

Financial Asset Returns, Market Timing, and Volatility Dynamics

Peter F. Christoffersen
Faculty of Management
McGill University
and CIRANO

Francis X. Diebold
Department of Economics
University of Pennsylvania
and NBER

March 2001

This Draft / Print: January 14, 2002

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 with implications for market timing, 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: (1) 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. (2) The standard finding of little or no conditional mean dependence is entirely consistent with a significant degree of sign dependence and volatility dependence. In particular, sign dependence does not imply market inefficiency. (3) Sign dependence is not likely to be found via analysis of sign autocorrelations, because the nature of sign dependence is highly nonlinear. (4) 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.

Acknowledgments: This work is supported by the National Science Foundation, the Wharton Financial Institutions Center, FCAR, IFM2 and SSHRC. For helpful comments we thank Michael Brandt, Michel Dacorogna, Graham Elliott, Christian Gouriéroux, Clive Granger, Joanna Jasiak, Blake LeBaron, Bruce Lehman, Nour Meddahi, Sergei Sarkissian, Frank Schorfheide, Allan Timmermann, Hal White, and seminar participants at McGill University, UC San Diego and at the Third Annual Conference on Financial Econometrics at the University of Waterloo. Sean Campbell and Clara Vega provided outstanding research assistance. All inadequacies are ours alone.

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 with implications for market timing, 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.¹

Second, consider dependence and hence forecastability of asset return signs, and the implications for market timing. The possibility of market timing focuses interest not on prediction of returns directly, but rather on prediction of *signs* of returns, on the grounds that profitable trading strategies may result if one is successful at forecasting return signs, quite apart from whether one is successful at forecasting the mean of returns. A well-known and classic example, even at the textbook level such as Levich (1998, Chapter 8), involves foreign exchange trading. 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.² Positive profits may be made when the sign forecast is correct.³ Interestingly, there is accumulating evidence that

¹ We emphasize the word “approximation,” as small deviations from conditional mean independence at short horizons are well-known (see for example Lo and MacKinlay, 1999), typically emerging as very weak but nonzero serial correlation in high-frequency returns. As we discuss below, the presence of time-varying risk premia would of course also imply some degree of conditional mean predictability.

² Generalizations to timing multiple asset classes, such as stock and bond markets, are immediate, in which case one bases the allocation strategy on a forecast of the sign of the return spread.

³ We say “may” for a number of reasons. First, of course, one must adjust for interest and other costs. Second, and importantly, one may be led astray by focusing exclusively on sign, as one must also consider the absolute magnitude of returns, and, for example, the timing of returns of very large magnitude. For additional discussion, see Cumby and Modest (1987).

sign forecasting – again the key ingredient in market timing – 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), and White (2000), among others.

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 contributions in the GARCH tradition, while Ghysels, Harvey and Renault (1996) survey the stochastic volatility literature, Franses and van Dijk (2000) survey models of regime-switching volatility, and Andersen, Bollerslev and Diebold (2002) survey models of realized volatility. Note that conditional volatility dependence—in sharp contrast to conditional mean dependence—need not be traded away in efficient markets. Thus evidence of conditional volatility dependence is much more prevalent and much less controversial than is evidence of conditional mean dependence. 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 models 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. Much less is known, however, about sign dependence in general, 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) 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.
- (2) The standard finding of little or no conditional mean dependence is entirely consistent with a significant degree of sign dependence and volatility dependence. In particular, sign dependence does not imply market inefficiency.
- (3) Sign dependence is not likely to be found via analysis of sign autocorrelations, because the nature of sign dependence is highly nonlinear.

- (4) 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 at intermediate return horizons of two or three months.

We derive the theoretical results in a general setting, and we illustrate them using a popular parametric model of asset prices and returns.

We proceed as follows. In sections 2 and 3 we build intuition by sketching the main results in simple contexts, focusing primarily on the conditionally Gaussian case. In section 2 we discuss the basic framework, volatility dynamics as a source of sign dynamics, and implications for market efficiency, and in section 3 we present more general results allowing for unconditional and conditional skewness and kurtosis. Next, in section 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. 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 offer concluding remarks and discuss directions for future work, including methods of exploiting sign forecastability for improved asset allocation.

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 $E(R_{t+1}|\Omega_t)$ is a nonconstant function of 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 $E(I(R_{t+1}>0)|\Omega_t)$ is a nonconstant function of 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 $\sigma_{t+1|t}^2 \equiv \text{Var}(R_{t+1}|\Omega_t)$ is a nonconstant function of Ω_t .

Having made the requisite definitions, we now proceed to characterize the relationship between sign and volatility dynamics, and the implications of sign dynamics for market efficiency.

2.1 Sign Dynamics Follow from Volatility Dynamics

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

Consider the prevalence of volatility dynamics in high-frequency asset returns, and the positive expected returns earned on risky assets. To take a very simple but illuminating example, assume the returns on a generic risky asset are distributed as

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

and therefore display conditional mean independence and conditional variance dependence. The probability of a positive return is

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

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 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 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.⁵

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 and in fact is significantly enriched, as we will show in subsequent detail. In particular,

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

⁵ 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 mean.

where $F(\bullet)$ denotes the relevant c.d.f. Hence the sign of returns is still forecastable so long as the mean return is non-zero and volatility is dynamic.

One might naturally wonder whether sign forecastability implies violation of weak market efficiency, in the sense of conditional mean dependence, in which case returns would be forecastable based upon their own past. The answer is of course no, as the above example of a process displaying both conditional mean independence and sign forecastability made clear. In the next sub-section, we elaborate on this important point.

2.2 Return Sign Forecastability is Consistent with Market Efficiency

Market timing is clearly of interest for active asset allocation, and recent sophisticated empirical work suggests that market timing *can* be done. Successful market timing requires sign forecastability, and if return signs are forecastable then returns must be somehow dependent. When sign forecastability is found empirically, it is tempting to conjecture that it is driven by subtle nonlinear dependence in conditional mean dynamics, which would be missed in standard analyses of (linear) dependence, such as those based on return autocorrelations.

The key insight, however, 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. Furthermore, sign dynamics do not imply violation of market efficiency in the mean-squared-error forecastability sense. We realize that general equilibrium definitions of market efficiency, which are different from the MSE-based definition used here, can be consistent with some degree of conditional mean forecastability, for example over the business cycle.⁶ One should of course not be surprised to find evidence of sign prediction if the conditional mean is predictable. Thus, definitions of market efficiency which allow, for example, for the possibility of time-varying risk-premia will generally lead to sign predictability. We deliberately restrict ourselves to working with return processes that are conditional mean independent and therefore satisfy a simple version of market efficiency, yet they are sign predictable. Our purpose is to assess if volatility dynamics alone can generate magnitudes of sign predictability similar to those observed empirically.

⁶ 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 Wadhwani (1991), and Sweeney (1986).

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

$$R_{t+1} = \text{sign}_{t+1} \cdot |R_{t+1}|.$$

As we have discussed, both of the right-hand-side components of returns are forecastable, yet the left-hand-side variable, returns themselves, are approximately unforecastable.⁷

3. Allowing for (Potentially Time-Varying) Conditional Skewness and Kurtosis

In the previous section 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, and in fact is significantly enriched. Here we defend that assertion by assessing the potential impact of time-varying higher-ordered conditional moments.

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

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

and assume they are approximately distributed according to the pdf

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

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

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

⁷ This is an example of a nonlinear “common feature,” in the terminology of Engle and Kozicki (1993): signs are conditional mean dependent and absolute returns are conditional mean dependent, yet their product is conditional mean *independent*.

and

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

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

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

Notice that, because

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

we have:

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

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

$$1 - F_{t+1|t}(-\mu/\sigma_{t+1|t}) = 1 - \Phi(-\mu/\sigma_{t+1|t}) + \varphi(-\mu/\sigma_{t+1|t}) \left[\frac{\gamma_{3,t+1|t}}{3!} (\mu^2/\sigma_{t+1|t}^2 - 1) + \frac{\gamma_{4,t+1|t}}{4!} (-\mu^3/\sigma_{t+1|t}^3 + 3\mu/\sigma_{t+1|t}) \right],$$

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

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

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

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

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.

4. Sign Forecasts, Sign Realizations, and the Link

Here we examine a number of issues relevant to the quantification of sign forecastability. How, if at all, does the sign forecast derivative with respect to volatility vary as a function of volatility, and in what volatility region is it maximized? What is the correlation between sign forecasts and realizations, and how, if at all, is the correlation related to the volatility of sign forecasts? What is the autocorrelation structure of sign realizations, and is it likely to be useful for identifying sign forecastability?

4.1 Measuring Sign forecastability I: The Responsiveness of the Sign Forecast to Changes in Volatility

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

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

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

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

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

$$\mathfrak{R} = -\varphi\left(\frac{\mu}{\sigma_{t+1|t}}\right)\left(\frac{\mu}{\sigma_{t+1|t}^2}\right).$$

In Figure 2, we work in a Gaussian situation and plot the values of the \mathfrak{R} measure for different

volatilities, keeping the mean at ten percent. The responsiveness has an interior optimum and is thus not monotonically decreasing in the standard deviation. 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 $\frac{1}{2}$, and hence responsiveness is again tiny. Intermediate values of σ , however, can produce greater responsiveness. Notice that the probability of a positive return is always decreasing (the derivative is always negative) in the conditional standard deviation, but one can show that the responsiveness is the highest when $\sigma_{t+1|t} = \mu/\sqrt{2} \approx 0.0707$.

4.2 Measuring Sign Forecastability II: The Correlation Between Sign Forecasts and Realizations

In the previous subsection we considered the issue of responsiveness of probability forecasts of 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.⁸

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

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

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

$$E[I_{t+1}P_{t+1|t}] = E[E_t\{I_{t+1}P_{t+1|t}\}] = E[E_t\{I_{t+1}\}P_{t+1|t}] = E[P_{t+1|t}^2].$$

Hence we have

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

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

Converting to correlation, we can write

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

Notice that the correlation between sign forecasts and realizations depends only on the standard deviation

⁸ For an interesting discussion of R^2 -type measures in binary regressions, see Estrella (1998).

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.

4.3 On the Difficulty of Detecting Sign Forecastability: The Autocorrelation of Sign Realizations

Turning attention now to the first-order autocovariance of the indicator sequence, we can write

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

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

$$E[I_{t+1}I_t] = E[E_t(I_{t+1}I_t)] = E[I_t E_t(I_{t+1})] = E[I_t P_{t+1|t}].$$

Hence

$$\text{Cov}[I_{t+1}, I_t] = E[P_{t+1|t}I_t] - P^2.$$

Converting now to autocorrelations, we have

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

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

$$\text{Corr}(P_{t+1|t}, I_t) = \frac{E[P_{t+1|t}I_t] - E[P_{t+1|t}]E[I_t]}{\text{Std}(P_{t+1|t}) \text{Std}(I_t)} = \frac{E[P_{t+1|t}I_t] - P^2}{\text{Std}(P_{t+1|t}) \sqrt{P(1-P)}}. \quad (3)$$

Substituting (3) into the expression for the autocorrelation of the indicator sequence (2), we get

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

Substituting the expression for the correlation between the ex-ante sign forecast and the ex-post sign-realization (1), we have:

$$\text{Corr}(I_{t+1}, I_t) = \text{Corr}(P_{t+1|t}, I_{t+1}) \text{Corr}(P_{t+1|t}, I_t). \quad (4)$$

If all three correlations in (4) are positive and bounded away from one (which is a realistic assumption),

we obtain

$$\text{Corr}(I_{t+1}, I_t) < \text{Corr}(P_{t+1|t}, I_{t+1}). \quad (5)$$

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.

The upshot for empirical work is clear: sign forecastability will generally be difficult to detect from sign autocorrelations, which can be small even when sign forecastability is large.

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 either expected returns approach zero or volatility forecastability approaches zero, sign forecastability approaches zero.⁹ 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 the shortest and longest extremes. In this section, we evaluate this conjecture under realistic stochastic processes for returns.

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.¹⁰ To circumvent this problem, we work with a continuous-time stochastic volatility model, focusing in particular on the convenient and popular model of Heston (1993).

⁹ Of course, higher-ordered conditional moment dynamics can contribute to sign predictability as well.

¹⁰ For penetrating insight into the difficulties involved in the temporal aggregation of discrete-time volatility models, see Meddahi and Renault (2000). Discrete-time volatility models with convenient temporal aggregation properties could potentially be obtained using the results of Meddahi (2001), Darolles, Gouriéroux, and Jasiak (2001), or Heston and Nandi (2000).

Working with the continuous-time stochastic volatility model has the added benefit that temporal aggregation and increasing horizons are interchangeable. We simply assume that the current volatility is observed at some frequency, say monthly, and calculate the one-month ahead sign prediction. This can be thought of as either working with monthly returns, or working with a one-month horizon, as in either case only the instantaneous volatility is needed.

5.1 Simulation Design

The stochastic volatility model parsimoniously captures many of the stylized facts of asset returns, including skewness, leptokurtosis and volatility persistence. Furthermore, the stochastic volatility model has the great advantage that the conditional density can be easily calculated at any forecast horizon. For all of these reasons, the stochastic volatility model 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), and Pan (2000), among others.

The stochastic volatility process is a bivariate diffusion defined as

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

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 denoted by μ , the long run variance by θ , the speed of adjustment of the variance by κ , and the volatility of volatility by η . Keeping the parameter values such that $\eta^2 \leq 2\theta\kappa$ ensures that the continuous-time variance process stays strictly positive almost surely.

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

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

Notice that although the instantaneous drift is simply a constant, the continuously compounded return has a slight time-variation 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.¹¹ In particular,

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

where $f(x, \sigma^2, \tau; \psi)$ 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

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

where

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

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

$$g = \frac{\kappa - \rho\eta\psi i + d}{\kappa - \rho\eta\psi i - d}$$

and

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

In the simulations below, we generate 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 values estimated in the empirical papers cited above. 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%, an unconditional skewness of about -0.1, and an unconditional excess

¹¹ As in the Gaussian case, the computation of the sign probability requires numerical integration, but the well-behaved integrand renders the integration straightforward.

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.

In Figure 3 we plot the sign forecast from a typical sample path of the simulated process. We show daily, weekly, monthly and annual conditional as well as unconditional sign probabilities. Notice the increase in the unconditional probability of a positive return as well as the decrease in the persistence of the conditional probability of a positive return as time is aggregated. In Figure 4 we plot the unconditional probability of a positive return across horizons.

5.2 Simulation Results

Figure 5 supports our conjecture that sign dynamics will be most prevalent at intermediate frequencies. The correlation between the ex ante sign forecast and the ex post sign realization is highest at around the 2-3 month frequency (corresponding to 40-60 trading days). The correlation is quite low at the highest frequencies and tapers off again as time is aggregated to the annual frequency.

Figure 6 shows the correlation between the ex ante sign forecast and the ex post sign realization when the drift of the instantaneous rate of return, μ , is 0.05 and zero, but all other parameters are kept at their benchmark values. Although signs are much less forecastable when μ is zero, we still get some forecastability at intermediate horizons, stemming from a non-zero leverage effect, ρ , which interacts with the volatility dynamics.

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.¹² 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

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

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

Figure 7 shows the impact of increasing κ to 10 and thereby lowering the volatility persistence. A κ of 10 corresponds to a daily volatility persistence of about 0.96.

Figure 8 shows another important result, also suggested but not conclusively established by our

¹² Meddahi and Renault (2000) find similar leverage effects arising under aggregation even if the innovations are uncorrelated.

earlier analytic work: the simple autocorrelation of the sign realization is relatively small at all horizons. This contrasts with the correlation between the signs and their forecasts. The magnitude of the autocorrelation function implies that the search for a simple linear forecast of signs from their own past is not likely to be fruitful. Figure 9 compares the correlation between the ex-ante probability forecasts and the ex-post realizations with the autocorrelations of the realization.

It appears that the nonlinearity of the volatility-based forecast renders the linear forecastability in the Markov chain small. In our setup, having a low conditional variance is key for getting a high probability of a positive return, regardless of the sign of yesterday's return. In a symmetric volatility model, *ceteris paribus*, both small positive and small negative returns (in absolute value) will generate a smaller conditional variance and thus a higher probability of a positive return tomorrow, whereas a large positive return (and a large negative return) will increase the conditional volatility and thus lower the probability of a positive return.¹³

Keeping in mind the standard practice in the literature of specifying Probit or Logit regressions in search of sign forecasts, the above discussion leads one to the following question: What is the implied Probit/Logit regression on an indicator variable of the sign of returns that one should run to detect sign dynamics? It is clear that the regressor should be a nonlinear function of conditional variance, and more specifically, our earlier analysis suggests that the obvious regressor is an estimate of $\mu / \sigma_{t+1,t}$. If, moreover, volatility is a function of predetermined variables (e.g., current and past absolute returns, volume, day of week, etc.), then Probit/Logit models which include these variables should show some significance in explaining the sign. Note well, in particular, that finding a significant explanatory variable in a Probit regression on return signs could simply be suggestive that the variable used is correlated with the volatility of the asset return; it does *not* necessarily imply linear *or* nonlinear conditional-mean return dependence and is entirely consistent with market efficiency.

5.5 Extensions

The simulation analysis above was much more than a mere illustration of various analytic results. Indeed, the nonlinearity associated with sign dynamics makes full analytic results very difficult to obtain. Hence, extending the simulation analysis to include richer data-generating processes is of interest. Obvious extensions include:

- (1) Multifactor models with possible jumps. We are currently working with one of the simplest dynamic one-factor volatility specifications available, but we could easily entertain the

¹³ Note that nonlinear models, such as the neural networks considered by Kuan and Liu (1995), might be useful for sign forecasting.

possibilities of multiple volatility factors and jumps as in Alizadeh, Brandt and Diebold (2002) and Chernov, Gallant, Ghysels and Tauchen (2001).

- (2) Long-memory models. So far we have worked only with short-memory volatility models, whereas there is ample evidence of long memory dynamics in asset return volatility (e.g., Andersen, Bollerslev, Diebold and Ebens, 2001, Andersen, Bollerslev, Diebold and Labys, 2001a and 2001b). Extensions to long-memory volatility models could of course easily be made.
- (3) Heteroskewness and heterokurtosis. El Babsiri and Zakoian (2001) argue for a richer class of GARCH models that incorporates conditional heteroskewness and heterokurtosis. As shown in Section 4, these models would have important implications for sign dynamics as well.
- (4) Multivariate models. We have confined attention to the univariate case, but portfolio management decisions are inherently multivariate in nature, and an investigation of multivariate interaction among return signs might provide useful insight.

Although we have not pursued these extensions here, they would be interesting to explore in future work.

6. Concluding Remarks

We have shown that, given the widely-accepted statistical properties of speculative prices, one should not be surprised to find forecastability in the sign of returns. Under the realistic assumptions of a constant but non-zero drift, and persistent but mean-reverting volatility dynamics, we show that sign forecastability is to be expected and does not violate traditional notions of market efficiency. We also consider higher-order conditional moment dynamics, showing for example that time-varying skewness can yield sign forecastability even with zero drift and no volatility dynamics.

We show analytically that the probability forecast of a positive return is most sensitive to changes in volatility when volatility is at an intermediate level. In exceptionally tranquil times, the sign outcome is virtually certain (it will match the sign of the drift), and in exceptionally volatile times, the sign outcome is unpredictable and the conditional probability forecast will simply be 0.5.

Finally, we show in a realistically calibrated simulation exercise that sign forecastability is strongest at horizons of two or three months. At short horizons the drift will be close to zero, and at long horizons volatility dynamics are largely absent; thus at neither end of the spectrum will one find sign forecastability. But at intermediate horizons, sign forecastability can be quite high. Because the sign forecastability arises from volatility dynamics, it is of a nonlinear nature, and a simple linear investigation of the autocorrelation of the signs is unlikely to reveal sign forecastability even if it is present.

Let us conclude with a brief discussion of a promising direction for future research: the exploitation of sign forecastability for improved portfolio management. Preliminary results indicate that the probability of a positive portfolio return can be significantly increased by exploiting the volatility dynamics in individual securities, even when assuming a constant, but non-zero, mean. One can essentially time the market by exploiting conditional variance rather than conditional mean dynamics, as in Fleming, Kirby and Ostdiek (2001a, 2001b).¹⁴ 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, say, simple Sharpe ratios, such as Stutzer’s (2001) portfolio performance index.

¹⁴ Earlier work on market timing based on conditional mean dynamics includes Merton (1981), Henriksson and Merton (1981), Cumby and Modest (1987), and Whitelaw (1997).

References

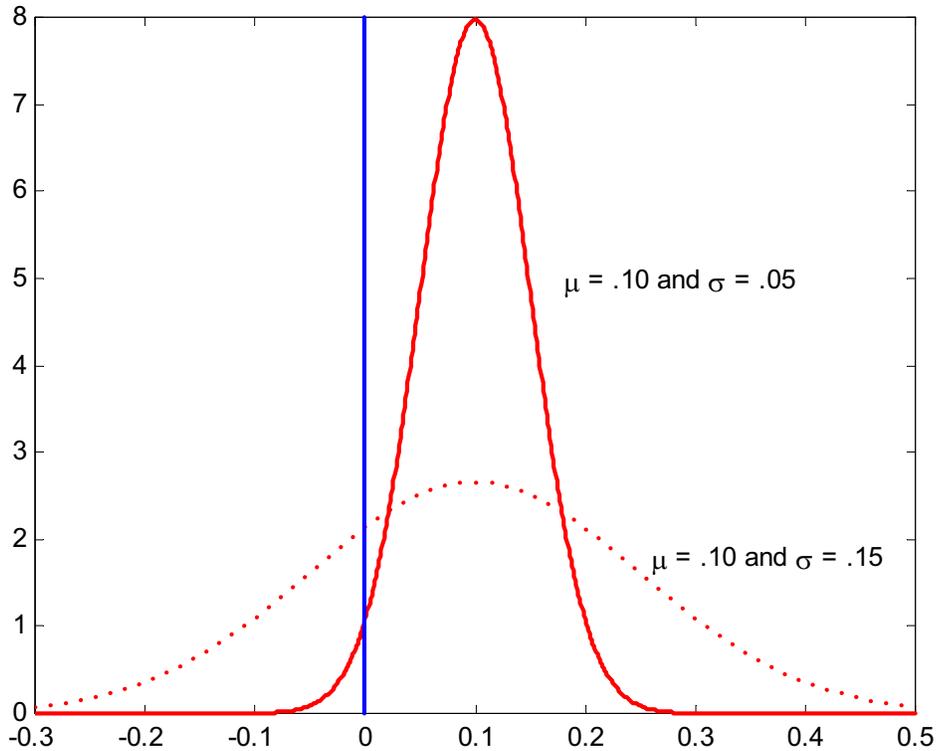
- Alizadeh, Sasan, Michael W. Brandt and Francis X. Diebold, 2002, Range-based estimation of stochastic volatility models, *Journal of Finance*, forthcoming.
- Andersen, Torben G., Luca Benzoni, and Jesper Lund, 2000, Estimating jump-diffusions for equity returns, Manuscript, Kellogg Graduate School of Management.
- Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold, 2002, "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, 2001a, 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, 2001b, Modeling and forecasting realized volatility, Manuscript, Northwestern University, Duke University, and University of Pennsylvania. NBER Working Paper No. 8160.
- 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.
- 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.
- 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.
- 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.
- 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 Gourieroux, and Joanna Jasiak, 2001, Compound autoregressive models, Manuscript, York University.
- El Babsiri, Mohamed, and Jean-Michel Zakoian, 2001, Contemporaneous asymmetry in GARCH processes, *Journal of Econometrics* 101, 257-294.
- Elliott, Graham, and Takatoshi Ito, 1995, 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 impact of jumps in volatility and returns, Manuscript, Graduate School of Business, University of Chicago.
- 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.
- 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, Kirby, Chris, and Ostdiek, Barbara, 2001a, The economic value of volatility timing, *Journal of Finance* 56, 329-352.
- Fleming, Jeff, Kirby, Chris, and Ostdiek, Barbara, 2001b, The Economic value of volatility timing using realized volatility, Manuscript.
- Fontnouvelle, Patrick de, 2000, Information dynamics in financial markets, *Macroeconomic Dynamics* 4, 139-169.
- 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 Kluger, Brian, 1987, Safety first: an alternative performance measure, *Journal of Portfolio Management* 13, 34-40.
- Hansen, B.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.
- 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.
- 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.
- 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.

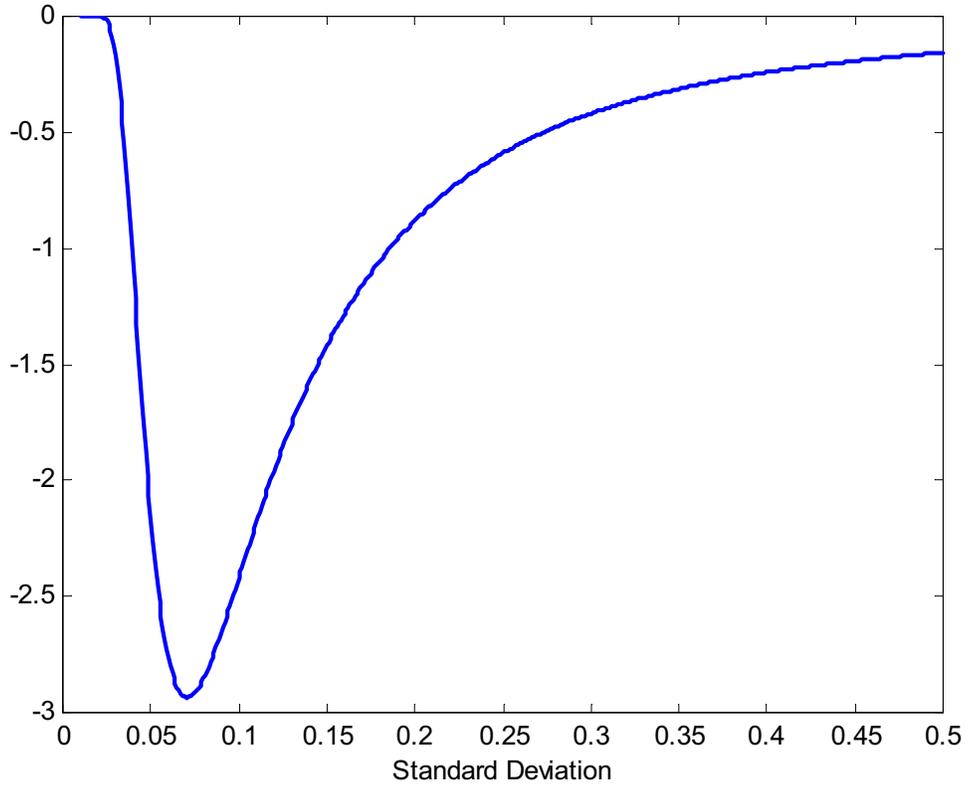
- Levich, Richard M., 1998, *International Financial Markets*. New York: McGraw-Hill.
- Lo, Andrew W., and A. Craig MacKinlay, 1999, *A Non-Random Walk Down Wall Street*. Princeton, NJ: Princeton University Press.
- Mankiw, N. Gregory, David Romer, and Matthew D. Shapiro, 1991, Stock market forecastability 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.
- 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, 1951-72.
- Pesaran, M. Hashem, and Allan G. Timmermann, 1995, Predictability of stock returns: robustness and economic significance, *Journal of Finance* 50, 1201-1228.
- 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, M., 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.
- 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 Probability of a Positive Return Depends on Volatility



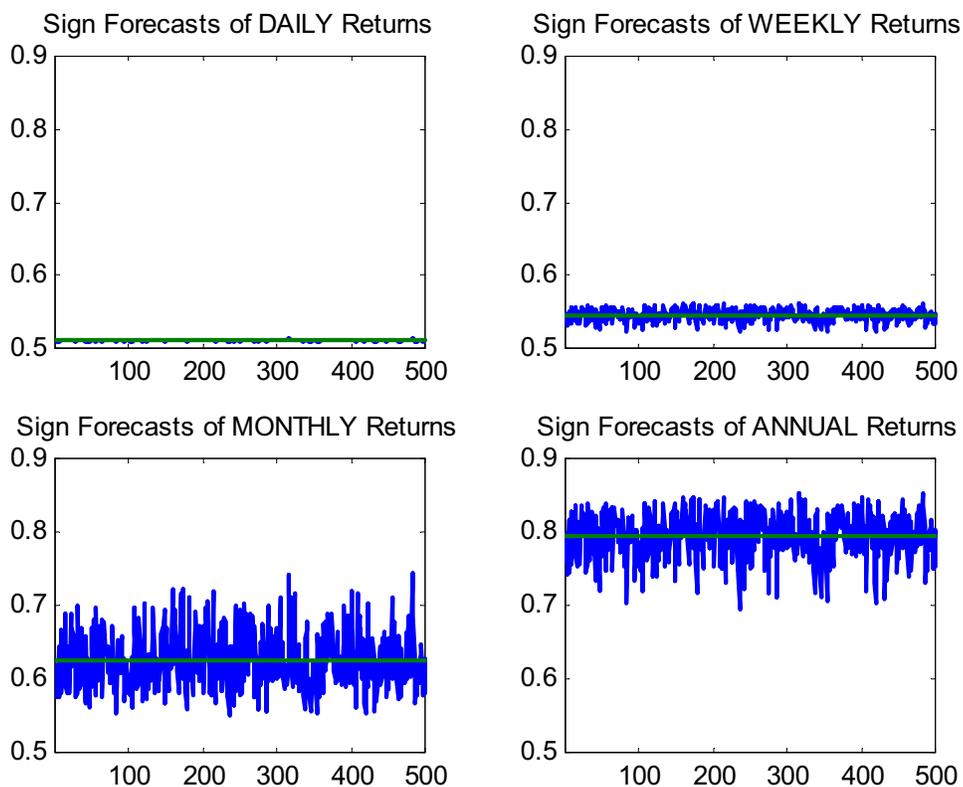
Notes to figure: We show two Gaussian return distributions, each with a mean of ten percent. The first has a standard deviation of five percent and a corresponding probability of positivity of 98 percent (the area to the right of zero under the peaked function). The second has a standard deviation of fifteen percent and a corresponding probability of positivity of only seventy-five percent (the area to the right of zero under the flatter function).

Figure 2
First Derivative of the Probability Forecast
with Respect to the Standard Deviation



Notes to figure: We plot \mathfrak{R} , the derivative of the probability of a positive return as a function of return volatility. We assume Gaussian returns with a mean of 0.10.

Figure 3
Conditional Probability of a Positive Sign at Various Horizons



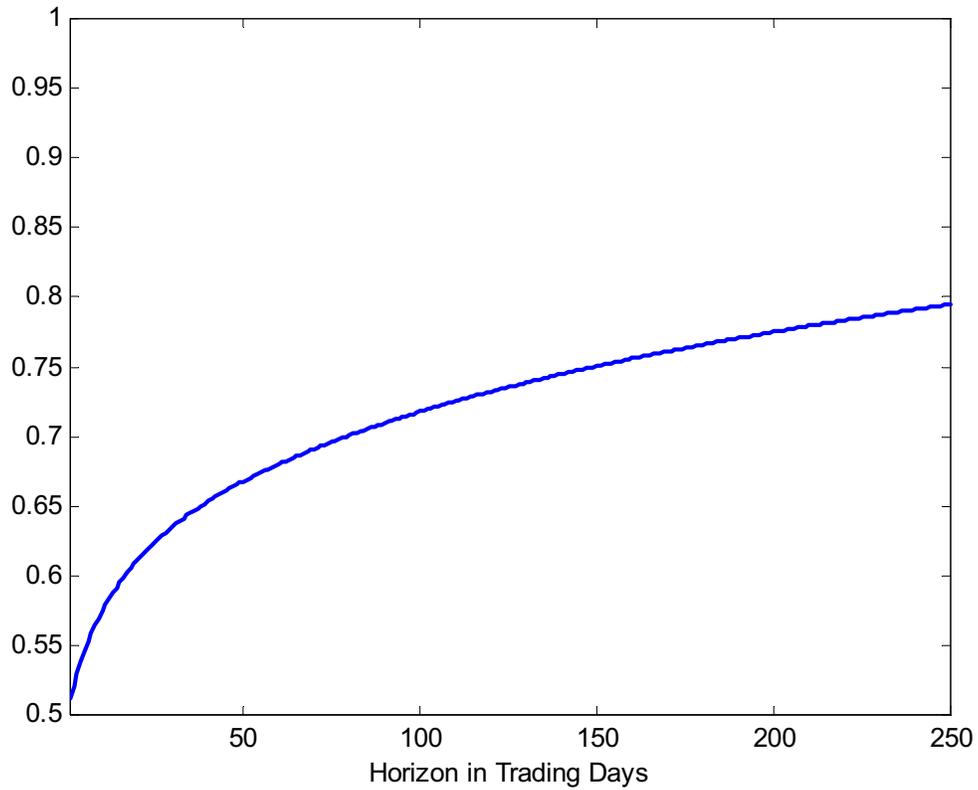
Notes to figure: We simulate asset prices from the stochastic volatility model parameterized as

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

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

with $\text{corr}(dz_1, dz_2) = -0.5$. See text for details. We then calculate 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.

Figure 4
Unconditional Probability of a Positive Return Across Horizons



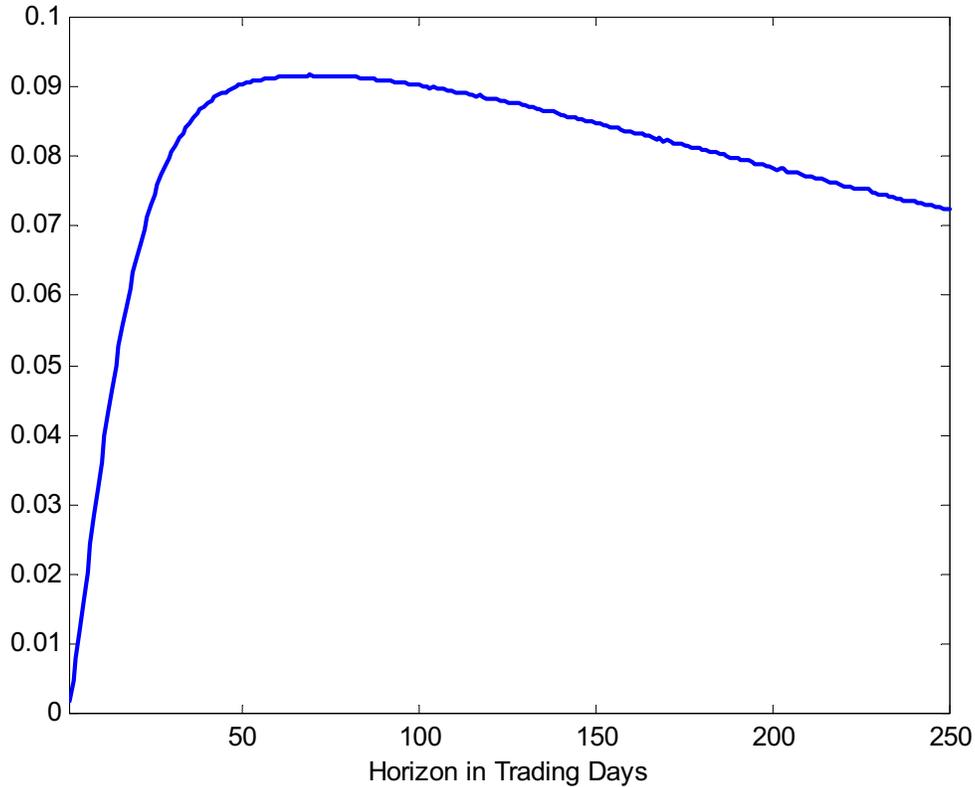
Notes to figure: We simulate asset prices from the stochastic volatility model parameterized as

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

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

with $\text{corr}(dz_1, dz_2) = -0.5$. See text for details. We then calculate the unconditional probability of a positive return at horizons ranging from one through 250 trading days (one year). We calculate the unconditional probabilities as simple averages of conditional probabilities, which we calculate using the inverse characteristic function technique described in the text.

Figure 5
Correlation Between Sign Forecasts and Realizations
Various Horizons, Benchmark Parameters



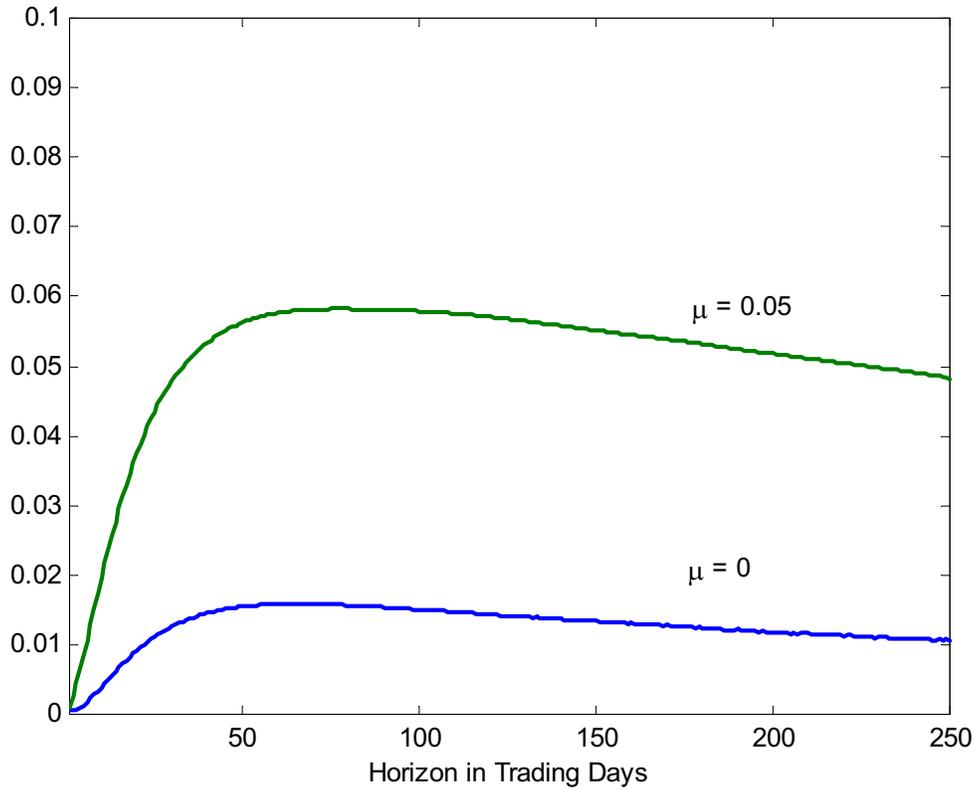
Notes to figure: We simulate asset prices from the stochastic volatility model parameterized as

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

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

with $\text{corr}(dz_1, dz_2) = -0.5$. We then 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 Return Drift



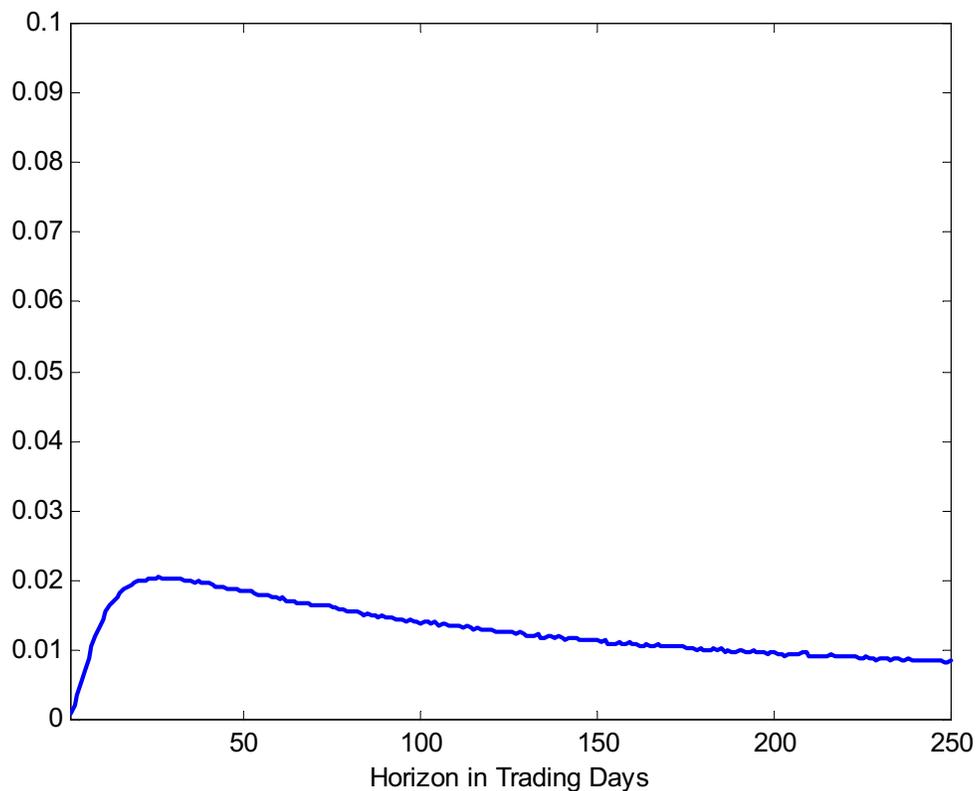
Notes to figure: We simulate asset prices from the stochastic volatility model

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

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

with $\text{corr}(dz_1, dz_2) = -0.5$, with drift of 0.05 (top line) and 0.00 (bottom line). See text for details. 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 7
Correlation between Sign Forecast and Realization
Various Horizons, Rapid Reversion of Volatility to its Mean



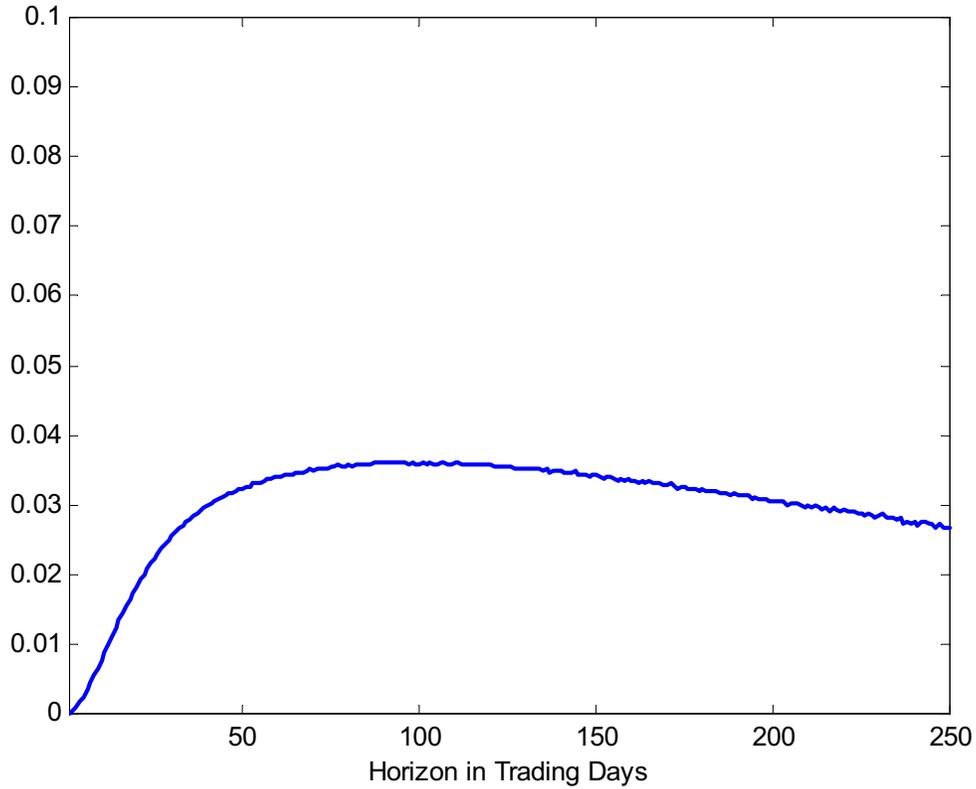
Notes to figure: We simulate asset prices from the stochastic volatility model parameterized as

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

$$d\sigma^2(t) = 10(0.015 - \sigma^2(t))dt + 0.15\sigma(t)dz_2,$$

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 using a large number of realizations, making use of the quasi-analytic result in the text.

Figure 8
First Autocorrelation of Return Signs
Various Horizons, Benchmark Parameters



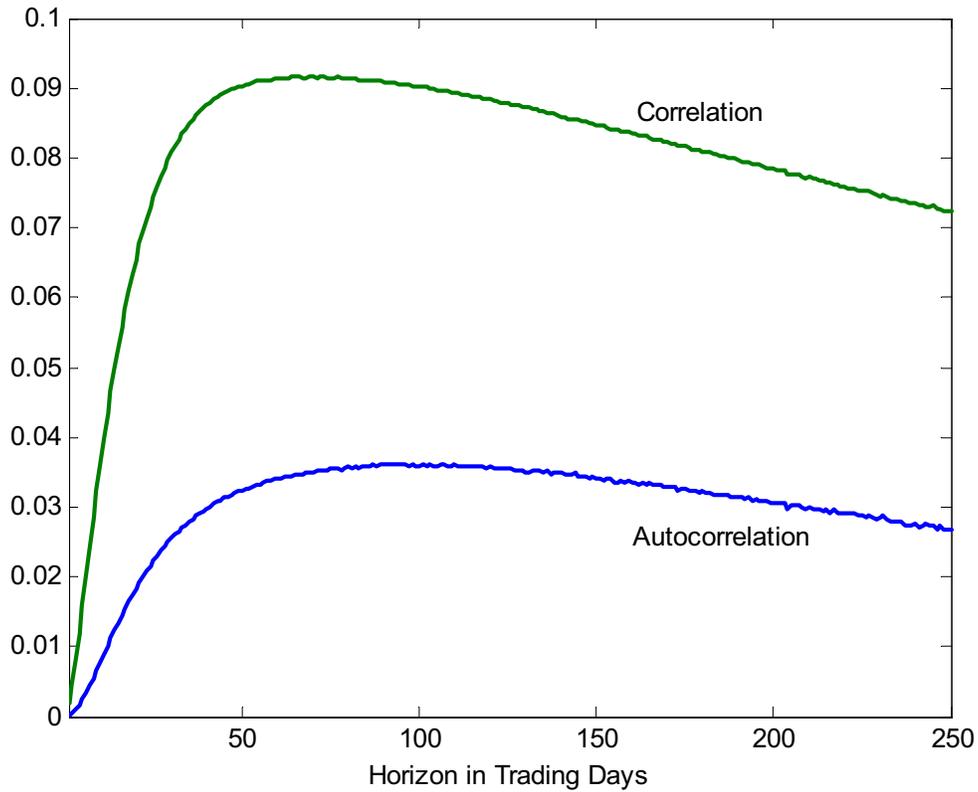
Notes to figure: We simulate asset prices from the stochastic volatility model parameterized as

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

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

with $\text{corr}(dz_1, dz_2) = -0.5$. We then construct an indicator sequence of return signs for each horizon. We calculate the sample autocorrelation from a long simulated sequence of returns, using the quasi-analytic result in the text.

Figure 9
Correlation between Sign Forecast and Realization
and First Autocorrelation of Return Signs
Various Horizons, Benchmark Parameters



Notes to figure: The correlation between the sign forecast and realization is as in Figure 5 and the first autocorrelation of the sign sequence is as in Figure 8.