahi, Sergei Sarkissian, Frank Schorfheide, Allan Timmermann, 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, and the Federal Reserve Bank of St. Louis. Sean Campbell and Clara Vega provided outstanding research assistance. All inadequacies are ours alone. The views expressed herein are those of the authors and are not necessarily those of the National Bureau of Economic Research.

©2003 by Peter F. Christoffersen and Francis X. Diebold. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Financial Asset Returns, Direction-of-Change Forecasting, and Volatility Dynamics
Peter F. Christoffersen and Francis X. Diebold
NBER Working Paper No. 10009
October 2003
JEL No. G1

ABSTRACT

We consider three sets of phenomena that feature prominently – and separately – 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) The standard finding of little or no conditional mean dependence is entirely consistent with a significant degree of sign dependence and volatility 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; (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) Sign dependence is very much present in actual U.S. equity returns, and its properties match closely our theoretical predictions; (f) 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.

Peter F. Christoffersen
McGill University
peter.christoffersen@mcgill.ca

Francis X. Diebold
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
and NBER
fdiebold@wharton.upenn.edu

1. Introduction

We consider three sets of phenomena that feature prominently – and separately – in the financial economics literature: conditional mean independence (and hence 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. Asset return forecasting is central to active asset allocation. Short-run return forecasting, however, is widely viewed as difficult, and perhaps even impossible. This view 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, conditional mean independence is reasonably viewed as a good working approximation to asset return dynamics. We do, however, emphasize the word “approximation,” as weak conditional mean dependence at very short horizons is well-known (see for example Lo and MacKinlay, 1999), typically emerging as weak serial correlation in ultra-high-frequency returns due to microstructure effects. Additionally, many believe that conditional-mean dependence is also operative at very long horizons (see for example Fama and French, 1988, 1989, and Campbell and Shiller, 1988), perhaps due to time-varying risk premia. This belief is, however, far from universally held, due to serious statistical complications including possibly spurious regressions (e.g., Kirby, 1997), data snooping (e.g., Foster, Smith and Whaley, 1997) and small sample biases (Nelson and Kim, 1993), which may grossly distort standard inference procedures when applied to long-horizon predictability regressions. However, the possible presence of very short-run or very long-run mean reversion is of little relevance to this paper, because (1) we focus on return horizons between one day and one year, and (2) our basic contention, namely that sign dependence exists and is interesting, would in general only be strengthened by conditional mean dependence.

Second, consider dependence and hence forecastability of market direction (the return *sign*). Profitable trading strategies may result if one is successful at forecasting 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.¹ Positive profits may be made when the sign forecast is correct.² This leads Levich, for example, to focus on what he terms *useful* forecasts, namely those that predict the direction of price change and hence “...lead to profitable speculative positions and correct hedging decisions.”

Interestingly, there is substantial evidence that sign forecasting can often be done with surprising success. Relevant literature includes 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 (1999), Elliott and Ito (1999), White (2000), Pesaran and Timmermann (2001), and Cheung, Chinn and Pascual (2003).

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. Note that strong conditional volatility dependence—in sharp contrast to conditional mean dependence—will not in general be reduced by market trading; hence the large amount of evidence of conditional volatility dependence is not at all controversial. 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, Andersen, Bollerslev and Diebold (2003) survey results from realized volatility modeling, and Franses and van Dijk (2000) survey results from regime-switching volatility models. 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 kurtosis model of Hansen (1994). The recent literature also contains intriguing theoretical work explaining the empirical phenomena, 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 entirely compatible, as occurs for example in a pure GARCH process. 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:

¹ Generalizations to multiple asset classes, such as stock and bond markets, involve basing allocation strategies on forecasts of the sign of the return spread.

² We say “may” because one must obviously adjust for interest and other costs.

- (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) The standard finding of little or no conditional mean dependence is entirely consistent with a significant degree of sign dependence and volatility dependence.
- (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 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.
- (f) Sign dependence is very much present in actual U.S. equity returns, and its properties match closely our theoretical predictions.

We derive the theoretical results in a general setting, and we illustrate them using a popular model of return dynamics.

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 some of its 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. In Section 6 we continue to extend our results, allowing for the possibility of unconditional and conditional skewness and kurtosis. In Section 7 we illustrate our methods and results by forecasting the 1-month-ahead direction of the S&P 100 index, using the VIX volatility index traded on the CBOE. We conclude in Section 8, offering concluding remarks and discussing directions for future work.

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 . 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 .³ 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

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

where $\Phi(\cdot)$ 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 probability of a positive return, as illustrated in Figure 1.

The surprising result that the sign of the return is forecastable even though the conditional mean

³ 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 $\Pr(R_{t+1} > 0 | \Omega_t) = E(I(R_{t+1}) > 0 | \Omega_t)$.

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. 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 Granger (1969) and 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, note that if returns are conditionally non-Gaussian (e.g., conditionally skewed), the result that volatility forecastability implies sign forecastability still holds, as we will explore in detail subsequently in Section 6. In the simple case at hand with all conditional moments constant except the conditional variance, but allowing for a non-Gaussian conditional density, we have that

$$\Pr_t(\mathbf{R}_{t+1} > 0) = 1 - \Pr_t(\mathbf{R}_{t+1} < 0) = 1 - \Pr\left(\frac{\mathbf{R}_{t+1} - \boldsymbol{\mu}}{\boldsymbol{\sigma}_{t+1|t}} < \frac{-\boldsymbol{\mu}}{\boldsymbol{\sigma}_{t+1|t}}\right) = 1 - \mathbf{F}\left(\frac{-\boldsymbol{\mu}}{\boldsymbol{\sigma}_{t+1|t}}\right), \quad (3)$$

where $\mathbf{F}(\bullet)$ denotes the relevant c.d.f. Hence market direction remains forecastable so long as the mean return is non-zero and volatility is dynamic.

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 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. Obviously such assertions are unfounded, as made clear by the above example of a process displaying both conditional mean independence and sign forecastability.

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 when returns are conditional mean independent. The upshot is that the standard finding of little or no conditional mean dynamics is entirely consistent with a significant degree of sign dynamics.

Moreover, 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, for example over the business cycle.⁴ 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 at least three reasons. First, as discussed earlier, a wealth of evidence suggests that conditional mean independence is a reasonable empirical approximation to asset return dynamics, despite the fact that theory does not require it.⁵ Second, we want to focus on the subtle and little-understood connection between sign dynamics and *volatility* dynamics. Third, we focus on fairly short horizons, certainly much shorter than the business cycle, ranging from one day to one year.

2.3 An Intriguing Decomposition

In closing this section, we note that it is interesting to interpret the phenomena at hand through the decomposition,

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

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 is conditional mean *independent* and hence unforecastable.

3. Measuring the Strength of Sign Forecastability

Here we examine a number of questions relevant to the 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

⁴ For a discussion of this and of return forecastability more generally, see, Balvers, Cosimano and McDonald (1990), Ferson and Harvey (1993), Glosten, Jaganathan, and Runkle (1993), Jegadeesh (1990), Jensen (1978), Mankiw, Romer, and Shapiro (1991), Patelis (1997), Sentana and Wadhvani (1991), and Sweeney (1986).

⁵ Moreover, recent evidence suggests that to the extent that conditional means and variances move jointly, they are *negatively* correlated, which would further increase sign predictability. See Lettau and Ludvigson (2003) and Brandt and Kang (2003).

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

Consider an obvious measure of sign probability responsiveness to changes in volatility,

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

where we choose the notation \mathfrak{R} for “responsiveness.”⁶ The motivation behind this measure is that in our simplest setup, we achieve probability forecastability only from volatility dynamics. A key issue is how much the probability forecast changes when the volatility changes. The \mathfrak{R} measure captures this. In general we have

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

where $f(\bullet)$ is the pdf of standardized returns. In the Gaussian case,

$$\mathfrak{R} = -\phi\left(\frac{\mu}{\sigma_{t+1|t}}\right)\left(\frac{\mu}{\sigma_{t+1|t}^2}\right), \quad (7)$$

where ϕ is a $N(0,1)$ density function.

In Figure 2, we work in a Gaussian situation and plot \mathfrak{R} as a function of volatility, keeping the mean at ten percent. 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 the standard deviation; instead, it achieves an interior optimum at $\sigma_{t+1|t} = \mu/\sqrt{2} \approx 0.071$. This makes sense: for tiny σ , the conditional probability can deviate little from 1, and hence responsiveness is tiny. Similarly, for huge σ , the conditional probability can deviate little from $1/2$, and hence responsiveness is again tiny. Intermediate values of σ , however, can produce greater responsiveness. Interestingly, Figure 2 indicates that forecastability as measured by \mathfrak{R} is maximized when σ is low (but not too low) relative to μ . With an expected return μ of 10 percent, $\sigma_{t+1|t} = 0.071$ implies quite a high conditional information ratio $\mu/\sigma_{t+1|t}$. 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

In the previous subsection we considered the issue of responsiveness of probability forecasts of

⁶ This is also known as the “marginal effect” in the binary response literature.

return signs to movements in the underlying volatility. If sign forecasts don't respond much to volatility movements, then we could not hope for close agreement between sign forecasts and realizations. This brings up the more general issue of how one might quantify such agreement (or lack thereof); hence in this subsection we consider aspects of the correlation between sign forecasts and realizations. This effectively amounts to something of an R^2 measure of sign forecastability.⁷

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

$$\text{Cov}(I_{t+1}, P_{t+1|t}) = E(I_{t+1} P_{t+1|t}) - E(I_{t+1}) E(P_{t+1|t}) = E(I_{t+1} P_{t+1|t}) - P^2, \quad (8)$$

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

$$E(I_{t+1} P_{t+1|t}) = E(E_t(I_{t+1} P_{t+1|t})) = E(E_t(I_{t+1}) P_{t+1|t}) = E(P_{t+1|t}^2). \quad (9)$$

Hence we have

$$\text{Cov}(I_{t+1}, P_{t+1|t}) = E(P_{t+1|t}^2) - P^2 = \text{Var}(P_{t+1|t}), \quad (10)$$

so the covariance between the forecast and the realization is equal to the variance of the forecast.⁸

Converting to correlation, we can write

$$\text{Corr}(I_{t+1}, P_{t+1|t}) = \frac{\text{Var}(P_{t+1|t})}{\text{Std}(I_{t+1}) \text{Std}(P_{t+1|t})} = \frac{\text{Std}(P_{t+1|t})}{\sqrt{P(1-P)}}. \quad (11)$$

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. In spite of its generality, the correlation expression furnishes considerable insight. A high standard deviation of the conditional probability forecast, which could arise from a high variance of the conditional variance, will increase the forecastability of the return sign.

4. Detecting Sign Forecastability

Now we examine questions relevant to the detection of sign forecastability. In particular, we examine the serial correlation structure of sign realizations, and we argue that is it not likely to be useful

⁷ For an interesting discussion of R^2 -type measures in binary regressions, see Estrella (1998).

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 $\text{Cov}(I_{t+1}, P_{t+1|t}) = \text{Var}(P_{t+1|t})$. We thank a referee for this insight.

for identifying sign forecastability. We also study the efficacy of runs tests and traditional market timing tests for the detection of sign forecastability.

4.1 Serial Correlation Coefficients

Here we consider whether sign forecastability is likely to be detected using the obvious and standard method: examination of the autocorrelations of the sign sequence. It turns out that the answer is no, because the nature of sign dependence induced by volatility dependence is highly nonlinear, so that standard measures of linear dependence are likely small even when sign forecastability is large.

Consider the first-order autocovariance of the indicator sequence. We can write

$$\text{Cov}(I_{t+1}, I_t) = E(I_{t+1}I_t) - E(I_{t+1})E(I_t) = E(I_{t+1}I_t) - P^2. \quad (12)$$

As before we can use the law of iterated expectations to get

$$E(I_{t+1}I_t) = E(E_t(I_{t+1}I_t)) = E(I_t E_t(I_{t+1})) = E(I_t P_{t+1|t}). \quad (13)$$

Hence

$$\text{Cov}(I_{t+1}, I_t) = E(P_{t+1|t}I_t) - P^2. \quad (14)$$

Converting now to autocorrelations, we have

$$\text{Corr}(I_{t+1}, I_t) = \frac{E(P_{t+1|t}I_t) - P^2}{\text{Std}(I_{t+1})\text{Std}(I_t)} = \frac{E(P_{t+1|t}I_t) - P^2}{P(1-P)}. \quad (15)$$

Although this expression again depends on the particular return process at hand, it is useful in calculating an upper bound on the autocorrelation. As an intermediate step, consider the correlation between today's sign realization and today's forecast of tomorrow's sign. We have

$$\text{Corr}(P_{t+1|t}, I_t) = \frac{E(P_{t+1|t}I_t) - E(P_{t+1|t})E(I_t)}{\text{Std}(P_{t+1|t})\text{Std}(I_t)} = \frac{E(P_{t+1|t}I_t) - P^2}{\text{Std}(P_{t+1|t})\sqrt{P(1-P)}}. \quad (16)$$

Substituting (16) into the expression for the autocorrelation of the indicator sequence (15), we get

$$\text{Corr}(I_{t+1}, I_t) = \frac{\text{Corr}(P_{t+1|t}, I_t)\text{Std}(P_{t+1|t})}{\sqrt{P(1-P)}}. \quad (17)$$

Substituting the expression for the correlation between the ex-ante sign forecast and the ex-post sign-realization (11), we have

$$\text{Corr}(I_{t+1}, I_t) = \text{Corr}(P_{t+1|t}, I_{t+1}) \text{Corr}(P_{t+1|t}, I_t). \quad (18)$$

If all three correlations in (18) are positive and bounded away from one (which is a realistic assumption), we obtain

$$\text{Corr}(I_{t+1}, I_t) < \text{Corr}(P_{t+1|t}, I_{t+1}). \quad (19)$$

Intuitively, the optimal time-t forecast, $P_{t+1|t}$, has a higher correlation with the time t+1 realization than anything else observed at time t, including the time t realization itself. The lower the correlation between today's sign and today's forecast of tomorrow's sign, the lower the autocorrelation in the sign. Conversely, notice that if the sign forecast were linear in the current realization, then the correlation between today's sign and today's forecast of tomorrow's sign would be one, and the autocorrelation would coincide with the correlation between the ex-ante predictor and ex-post realization. Thus it is the nonlinearity in the dynamic process of the indicator sequence that lowers the autocorrelation relative to the cross correlation between the ex-ante predictor and the ex-post realization.

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.⁹ 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,

$$N_{\text{runs}} = 1 + \sum_t (1 - I_t) I_{t+1} + \sum_t I_t (1 - I_{t+1}), \quad (20)$$

which can be written as

$$N_{\text{runs}} = 1 + 2\hat{P}T - 2\sum I_t I_{t+1}, \quad (21)$$

where we have used $\hat{P} = \frac{1}{T} \sum I_t$. Solving the N_{runs} equation for $\sum I_t I_{t+1}$ and substituting into the formula for $\text{Corr}(I_{t+1}, I_t)$ derived earlier, we obtain an estimator

$$\hat{\text{Corr}}(I_{t+1}, I_t) = \frac{\frac{1}{T} \sum I_t I_{t+1} - \hat{P}^2}{\hat{P}(1 - \hat{P})} = 1 - \frac{N_{\text{runs}} - 1}{2T\hat{P}(1 - \hat{P})}. \quad (22)$$

⁹ See, for example, Campbell, Lo and MacKinlay (1997, Chapter 2) for an excellent overview.

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. Thus our finding above that sign predictability is unlikely to be found using autocorrelations of the hit sequence implies that it is also unlikely to be found using runs tests. In general, tests that rely only on the hit 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.

The literature on market timing is intimately concerned with signs and sign forecasting. For example, Hendriksson and Merton (1981) argue that $\mathbf{p}_1 + \mathbf{p}_2 - 1$ can be interpreted as the value of market timing, where \mathbf{p}_1 is the probability of a correctly forecasted negative return and \mathbf{p}_2 is the probability of a correctly forecasted positive return. Breen, Glosten, and Jagannathan (1989) show that in the regression,

$$\mathbf{I}(\mathbf{R}_{t+1} > 0) = \mathbf{a} + \mathbf{b} \mathbf{I}(\mathbf{P}_{t+1|t} > 0.5) + \mathbf{e}_{t+1}, \quad (23)$$

we have that $\mathbf{b} = \mathbf{p}_1 + \mathbf{p}_2 - 1$. Hence the absence of Hendriksson-Merton (1981) market timing ability corresponds to $\mathbf{b} = 0$, which is easily tested using standard methods. Cumby and Modest (1987) examine the closely-related regression,

$$\mathbf{R}_{t+1} = \mathbf{a} + \mathbf{b} \mathbf{I}(\mathbf{P}_{t+1|t} > 0.5) + \mathbf{e}_{t+1}, \quad (24)$$

and similarly test the significance of \mathbf{b} .

Note that all of the tests above process $\mathbf{P}_{t+1,t}$ in a particular way: $\mathbf{P}_{t+1,t}$ enters the tests only through the indicator function $\mathbf{I}(\mathbf{P}_{t+1,t} > 0.5)$.¹⁰ Hence, for example, a $\mathbf{P}_{t+1,t}$ of 0.5001 is treated fundamentally differently from a $\mathbf{P}_{t+1,t}$ of 0.4999, whereas a $\mathbf{P}_{t+1,t}$ of 0.9999 is treated no differently from a $\mathbf{P}_{t+1,t}$ of 0.5001. This is particularly unfortunate because, although volatility dynamics lead to sign forecastability (i.e., variation in $\mathbf{P}_{t+1,t}$), the time-varying $\mathbf{P}_{t+1,t}$ 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 (and in the empirical example that we will present shortly), 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

¹⁰ The same is true for generalizations of the Hendriksson-Merton (1981) test, such as that of Pesaran and Timmermann (1992).

than from variation in volatility (or higher-ordered conditional moments).¹¹

5. Sign Forecasting for Various Data Frequencies and Forecast Horizons

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 approaches zero, sign forecastability approaches zero.¹² 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 virtually no discrete-time dynamic model with time-varying volatility is closed in distribution under increasing horizons.¹³ To circumvent this problem, we work with a continuous-time stochastic volatility model, in particular the convenient and popular model of Heston (1993). Working in continuous time has the additional benefit that temporal aggregation and forecasting at increasing horizons are interchangeable.

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. It has been estimated by Andersen, Benzoni and Lund (2000), Bakshi, Cao and Chen (1997), Benzoni (1999), Chernov, Gallant, Ghysels and Tauchen (2001), Chernov and Ghysels (2000), Eraker, Johannes, and Polson (2000), Jones (1999), and

¹¹ In addition to the papers already mentioned, see Merton (1981), Whitelaw (1997), and Busse (1999).

¹² Of course, higher-ordered conditional moment dynamics can also contribute to sign predictability, as we will discuss in detail subsequently.

¹³ For penetrating insight into the difficulties involved in the temporal aggregation of discrete-time volatility models, see Meddahi and Renault (2000), Meddahi (2001), Darolles, Gouriéroux, and Jasiak (2001), and Heston and Nandi (2000).

Pan (2000), among others.¹⁴

The Heston (1993) stochastic volatility model is

$$\begin{aligned}dS(t) &= \mu S dt + \sigma(t) S dz_1 \\d\sigma^2(t) &= \kappa(\theta - \sigma^2(t)) dt + \eta \sigma(t) dz_2,\end{aligned}\tag{25}$$

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

$$\begin{aligned}dx(t) &= (\mu - \sigma^2(t)/2) dt + \sigma(t) dz_1 \\d\sigma^2(t) &= \kappa(\theta - \sigma^2(t)) dt + \eta \sigma(t) dz_2.\end{aligned}\tag{26}$$

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.¹⁵ In particular,

¹⁴ 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 (2001), 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).

¹⁵ As in the Gaussian case, computation of the sign probability requires numerical integration, but the well-behaved integrand renders the integration straightforward.

$$P_{t+\tau|t} \equiv \Pr(\mathbf{x}(t+\tau) \geq \mathbf{x}(t) \mid \mathbf{x}(t)=\mathbf{x}, \sigma^2(t)=\sigma^2) = \frac{1}{2} + \frac{1}{\pi} \int_0^{\infty} \operatorname{Re} \left(\frac{\exp(-i\psi\mathbf{x}) f(\mathbf{x}, \sigma^2, \tau; \Psi)}{i\psi} \right) d\psi, \quad (27)$$

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

$$f(\mathbf{x}, \sigma^2, \tau; \Psi) = \exp(C(\tau, \Psi) + D(\tau, \Psi)\sigma^2 + i\psi\mathbf{x}), \quad (28)$$

where

$$C(\tau, \Psi) = \mu\psi i\tau + \frac{\kappa\theta}{\eta^2} \left((\kappa - \rho\eta\psi i + d)\tau - 2\log \left(\frac{1 - g \exp(d\tau)}{1 - g} \right) \right) \quad (29)$$

$$D(\tau, \Psi) = \left(\frac{\kappa - \rho\eta\psi i + d}{\eta^2} \right) \left(\frac{1 - \exp(d\tau)}{1 - g \exp(d\tau)} \right) \quad (30)$$

$$g = \frac{\kappa - \rho\eta\psi i + d}{\kappa - \rho\eta\psi i - d} \quad (31)$$

and

$$d = \sqrt{(\rho\eta\psi i - \kappa)^2 + \eta^2(\psi i + \psi^2)}. \quad (32)$$

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 key features of optimal sign forecasting, including the existence of a nontrivial optimal horizon for sign prediction. Before doing so, however, we note the apparently little-known fact that the Heston model actually *limits* the possible amount of sign predictability. In particular, to ensure that the continuous-time variance process stays

strictly positive almost surely in the Heston model, it is necessary to restrict the volatility of volatility (η), such that $\eta^2 \leq 2\theta\kappa$. Notice, however, that the restriction also 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.¹⁶ In this sense, then, it would seem that our simulation results below on the appearance of sign predictability using the Heston model as data-generating process are, if anything, conservative relative to those that could be 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. The annualized mean reversion parameter $\kappa = 2$ implies a daily persistence of about $1 - 2/250 = .992$ 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. We show daily, weekly, monthly 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 (11), 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. 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

¹⁶ 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.

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.

In Figure 5 we explore the effects of smaller expected return values. We show the correlation between the ex ante sign forecast and the ex post sign realization when $\mu=.05$ and when $\mu=0$, 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 , interacting with the volatility dynamics.¹⁷

In Figure 6 we explore the effects of lower volatility persistence. We again show the correlation between sign forecasts and realizations when $\kappa=10$ (corresponding to a daily volatility persistence of about .96) and when $\kappa=5$ (corresponding to a daily volatility persistence of about .98), 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 7 we show another important result, also suggested but not conclusively established by our earlier analytic work: the simple autocorrelation of the sign realization is relatively small at all horizons, in contrast to the correlation between the forecasted and realized signs, suggesting that attempts to detect, model, or forecasts signs by simple linear autoregressive models are not likely to be fruitful. It appears that the nonlinear volatility dynamics that drive sign dynamics render small the linear forecastability in signs. 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.

6. Allowing for Potentially Time-Varying Conditional Skewness and Kurtosis

¹⁷ 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 (2000) find similar leverage effects arising under aggregation even if the innovations are uncorrelated. By defining $dS(t)/S(t)$ to have a symmetric distribution, an asymmetry is automatically introduced in the distribution of $S(t+\tau)$. Alternatively, we could have assumed a stock price process of

$$dS(t) = (\mu + \sigma^2(t)/2)S dt + \sigma(t)S dz_1$$

$$d\sigma^2(t) = \kappa(\theta - \sigma^2(t))dt + \eta\sigma(t)dz_2,$$

which would have eliminated sign forecastability when $\mu=\rho=0$.

Thus far we have 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 potential 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 the Gram-Charlier expansion, which can be viewed as an approximation to any density with nontrivial higher-ordered moments. First define standardized returns as

$$z_{t+1} \equiv \frac{R_{t+1} - \mu}{\sigma_{t+1|t}}, \quad (33)$$

and assume they are approximately distributed according to the pdf

$$f_{t+1|t}(z) = \phi(z) - \gamma_{3,t+1|t} \frac{1}{3!} D^3 \phi(z) + \gamma_{4,t+1|t} \frac{1}{4!} D^4 \phi(z), \quad (34)$$

where $\phi(x)$ 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

$$\gamma_{3,t+1|t} = \frac{E_t(R_{t+1} - \mu)^3}{\sigma_{t+1|t}^3} \quad (35)$$

and

$$\gamma_{4,t+1|t} = \frac{E_t(R_{t+1} - \mu)^4}{\sigma_{t+1|t}^4} - 3. \quad (36)$$

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

$$F_{t+1|t}(z) = \Phi(z) - \gamma_{3,t+1|t} \frac{1}{3!} D^2 \phi(z) + \gamma_{4,t+1|t} \frac{1}{4!} D^3 \phi(z). \quad (37)$$

Notice that, because

$$\begin{aligned} D\phi(z) &= -z\phi(z) \\ D^2\phi(z) &= (z^2-1)\phi(z) \\ D^3\phi(z) &= -(z^3-3z)\phi(z), \end{aligned} \quad (38)$$

we have

$$F_{t+1|t}(z) = \Phi(z) - \phi(z) \left(\frac{\gamma_{3,t+1|t}}{3!} (z^2-1) + \frac{\gamma_{4,t+1|t}}{4!} (z^3-3z) \right). \quad (39)$$

In general, the probability of a positive return is equivalent to the probability of the standardized return being bigger than $-\mu/\sigma_{t+1|t}$; hence the probability of a positive return is

$$\begin{aligned} 1 - F_{t+1|t}(-\mu/\sigma_{t+1|t}) &= 1 - \Phi(-\mu/\sigma_{t+1|t}) \\ &+ \phi(-\mu/\sigma_{t+1|t}) \left(\frac{\gamma_{3,t+1|t}}{3!} (\mu^2/\sigma_{t+1|t}^2 - 1) + \frac{\gamma_{4,t+1|t}}{4!} (-\mu^3/\sigma_{t+1|t}^3 + 3\mu/\sigma_{t+1|t}) \right), \end{aligned} \quad (40)$$

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

$$1 - F_{t+1|t}(-\mu/\sigma) = 1 - \Phi(-\mu/\sigma) + \phi(-\mu/\sigma) \left(\frac{\gamma_{3,t+1|t}}{3!} (\mu^2/\sigma^2 - 1) + \frac{\gamma_{4,t+1|t}}{4!} (-\mu^3/\sigma^3 + 3\mu/\sigma) \right), \quad (41)$$

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

$$1 - F_{t+1|t}(0) = 1/2 - \frac{\gamma_{3,t+1|t}}{\sqrt{2\pi}3!} \approx 1/2 - 0.0665\gamma_{3,t+1|t}, \quad (42)$$

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.

7. Forecasting S&P 100 Market Direction Using VIX Volatility

We now present a simple empirical example, both to illustrate our ideas and methods and to provide some preliminary evidence as to their empirical relevance. As we have shown, a key ingredient in any successful equity sign forecast is a successful volatility forecast, which could be obtained either from an econometric volatility model (see Andersen, Bollerslev, Diebold and Labys, 2003, for a comparison of several) or from the marketplace via option-implied volatilities. Here we take the latter route, forecasting return signs using the VIX index of S&P 100 volatility, which is traded on the Chicago Board Options Exchange (CBOE), and which is widely viewed as a good indicator of market sentiment (an “investor fear gauge” in the colorful language of Whaley, 2000). It is calculated as a weighted average of the implied volatility from eight S&P 100 options (four calls and four puts nearest to the money), with the observed option quotes interpolated so as to obtain a synthetic volatility from an at-the-money option with a maturity of exactly thirty days.¹⁸

To match the maturity of the VIX, we forecast signs of 30-day returns. More precisely, let $R_{t+1:t+30}$ be the S&P 100 30-day excess return (relative to the 30-day return on a 3-month Treasury bill), and define the sign indicator sequence as $I_{t+30} = 1$ if $R_{t+1:t+30} > 0$. We want to forecast I_{t+30} , and the analysis of this paper strongly suggests a model of the form

$$I_{t+30} = F\left(\frac{\mu}{\sigma_t}\right) + e_{t+30}, \quad (43)$$

where σ_t denotes a forecast of 30-day excess return volatility and $F(\bullet)$ is a monotone function with a left limit of zero and a right limit of one. An obvious choice of $F(\bullet)$ is the logistic,

$$F(x) = \frac{e^x}{1+e^x}, \quad (44)$$

¹⁸ See Whaley (2000) for details of VIX construction.

which produces the popular logit model. Finally, for σ_t we insert VIX_t .

Logit models are traditionally estimated by the method of maximum likelihood (ML) under an iid assumption. However, the 29-day overlap in the daily cumulative 30-day returns introduces dependence, biasing the traditional ML standard errors. We therefore estimate the logit model using GMM instead. We use the ML score equation as the orthogonality condition, and we compute standard errors allowing for autocorrelation using a quadratic kernel with a bandwidth of 29 days.¹⁹

In Figure 8 we plot the daily 30-day VIX volatilities for January 2, 1986 through May 15, 2003, reported as annualized percentages. Notice the very high VIX during the 1987 crash and the very low VIX during the mid-1990s bull market. The increase in volatility during the new millennium is also evident.

To get a feel for the real-world importance (or lack thereof) of the time-varying market sign probabilities induced by time-varying market volatility, we can examine directly the probabilities implied by the estimated logit regression of I_{t+30} on $1/VIX_t$. Near-constancy of the probabilities would indicate negligible sign forecastability. We plot the probabilities in Figure 9; their range is strikingly wide, from a low of 51% during the 1987 crash to a high of 72% during the mid-1990s. In addition, the range of probabilities matches well the range of the monthly probabilities in the simulated stochastic volatility model (Figure 3, bottom left panel).²⁰ Finally, note that the traditional market timing tests such as Hendriksson-Merton (1981) and Pesaran-Timmermann (1992) have no power to detect this strong sign dependence, because the sign probability never drops below 0.5.

The responsiveness of the sign probability forecast to changes in the VIX is given by the marginal effect, \mathfrak{R}_t , as defined in Section 3 and plotted in Figure 10. As predicted theoretically, the marginal effect is always negative, and it is largest (in absolute value) when the level of volatility is relatively low, as it was in the mid-1990s. We highlight this relationship in Figure 11, in which we display a scatterplot of \mathfrak{R}_t against VIX. The empirical relationship (Figure 11) matches closely the relevant portion of the theoretical relationship (Figure 2).

8. Concluding Remarks and Directions for Future Research

Our contribution is two-fold. First, we have shown that, given the widely-accepted statistical properties of returns, one should not be surprised to find sign forecastability, and that under the realistic

¹⁹ Under our assumptions, the logit coefficient on $1/VIX$ is the expected return, μ , which we estimate to be 8.70 with a standard error of 1.93.

²⁰ The dynamic patterns are quite different, however, but that is simply because of the different number of observations and the 29-day overlap in the data.

assumption of persistent volatility dynamics, sign forecastability is to be expected even when 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, we show in a realistically calibrated simulation exercise that sign forecastability indeed appears strongest at intermediate horizons of two or three months, and we show in an empirical analysis that sign dependence is very much present in actual U.S. equity returns and has properties that match closely our theoretical predictions. In addition, we show that the nonlinear nature of sign dependence makes it unlikely to be found via traditional approaches such as analysis of sign autocorrelations, runs tests, or market timing tests.

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, for example, simple Sharpe ratios, such as Stutzer’s (2001) portfolio performance index.

References

- 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, 2000, Estimating jump-diffusions for equity returns, Manuscript, Kellogg Graduate School of Management.
- Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold, 2003, 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.
- 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.
- Balvers, Ronald J., Thomas F. Cosimano, and Bill McDonald, 1990, Predicting stock returns in an efficient market, *Journal of Finance* 45, 1109-1128.
- 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, 1999, Pricing options under stochastic volatility: an econometric analysis, Manuscript, Carlson School of Management, 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, 2003, On the relationship between the conditional mean and volatility of stock returns: a latent var approach, Qiang Kang, *Journal of Financial Economics*, forthcoming.
- 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, 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.
- Chernov, Mikhail, A. Ronald Gallant, Eric Ghysels, and George Tauchen, 2001, Alternative models for stock price dynamics, Manuscript, University of North Carolina.
- 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.
- 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.
- Eraker, Bjorn, Michael Johannes, and Nicholas Polson, 2000, The role of jumps in returns and volatility, *Journal of Finance*, forthcoming.
- Estrella, Arturo, 1998, A new measure of fit for equations with dichotomous dependent variables, *Journal of Business and Economic Statistics* 16, 198-205.

- 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, 1993, The risk and predictability of international equity returns, *Review of Financial Studies* 6, 527-566.
- 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.
- Gencay, Ramazan, 1998, Optimization of technical trading strategies and the profitability in security markets, *Economics Letters* 59, 249-254.
- 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.
- Glosten Lawrence R., Ravi Jaganathan, 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.
- Granger, Clive W.J., 1969, Prediction with a generalized cost of error function, *Operational Research Quarterly* 20, 199-207.
- Hagigi, Moshe, and Brian Kluger, 1987, Safety first: an alternative performance measure, *Journal of Portfolio Management* 13, 34-40.
- Hansen, Bruce E., 1994, Autoregressive conditional density estimation, *International Economic Review* 35, 705-730.
- Harvey, Campbell R., and Akhtar Siddique, 2000, Conditional skewness in asset pricing tests, *Journal of*


Finance 55, 1263-1296.

- Harvey, Campbell R., and Robert Whaley, 1992, Market volatility prediction and the efficiency of the S&P 100 index option market, *Journal of Financial Economics* 31, 43-73.
- 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.
- Jegadeesh, Narasimhan, 1990, Evidence of predictable behavior of security returns, *Journal of Finance* 45, 881-898.
- Jensen, Michael C., 1978, Some anomalous evidence regarding market efficiency, *Journal of Financial Economics* 6, 95-101.
- Johannes, Michael, Nicholas Polson and Jon Stroud, 2002, Sequential optimal portfolio performance: market and volatility timing, Manuscript, Graduate School of Business, Columbia University.
- Jones, Christopher, 1999, The dynamics of stochastic volatility, Manuscript, University of Rochester.
- 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, 1999, Forecasting stock indices: a comparison of classification and level estimation models, working paper, Indiana University.
- 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.
- Mankiw, N. Gregory, David Romer, and Matthew D. Shapiro, 1991, Stock market forecastability and

- volatility: a statistical appraisal, *Review of Economic Studies* 58, 455-77.
- Meddahi, Nour, 2001, An eigenfunction approach for volatility modeling, Manuscript, University of Montreal.
- Meddahi, Nour, and Eric Renault, 2000, Temporal aggregation of volatility models, working paper, CIRANO.
- 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.
- Pan, Jun, 2000, Integrated time-series analysis of spot and option prices, Manuscript, Sloan School of Management, MIT.
- Patelis, Alex D., 1997, Stock return predictability and the role of monetary policy, *Journal of Finance* 52, 1951-1972.
- 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, 2001, How costly is it to ignore breaks when forecasting the direction of a time series?, Manuscript, Cambridge University and UCSD.
- Roy, A.D. 1952, Safety-first and the holding of assets, *Econometrica* 20, 431- 449.
- Sentana, Enrique, and Sushil Wadhvani, 1991, Semi-parametric estimation and the predictability of stock market returns: some lessons from Japan, *Review of Economic Studies* 58, 547-563.
- Stutzer, Michael, 2001, A portfolio performance index and its implications, Manuscript, University of Iowa.
- Sweeney, Richard J., 1986, Beating the foreign exchange market, *Journal of Finance* 41, 163-182.
- 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.
- Whaley, Robert, 2000, The investor fear gauge, *Journal of Portfolio Management* 26, 12-27.
- 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, NYU, Stern School of Business.
- Womack, Kent L., 1996, Do brokerage analysts' recommendations have investment value?, *Journal of*

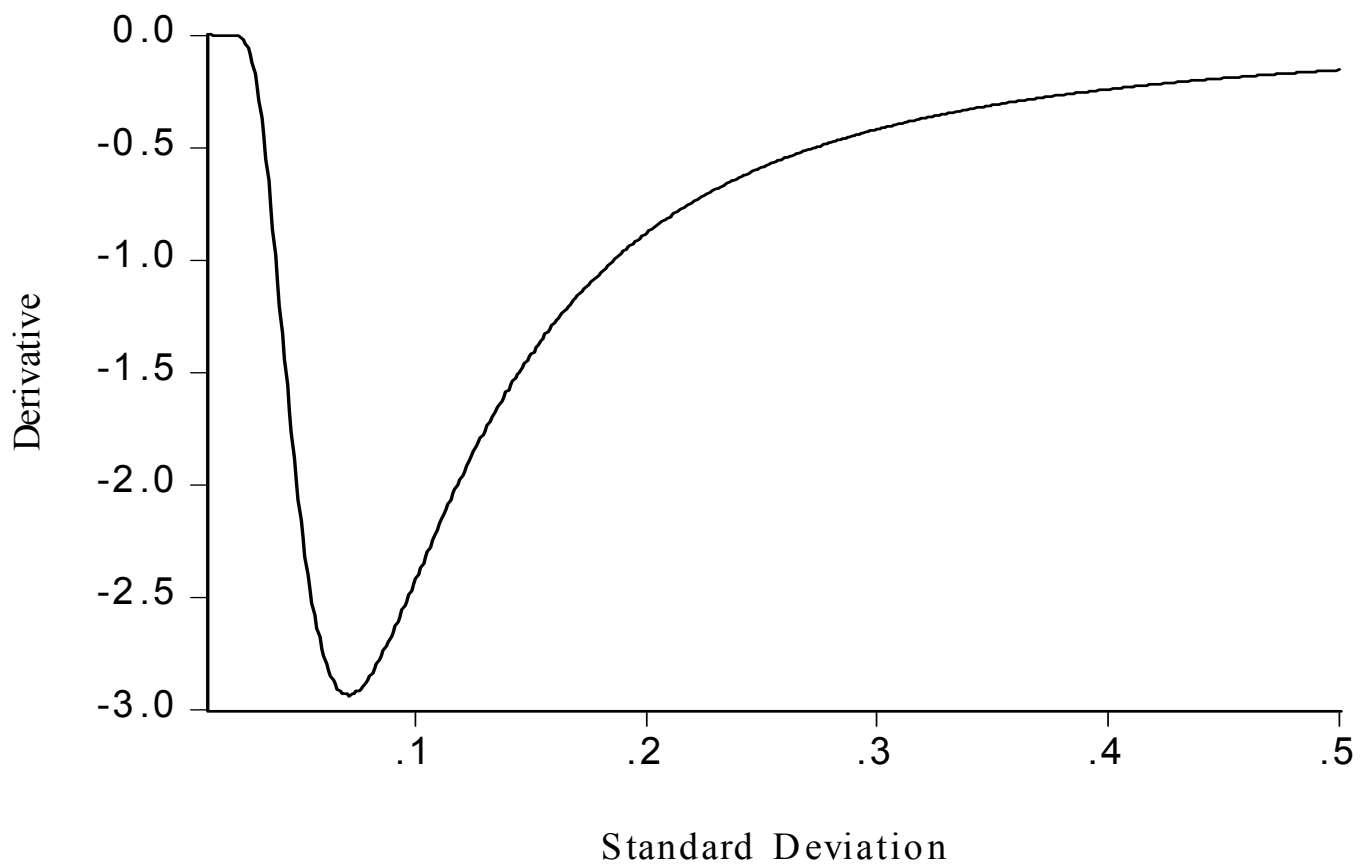
Finance 51, 137-167.

Figure 1
The Dependence of Sign Probability on Volatility



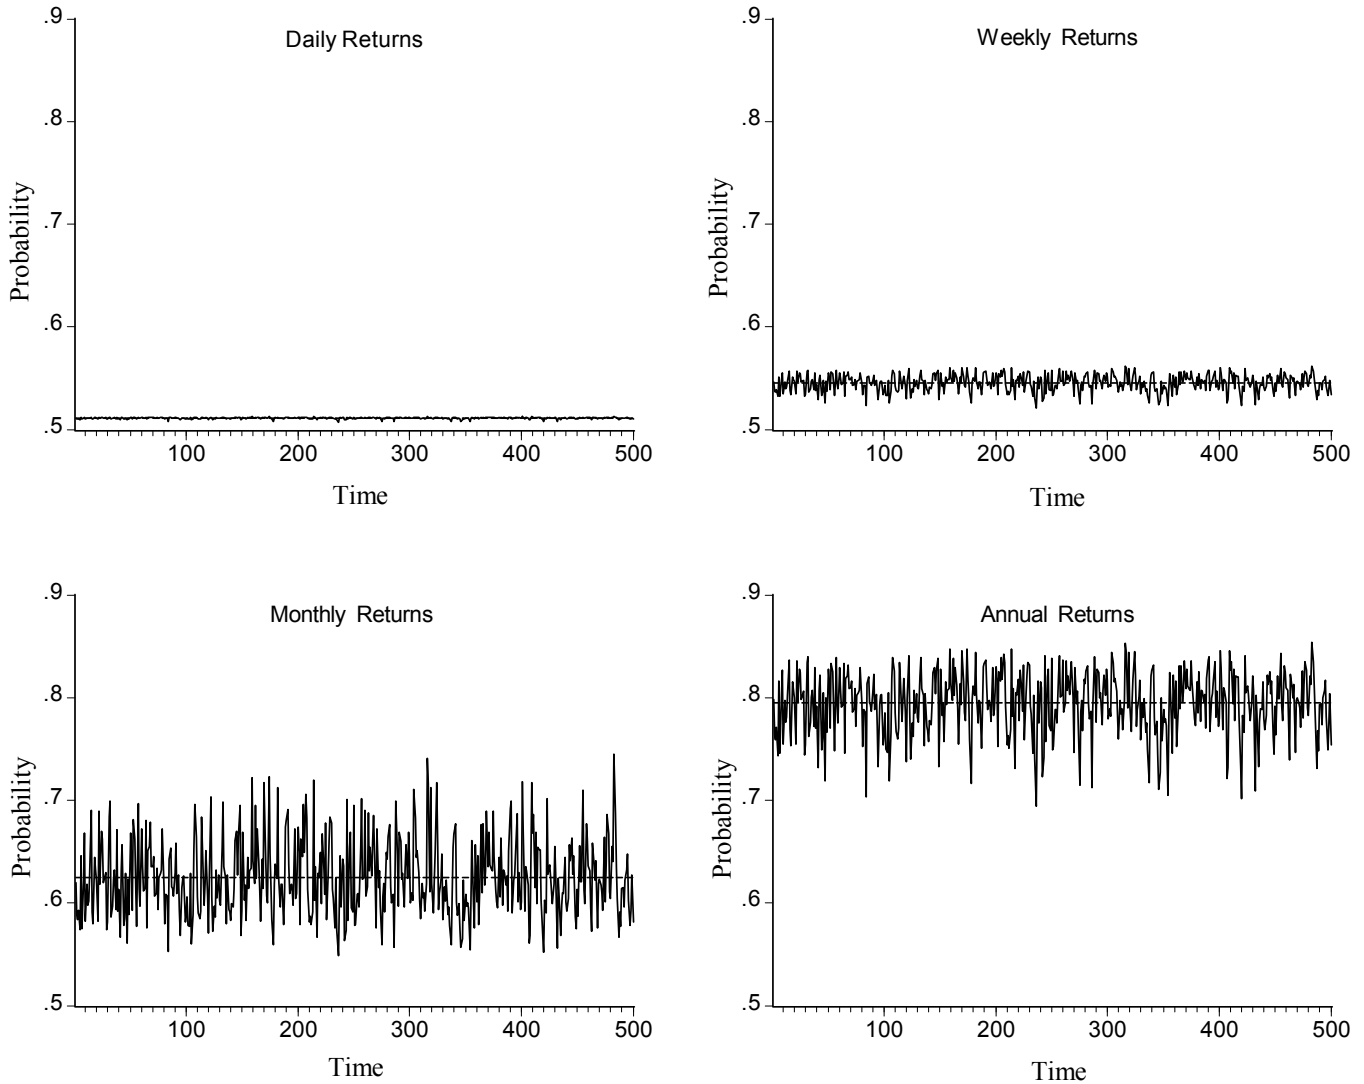
Notes to figure: We show two Gaussian return densities, each with expected return of ten percent. The first return has a standard deviation of five percent and hence is positive with probability 0.98 (the area to the right of zero under the more peaked density function). The second return has a standard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



Notes to figure: We plot \mathfrak{R}_t , the derivative of the probability of a positive return with respect to return volatility, as a function of volatility. We assume Gaussian returns with ten percent expected return.

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

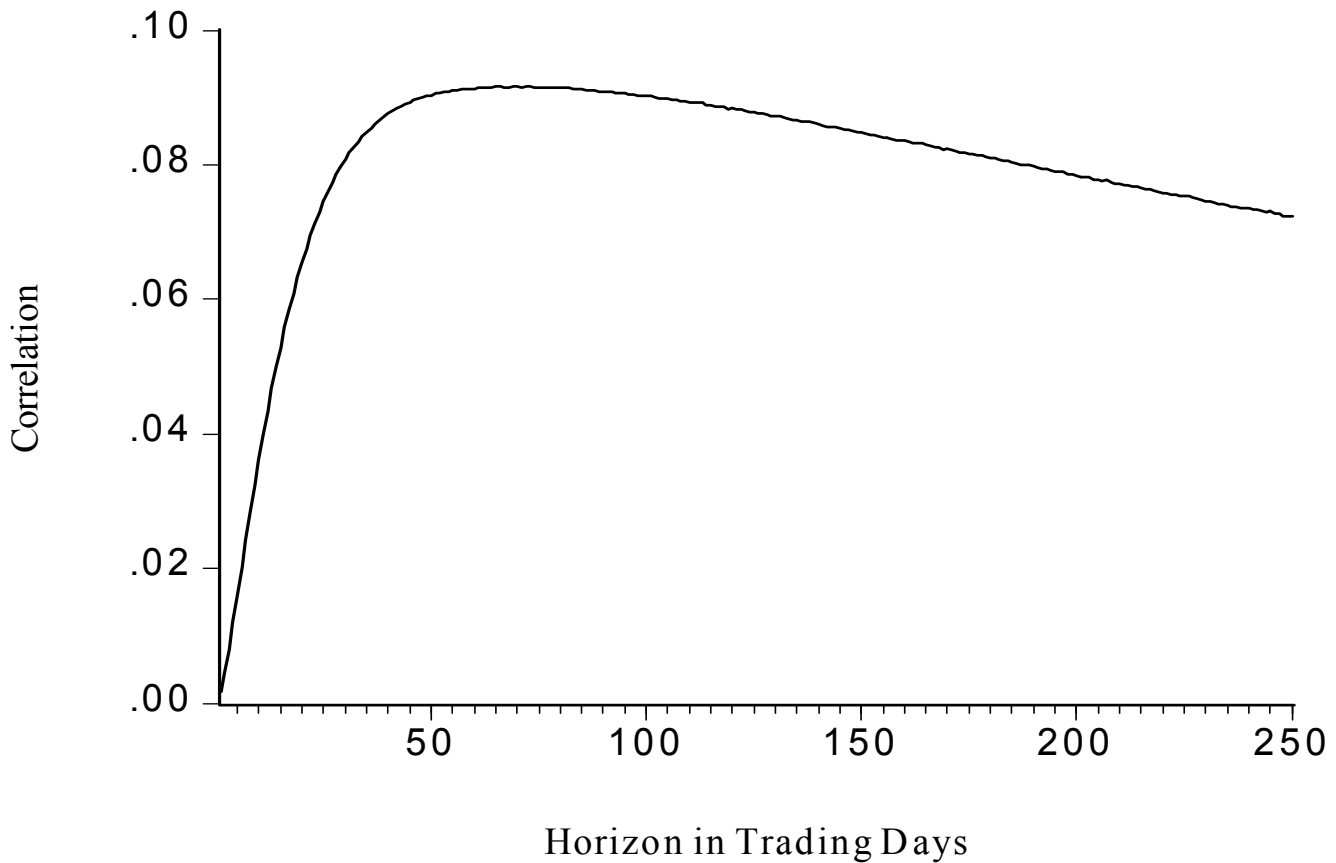


Notes to figure: We simulate asset prices from a stochastic volatility model parameterized as

$$\begin{aligned} dS(t) &= 0.10S dt + \sigma(t)S dz_1 \\ d\sigma^2(t) &= 2(0.015 - \sigma^2(t))dt + 0.15\sigma(t)dz_2, \end{aligned}$$

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
Various Horizons, Benchmark Parameters




Notes to figure: We simulate asset prices from a stochastic volatility model parameterized as

$$dS(t) = 0.10S dt + \sigma(t)S dz_1$$

$$d\sigma^2(t) = 2(0.015 - \sigma^2(t))dt + 0.15\sigma(t)dz_2,$$

with $\text{corr}(dz_1, dz_2) = -0.5$. We then calculate 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 results in the text.

Figure 5
Correlation between Sign Forecasts and Realizations
Various Horizons, Small Return Drift




Notes to figure: We simulate asset prices from the a stochastic volatility model parameterized as

$$dS(t) = \mu S dt + \sigma(t) S dz_1$$

$$d\sigma^2(t) = 2(0.015 - \sigma^2(t)) dt + 0.15 \sigma(t) dz_2,$$

with $\text{corr}(dz_1, dz_2) = -0.5$, with drift of 0.05 (top line) and 0.00 (bottom line). 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 results in the text.

Figure 6
Correlation between Sign Forecasts and Realizations
Various Horizons, Small Volatility Persistence




Notes to figure: We simulate asset prices from a stochastic volatility model parameterized as

$$dS(t) = 0.10S dt + \sigma(t)S dz_1$$

$$d\sigma^2(t) = \kappa(0.015 - \sigma^2(t))dt + 0.15\sigma(t)dz_2,$$

with $\text{corr}(dz_1, dz_2) = -0.5$, and with volatility persistence parameter $\kappa=5$ (top line) and $\kappa=10$ (bottom line). 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 in the text.

Figure 7
Correlation between Sign Forecasts and Realizations
and First Autocorrelation of Return Signs
Various Horizons, Benchmark Parameters




Notes to figure: 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

$$dS(t) = 0.10Sdt + \sigma(t)Sdz_1$$

$$d\sigma^2(t) = 2(0.015 - \sigma^2(t))dt + 0.15\sigma(t)dz_2,$$


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 in the text.

Figure 8
S&P 100 Volatility (VIX)




Notes to figure: We show the time series of S&P 100 volatility, as summarized in the VIX volatility index, traded on the Chicago Board Options Exchange.

Figure 9
Probability of a Positive Thirty-Day Excess S&P 100 Return




Notes to figure: We show the time series of conditional probability of a positive thirty-day excess S&P 100 return, as assessed using a logit model in conjunction with the VIX volatility forecast.

Figure 10
Time Series of Responsiveness of Sign Forecasts to Changes in VIX



Notes to figure: We show the time series of the derivative of the conditional probability of a positive return with respect to VIX, which provides a measure of the responsiveness of the sign probability to changes in underlying volatility conditions.

Figure 11
Scatterplot of Responsiveness of Sign Forecasts to Changes in VIX vs. VIX



Notes to figure: We show the derivative of the conditional probability of a positive return with respect to VIX, graphed against VIX.