Realized Beta:
Persistence and Predictability*

Torben G. Andersen[a], Tim Bollerslev[b], Francis X. Diebold[c] and Jin Wu[d]

January 2003
This Draft/Print: December 19, 2003

Abstract: A large literature over several decades reveals both extensive concern with the question of time-varying betas and an emerging consensus that betas *are* in fact time-varying, leading to the prominence of the conditional CAPM. Set against that background, we assess the dynamics in realized betas, vis-à-vis the dynamics in the underlying realized market variance and individual equity covariances with the market. Working in the recently-popularized framework of realized volatility, we are led to a model of nonlinear fractional cointegration: although realized variances and covariances are very highly persistent and fractionally integrated, realized betas, which are simple nonlinear functions of those realized variances and covariances, are less persistent, and arguably best modeled as a standard stationary $I(0)$ process. We conclude by drawing implications for asset pricing and portfolio management.

Key Words: quadratic variation and covariation, realized volatility, realized variance, realized covariance, capital asset pricing model, CAPM, asset pricing, high-frequency data, equity, stock, long memory, nonlinear fractional cointegration, continuous-time methods

JEL Codes: C1, G1

[*] This work was supported by the National Science Foundation and the Guggenheim Foundation. For useful discussion we thank seminar participants at the University of Pennsylvania and the London School of Economics' Financial Market Group, as well as Andrew Ang, Michael Brandt, Graham Elliott, Eric Ghysels, Rich Lyons, Norman Swanson, and Mark Watson.

[a] Department of Finance, Kellogg School, Northwestern University, and NBER, t-andersen@kellogg.northwestern.edu

[b] Departments of Economics and Finance, Duke University, and NBER, boller@econ.duke.edu

[c] Departments of Economics, Finance and Statistics, University of Pennsylvania, and NBER, fdiebold@sas.upenn.edu

[d] Department of Economics, University of Pennsylvania, jinw@ssc.upenn.edu

1. Introduction

One of the key insights of asset pricing theory is also one of the simplest: only systematic risk should be priced. Perhaps not surprisingly, however, there is disagreement as to the sources of systematic risk. In the one-factor capital asset pricing model (CAPM), for example, systematic risk is determined by covariance with the market (Sharpe, 1963; Lintner, 1965a, b), whereas in more elaborate pricing models additional empirical characteristics, such as firm size and book-to-market, are seen as proxies for another set of economy-wide risk factors that are also priced (Fama and French, 1993).[1]

As with almost all important scientific models, the CAPM has been subject to substantial criticism (e.g., Fama and French, 1992). Nevertheless, to paraphrase Mark Twain, the reports of its death are greatly exaggerated. In fact, the one-factor CAPM remains alive and well at the frontier of both academic research and industry applications. In part this is because recent work reveals that – despite its warts and wrinkles – it often works well, whether in traditional incarnations (e.g., Ang and Chen, 2003) or more novel variants (e.g., Cohen, Polk and Voulteenaho, 2002; Campbell and Vuolteenaho, 2002), and in part it is because competing multi-factor pricing models, although providing improved statistical fit, involve factors whose economic interpretations in terms of systematic risks remain unclear.[2] Moreover, the stability of many empirically-motivated asset pricing relationships appears tenuous when explored with true out-of-sample data, suggesting an element of data mining.

In this paper, then, we study the one-factor CAPM, which remains central to financial economics nearly a half century after its introduction. A key question within this setting is whether a stock's systematic risk, as assessed by its correlation with the market, is constant over time. Put bluntly, the question is "Are betas constant?" And if not, what is the degree of predictable time-variation in betas? The evolution of a large literature over several decades reveals both extensive concern with this question and, we contend, an eventual consensus that betas *are* likely time-varying.[3]

Several pieces of evidence support our contention. First, leading texts echo it. For example, Huang and Litzenberger (1988) assert that "It is unlikely that risk premiums and betas on individual assets are stationary over time" (p. 303).

Second, explicitly dynamic betas are often modeled nonstructurally via time-varying parameter regression, in a literature tracing at least to the early "return to normality" model of Rosenberg (1973), as

[1] The Roll (1977) critique is also relevant. That is, even if we somehow knew what factor(s) should be priced, it is not clear that the factor(s) as measured in practice would correspond to the factor(s) required by the theory.

[2] On the difficulty of interpreting additional factors in terms of systematic risk, see Keim and Hawawini (1999).

[3] There are of course qualifications, notably Ghysels (1998), which we discuss subsequently.

implemented in the CAPM by Schaefer, Brealey, Hodges and Thomas (1975).

Third, even in the absence of explicit allowance for time-varying betas, the CAPM is typically estimated using moving estimation windows, typically of five to ten years, presumably to guard against beta variation (e.g., Fama, 1976; Campbell, Lo and MacKinlay, 1997).

Fourth, theoretical and empirical inquiries in asset pricing are often undertaken in conditional, as opposed to unconditional, frameworks, the essence of which is to allow for time-varying betas, presumably because doing so is viewed as necessary for realism. The motivation for the conditional CAPM comes from at least two sources. Financial economic considerations suggest that betas may vary with conditioning variables, an idea developed theoretically and empirically in a large literature that includes, among many others, Dybvig and Ross (1985), Hansen and Richard (1987), Ferson, Kandel and Stambaugh (1987), Ferson and Harvey (1991), Jagannathan and Wang (1996), and Wang (2003).[4] From a different perspective, the financial econometric volatility literature (see Andersen, Bollerslev and Diebold, 2004, for a survey) has provided extensive evidence of wide fluctuations and high persistence in asset market conditional variances, and in individual equity conditional covariances with the market. Thus, even from a purely statistical viewpoint, market betas, which are ratios of time-varying conditional covariances and variances, might be expected to display persistent fluctuations, as in Bollerslev, Engle and Wooldridge (1988). In fact, unless some special cancellation occurs – in a way that we formalize – betas would inherit the persistence features that are so vividly present in their constituent components.

Set against this background, we seek to assess the dynamics in betas vis-à-vis the widely documented persistent dynamics in the underlying variance and covariances. We proceed as follows. In section 2 we sketch the framework, both economic and econometric, in which our analysis is couched. In section 3 we present the empirical results with an emphasis on analysis of persistence and predictability. In section 4 we consider the incorporation of high-frequency data. In section 5 we formally assess the uncertainty in our beta estimates. In section 6 we offer summary, conclusions and directions for future research.

2. Theoretical Framework

Our approach has two key components. First, in keeping with the recent move toward nonparametric volatility measurement, we cast our analysis within the framework of realized variances and covariances, or equivalently, empirical quadratic variation and covariation. That is, we do not entertain a null hypothesis of period-by-period constant betas, but instead explicitly allow for continuous evolution in the betas. The realized beta measures are (continuous-record) consistent for realizations of the underlying ratio between the

---

[4] The idea of conditioning in the CAPM is of course not unrelated to the idea of multi-factor pricing mentioned earlier.

integrated stock and market return covariance and the integrated market variance.[5] Second, we work in a flexible econometric framework that allows for, without imposing, fractional integration and/or cointegration between the market variance and individual equity covariances with the market.

Realized Quarterly Variances, Covariances, and Betas

We provide estimates of quarterly betas, based on nonparametric realized quarterly market variances and individual equity covariances with the market. The quarterly frequency is appealing from a substantive financial economic perspective, and it also provides a reasonable balance between efficiency and robustness to microstructure noise. Specifically, we produce our quarterly estimates using underlying daily returns, as in Schwert (1989), so that the sampling frequency is quite high relative to the quarterly horizon of interest, yet low enough so that contamination by microstructure noise is not a serious concern for the highly liquid stocks that we study. The daily frequency further allows us to utilize a long sample period in our study which is not feasible for higher return sampling frequencies.

Suppose that the logarithmic $N\times 1$ vector price process, p_t, follows a multivariate continuous-time stochastic volatility diffusion,

$$dp_t \;=\; \mu_t\, dt \;+\; \Omega_t\, dW_t\,, \tag{1}$$

where W_t denotes a standard N-dimensional Brownian motion, and both the process for the $N\times N$ positive definite diffusion matrix, Ω_t, and the N-dimensional instantaneous drift, μ_t, are strictly stationary and jointly independent of the W_t process. For our purposes it is helpful to think of the N'th element of p_t as containing the log price of the market and the i'th element of p_t as containing the log price of the i'th individual stock included in the analysis, so that the corresponding covariance matrix contains both the market variance, say $\sigma^2_{M,t} = \Omega_{(NN),t}$, and the individual equity covariance with the market, $\sigma_{iM,t} = \Omega_{(iN),t}$. Then, conditional on the sample path realization of μ_t and Ω_t, the distribution of the continuously compounded h-period return, $r_{t+h,h} \equiv p_{t+h} - p_t$, is

$$r_{t+h,h} \;/\; \sigma\{\mu_{t+\tau}, \Omega_{t+\tau}\}_{\tau=0}^{h} \;\sim\; N\!\left(\int_0^h \mu_{t+\tau}\, d\tau,\; \int_0^h \Omega_{t+\tau}\, d\tau\right), \tag{2}$$

where $\sigma\{\mu_{t+\tau}, \Omega_{t+\tau}\}_{\tau=0}^{h}$ denotes the σ-field generated by the sample paths of $\mu_{t+\tau}$ and $\Omega_{t+\tau}$ for $0\leq\tau\leq h$.[6] The integrated diffusion matrix $\int_0^h \Omega_{t+\tau}\, d\tau$ therefore provides a natural measure of the true latent h-period

---

[5] The underlying theory and related empirical strategies are developed in Andersen, Bollerslev, Diebold and Labys (2001, 2003), Andersen, Bollerslev, Diebold and Ebens (2001), and Barndorff-Nielsen and Shephard (2003). Here we sketch the basics; for a more rigorous treatment in the framework of special semimartingales, see the survey and unification by Andersen, Bollerslev and Diebold (2004).

[6] We normalize such that the unit time interval, or $h=1$, represents one trading day.

volatility.[7] The requirement that the innovation process, W_t, is independent of the drift and diffusion processes is unfortunately rather strict and precludes, for example, the asymmetric relations between return innovations and volatility captured by the so-called leverage or volatility feedback effects. However, from the results in Meddahi (2002) and Andersen, Bollerslev and Meddahi (2004), we know that the continuous-record asymptotic distribution theory for the realized covariation continues to provide an excellent approximation for empirical high-frequency realized volatility measures. As such, even if the conditional return distribution result (2) does not apply in full generality, the evidence presented below, based exclusively on the realized volatility measures, is also valid in the presence of asymmetries in return innovation-volatility relations.

By the theory of quadratic variation, we have that under weak regularity conditions, and regardless of the presence of leverage or volatility feedback effects, that

$$\Sigma_{j=1,\ldots,[h/\Delta]} r_{t+j\cdot\Delta,\Delta} \cdot r'_{t+j\cdot\Delta,\Delta} - \int_0^h \Omega_{t+\tau} d\tau \rightarrow 0, \tag{3}$$

almost surely for all t as the sampling frequency of the returns increases, or $\Delta \rightarrow 0$. Thus, by summing sufficiently finely-sampled high-frequency returns, it is possible to construct ex-post *realized* volatility measures for the integrated latent volatilities that are asymptotically free of measurement error.[8] This contrasts sharply with the common use of the cross-product of the h-period returns, $r_{t+h,h} \cdot r'_{t+h,h}$, as a simple ex-post (co-)variability measure. Although the squared return (innovation) over the forecast horizon provides an unbiased estimate for the realized integrated volatility, it is an extremely noisy estimator, and predictable variation in the true latent volatility process is typically dwarfed by measurement error.[9] Moreover, for longer horizons any conditional mean dependence will tend to contaminate this variance measure. In contrast, as the sampling frequency is lowered the impact of the drift term vanishes, thus effectively annihilating the mean.

These assertions remain valid if the underlying continuous time process in equation (1) contains jumps, so long as the price process is a special semimartingale, which will hold if it is arbitrage-free. Of course, in this case the limit of the summation of the high-frequency returns will involve an additional jump component, but the interpretation of the sum as the realized h-period return volatility remains intact.

Finally, with the realized market variance and realized covariance between the market and the

---

[7] This notion of integrated volatility already plays a central role in the stochastic volatility option pricing literature, in which the price of an option typically depends on the distribution of the integrated volatility process for the underlying asset over the life of the option. See, for example, the well-known contribution of Hull and White (1987).

[8] Consider the simple case of univariate discretely-sampled zero-mean i.i.d. Gaussian returns; i.e., $N=1$, $\mu_t=0$, and $\Omega_t = \sigma^2$. It follows by standard arguments that $E(h^{-1}\cdot\Sigma_{j=1,\ldots,[h/\Delta]} r^2_{t+j\cdot\Delta,\Delta}) = \sigma^2$, while $Var(h^{-1}\cdot\Sigma_{j=1,\ldots,[h/\Delta]} r^2_{t+j\cdot\Delta,\Delta}) = (\Delta/h)\cdot 2\cdot\sigma^4 \rightarrow 0$, as $\Delta \rightarrow 0$.

[9] In empirically realistic situations, the variance of $r_{t+1,1} r'_{t+1,1}$ is easily twenty times the variance of the true daily integrated volatility, $\int_0^1 \Omega_{t+\tau} d\tau$; see Andersen and Bollerslev (1998) for some numerical results along these lines.

individual stocks in hand, we can readily define and empirically construct the individual equity betas. Towards that end, we introduce some formal notation. Using an initial subscript to indicate the corresponding element of a vector, we denote the realized market volatility by

$$\hat{v}^2_{M,t,t+h} \quad = \quad \Sigma_{j=1,\ldots,[h/\Delta]} \; r^2_{(N),t+j\cdot\Delta,\Delta} \,, \tag{4}$$

and we denote the realized covariance between the market and the ith individual stock return by

$$\hat{v}_{iM,t,t+h} \quad = \quad \Sigma_{j=1,\ldots,[h/\Delta]} \; r_{(i),t+j\cdot\Delta,\Delta} \cdot r_{(N),t+j\cdot\Delta,\Delta} \,. \tag{5}$$

We then define the associated realized beta as

$$\hat{\beta}_{i,t,t+h} \quad = \quad \frac{\hat{v}_{iM,t,t+h}}{\hat{v}^2_{M,t,t+h}} \,. \tag{6}$$

Under the assumptions invoked for equation (1), this empirical beta measure is consistent for the true underlying realized beta in the following sense,

$$\hat{\beta}_{i,t,t+h} \quad \rightarrow \quad \beta_{i,t,t+h} \quad = \quad \frac{\int_0^h \Omega_{(iN),t+\tau}\, d\tau}{\int_0^h \Omega_{(NN),t+\tau}\, d\tau} \,. \tag{7}$$

almost surely for all t as the sampling frequency increases, or $\Delta \rightarrow 0$.

A number of comments are in order. First, the integrated return covariance matrix, $\int_0^h \Omega_{t+\tau} d\tau$, is treated as stochastic, so both the integrated market variance and the integrated covariances of individual equity returns with the market over $[t, t+h]$ are ex-ante, as of time t, unobserved and governed by a non-degenerate (and potentially unknown) distribution. Moreover, the covariance matrix will generally vary continuously and randomly over the entire interval, so the integrated covariance matrix should be interpreted as the average realized covariation among the return series. Second, equation (3) makes it clear that the realized market volatility in (4) and the realized covariance in (5) are continuous-record consistent estimators of the (random) realizations of the underlying integrated market volatility and covariance. Thus, as a corollary, the realized beta will be consistent for the integrated beta, as stated in (7). Third, the general representation here encompasses the standard assumption of a constant beta over the measurement or estimation horizon, which is attained for the degenerate case of the Ω_t process being constant throughout each successive h-period measurement interval, or $\Omega_t = \Omega$. Fourth, the realized beta estimation procedure in equations (4)-(6) is implemented through a simple regression (without a constant term) of individual high-frequency stock returns on the corresponding market return. Nonetheless, the interpretation is very different from a standard regression, as the OLS point estimate now represents a consistent estimator of the ex-post realized regression coefficient obtained as the ratio of unbiased estimators of the average realized covariance and the realized market variance. The associated continuous record asymptotic theory developed by Barndorff-Nielsen and Shephard (2003) explicitly recognizes the diffusion setting underlying the regression and hence facilitates the construction of standard errors for our beta estimators.

Nonlinear Fractional Cointegration: A Common Long-Memory Feature in Variances and Covariances

The possibility of common persistent components is widely recognized in modern multivariate time-series econometrics. It is also important for our analysis, because there may be common persistence features in the underlying variances and covariances from which betas are produced.

The idea of a common feature is a simple generalization of the well-known cointegration concept. If two variables are integrated but there exists a function *f* of them that is not, we say that they're cointegrated, and we call *f* the conintegrating function. More generally, if two variables have property X but there exists a function of them that does not, we say that they have common feature X. A key situation is when X corresponds to *persistence*, in which case we call the function of the two variables that eliminates the persistence the *copersistence function*. It will prove useful to consider linear and nonlinear copersistence functions in turn.

Most literature focuses on linear copersistence functions. The huge cointegration literature pioneered by Granger (198*) and Engle and Granger (1987) deals with linear common long-memory $I(1)$ persistence features. The smaller copersistence literature started by Engle and Kozicki (1993) deals mostly with linear common short-memory $I(0)$ persistence features. The idea of fractional cointegration, suggested by Engle and Granger (1987) and developed by Cheung and Lai (1993) and Robinson and Marinucci, (2001), among others, deals with linear common long-memory $I(d)$ persistence features, $0<d<1/2$.

Our interest is closely related but different. First, it centers on *nonlinear* copersistence functions, because betas are ratios. There is little literature on nonlinear common persistence features, although they are implicitly treated in Granger (1995). We will be interested in nonlinear common long-memory $I(d)$ persistence features, $0<d<1/2$, effectively corresponding to nonlinear fractional cointegration.[10]

Second, we are interested primarily in the case of *known* cointegrating relationships. That is, we may not know whether a given stock's covariance with the market is fractionally cointegrated with the market variance, but if it is, then there is a good financial economic reason (i.e., the CAPM) to suspect that the cointegrating function is the *ratio* of the covariance to the variance. This provides great simplification. In the integer-cointegration framework with known cointegrating vector under the alternative, for example, one could simply test the cointegrating combination for a unit root, or test the significance of the error-

[10] One could of course attempt a linear cointegration approach by taking logs of the realized volatilities and covariances, but there is no theoretical reason to expect all covariances to be positive, and our realized covariance measures are indeed sometimes negative, making logarithmic transformations problematic.

correction term in a complete error-correction model, as in Horvath and Watson (1995). We proceed in analogous fashion, examining the integration status (generalized to allow for fractional integration) of the realized market variance, realized individual equity covariances with the market, and realized market betas.

Our realized beta series are unfortunately relatively short compared to the length required for formal testing and inference procedures regarding (fractional) cointegration, as the fractional integration and cointegration estimators proposed by Geweke-Porter Hudak (1983), Robinson and Marinucci (2001) and Andrews and Guggenberger (2003) tend to behave quite erratically in small samples. In addition, there is considerable measurement noise in the individual beta series so that influential outliers may have a detrimental impact on our ability to discern the underlying dynamics. Hence we study the nature of the long range dependence and short-run dynamics in the realized volatility measures and realized betas through intentionally less formal but more informative graphical means and via some robust procedures that utilize the joint information across many series, to which we now turn.

3. Empirical Analysis

We examine primarily on realized quarterly betas constructed from daily returns. We focus on the dynamic properties of market betas vis-à-vis the dynamic properties of their underlying covariance and variance components. We quantify the dynamics in a number of ways, including explicit measurement of the degree of predictability in the tradition of Granger and Newbold (1986). Finally, we further explore the dynamic nature of realized betas and their components by studying monthly betas constructed from high-frequency intraday returns. Here, the time span covered is shorter, as the intraday data are limited, but the improvement in the quality of the measurement still allows a precise characterization of the dynamics.

Dynamics of Quarterly Realized Variance, Covariances and Betas

We will focus primarily, but not exclusively, on realized quarterly betas constructed from daily returns obtained from the Center for Research in Security Prices from July 1962 to September 1999. We take the market return $r_{m,t}$ to be the Dow Jones Industrial Average (DJIA) of thirty industrial stocks, and we study the subset of twenty-five DJIA stocks as of March 1997 with complete data from July 2, 1962 to September 17, 1999, as detailed in Table 1. We then construct quarterly realized DJIA variances, individual equity covariances with the market, and betas, 1962:3 - 1999:3 (149 observations).

In Figure 1 we provide a time series plot of the quarterly realized market variance, with fall 1987 included (top panel) and excluded (bottom panel). It is clear that the realized variance is quite persistent and, moreover, that the fall 1987 volatility shock is unlike any other ever recorded, in that volatility reverts to its mean almost instantaneously following that shock. In addition, our subsequent computation of asymptotic standard error bands reveals that the uncertainty associated with this particular quarterly beta estimate is enormous, to the point of rendering it entirely uninformative. In sum, it is an exceptional outlier with potential enormous influence on the analysis, and it is measured with a huge amount of imprecision. Hence, following many other authors, we drop the fall 1987 observation from this point onward.

In Figures 2 and 3 we show time series plots of the twenty-five quarterly realized covariances and realized betas.[11] Like the realized variance, the realized covariances appear highly persistent. The realized betas, in contrast, appear noticeably less persistent. This impression is confirmed by the statistics presented in Table 2: the mean Ljung-Box Q-statistic (through displacement 12) is 87 for the realized covariance, but only 53 for the realized beta.[12]

The impression of reduced persistence in realized betas relative to realized covariances is also confirmed by the sample autocorrelation functions shown in Figure 4 (market variance), Figure 5 (covariances with the market), and Figure 7 (betas). In particular, the sample autocorrelation functions of the betas generally decay more quickly than those for the covariances, as clearly revealed in the median sample autocorrelation functions of the covariances and the betas shown in Figures 6 and 8, respectively.

The work of Andersen, Bollerslev, Diebold and Ebens (2001) and Andersen, Bollerslev, Diebold and Labys (2003), as well as that of numerous other authors, indicates that asset return volatilities are well-described by a pure fractional noise process, typically with the degree of integration around $d \approx .4$.[13] That style of analysis is mostly conducted on high-frequency data. Very little work has been done on long memory in equity variances, market covariances, and market betas at the quarterly frequency, and it is hard to squeeze accurate information about *d* directly from the fairly limited quarterly sample. It is well-known, however, that if a flow variable x is *I(d),* then it remains *I(d)* under temporal aggregation. Hence, we can use the results of analyses of high-frequency data, such as Andersen, Bollerslev, Diebold and Labys (2003), to help us analyze the quarterly data. Inspired by their results we settle on the point estimate of $d \approx .42$.

In Figure 9 we graph the sample autocorrelations of the quarterly realized market variance prefiltered by $(1-L)^{.42}$, in Figure 10 we show the sample autocorrelations of quarterly realized covariances prefiltered by $(1-L)^{.42}$, and in Figure 12 we display the sample autocorrelations of quarterly realized betas prefiltered by $(1-L)^{.42}$. All dynamics in the realized variance and covariances are eliminated by filtering with $(1-L)^{.42}$, indicating that the pure fractional noise process with $d = .42$ is indeed a good approximation

[11] The quarterly realized variance, covariances and betas are computed from slightly different numbers of observations due to the different numbers of trading days across the quarters.

[12] Note also that the Dickey-Fuller statistics indicate that unit roots are not present in the market variance, individual equity covariances with the market, or market betas, despite their persistent dynamics.

[13] A partial list of references not written by the present authors includes Breidt and de Lima (1998), Comte and Renault (1998), Harvey (1998), and Robinson (2001), as well as many of the earlier papers cited in Baillie (1996).

to their dynamics. Interestingly, however, filtering the realized *betas* with $(1-L)^{.42}$ appears to produce *overdifferencing*, as evidenced by the fact that the first autocorrelation of the betas fractionally differenced using $(1-L)^{.42}$ often is negative. Compare, in particular, Figure 11's median sample autocorrelation function for the realized covariances prefiltered by $(1-L)^{.42}$ to Figure 13's median sample autocorrelation function for the realized betas prefiltered by $(1-L)^{.42}$. The difference is striking in the sense that the first autocorrelation coefficient is negative and much larger than the remainder in the latter case.

If fractional differencing of the realized betas by $(1-L)^{.42}$ may be "too much," then the question naturally arises as to how much differencing is "just right." Some experimentation revealed that differencing the betas by $(1-L)^{.20}$ was often adequate for elimination of dynamics. However, for short samples it is almost impossible to distinguish low order fractional integration from persistent but strictly stationary dynamics. We are particularly interested in the latter alternative where the realized betas are $I(0)$. To explore this possibility, we fit simple $AR(p)$ processes to realized betas, with p selected by the AIC. We show the estimated roots in Table 3, all of which are indicative of covariance stationarity. In Figure 14 we show the sample autocorrelation functions of quarterly realized betas prefiltered by the estimated $AR(p)$ lag-operator polynomials, and in Figure 15 we show the median such autocorrelation function across all stocks. The autocorrelation functions are indistinguishable from those of white noise.

Taken as a whole, the results suggest that realized betas are integrated of noticeably lower order than are the market variance and the individual equity covariances with the market, corresponding to a situation of nonlinear fractional cointegration. $I(d)$ behavior, with $d \in [0,\ .25]$, appears accurate for betas, whereas the market variance and the individual equity covariances with the market are better approximated as $I(d)$ with $d \in [.35,\ .45]$. Indeed, there is little evidence against an assertion that betas are $I(0)$, whereas there is strong evidence against such an assertion for the variance and covariance components.

Predictability

Examination of the *predictability* of realized beta and its components provides a complementary perspective and additional insight. Granger and Newbold (1986) propose a natural measure of the predictability of covariance stationary series under squared-error loss, patterned after the familiar regression R^2,

$$G(j) = \frac{var(\hat{x}_{t+j,t})}{var(x_t)} = 1 - \frac{var(e_{t+j,t})}{var(x_t)}, \tag{8}$$

where j is the forecast horizon of interest, $\hat{x}_{t+j,t}$ is the optimal (i.e., conditional mean) forecast, and

$e_{t+j,t} = x_{t+j} - \hat{x}_{t+j,t}$. Diebold and Kilian (2001) define a generalized measure of predictability, building on the Granger-Newbold measure, as

$$P(L,\ \Omega,\ j,\ k) = 1 - \frac{E(L(e_{t+j,t}))}{E(L(e_{t+k,t}))}, \tag{9}$$

where L denotes the relevant loss function, Ω is the available univariate or multivariate information set, j is the forecast horizon of interest, and k is a long but not necessarily infinite reference horizon.

Regardless of the details, the basic idea of predictability measurement is simply to compare the expected loss of a short-horizon forecast to the expected loss of a very long-horizon forecast. The former will be much smaller than the latter if the series is highly predictable, as the available conditioning information will then be very valuable. The Granger-Newbold measure, which is the canonical case of the Diebold-Kilian measure (corresponding to $L(e) = e^2$, univariate Ω, and $k = \infty$) compares the 1-step-ahead forecast error variance to that of the ∞-step-ahead forecast error variance, i.e., the unconditional variance of the series being forecast (assuming that it is finite).

In what follows, we use predictability measures to provide additional insight into the comparative dynamics of the realized variances and covariances versus the realized betas. Given the earlier-discussed strong evidence of fractional integration in the realized market variance and covariances, we maintain the pure fractional noise process for the quarterly realized market variance and the quarterly realized covariances, namely *ARFIMA*(0, .42, 0). We then calculate the Granger-Newbold predictability $G(j)$ analytically, conditional upon the *ARFIMA*(0, .42, 0) dynamics, and we graph it in Figure 16 for ~j~=~1, ~..., ~7 quarters.[14] The graph starts out as high as .4 and decays only slowly over the first seven quarters. If the realized beta likewise follows a pure fractional noise process but with a smaller degree of integration, say *ARFIMA*(0, .20, 0), which we argued was plausible, then the implied predictability is much lower, as shown in Figure 17. As we also argued, however, the integration status of the realized betas is difficult to determine. Hence, for the realized betas we also compute Granger-Newbold predictability using an estimated *AR*(p) sieve approximation to produce estimates of $var(e_{t+j,t})$ and $var(x_t)$; this approach is valid regardless of whether the true dynamics are short-memory or long-memory. In Figure 18 we plot the beta predictabilities, which remain noticeably smaller and more quickly-decaying than the covariance predictabilities, as is further clarified by comparing the median beta predictability, shown in Figure 19, to the market variance and equity covariances predictability in Figure 16.

[14] Note that only one figure is needed, despite the many different realized covariances, because all $G(j)$ are identical, as all processes are assumed to be *ARFIMA*(0, .42, 0).

4. Monthly Realized Betas Computed Using High-Frequency Intraday Returns

In order to further examine the dynamics of betas, we also study the realized betas computed from high-frequency intraday returns. Our analysis is based on data from the TAQ (Trade And Quotation) database. The TAQ data files contain continuously recorded information on the trades and quotations for the securities listed on the New York Stock Exchange (NYSE), American Stock Exchange (AMEX), and the National Association of Security Dealers Automated Quotation system (NASDAQ). The database is published monthly, and has been available on CD-ROM from the NYSE since January 1993; we refer the reader to the corresponding data manual for a more complete description of the actual data and the method of data-capture. Our sample extends from January 4, 1993 until September 30, 1999, for a total of 1,366 trading days. A complete analysis based on all trades for all stocks, although straightforward conceptually, is infeasible in practice. We therefore restrict our analysis to the thirty DJIA firms, which also helps to ensure a reasonable degree of liquidity. Finally, the corresponding market return is obtained from the S&P500 futures index at the Chicago Board of Trade.

Although the DJIA stocks are among the most actively traded U.S. equities, the median inter-trade duration for all stocks across the full sample is 23.1 seconds, ranging from a low of 7 seconds for Merck & Co. Inc. (MRK) to a high of 54 seconds for United Technologies Corp. (UTX). As such, it is not practically feasible to push the continuous record asymptotics and the length of the observation interval Δ in equation (3) beyond this level. Moreover, because of the organizational structure of the market, the available quotes and transaction prices are subject to discrete clustering and bid-ask bounce effects. Such market microstructure features are generally not important when analyzing longer horizon interdaily returns but can seriously distort the distributional properties of high-frequency intraday returns; see, e.g., the textbook treatment by Campbell, Lo and MacKinlay (1997). Thus, following the analysis in Andersen and Bollerslev (1997), we rely on artificially constructed equally-spaced intraday returns.[15] It turns out that a 15-minute sampling frequency is sufficient to alleviate the non-synchronous trading problem so that the individual stock return series are approximately serially uncorrelated and orthogonal to the lagged and leading 15-minute market returns. With the daily transaction record extending from 9:30 EST until 16:00 EST, there are a total 26 fifteen-minute returns for each day, corresponding to $\Delta = 1/26 \approx 0.0385$ in the notation above, and thus about 570 underlying intraday observations for the monthly beta calculations. This frequency should be sufficiently high to reduce the measurement errors substantially while also avoiding the confounding influences from market microstructure frictions; see ABDL (1999b) for further discussion along these lines. The main advantage of using the intraday returns is that we should obtain less noisy measures of the relevant variance and covariances

[15] An alternative, and much more complicated approach, would be to utilize all of the observations by explicitly modeling the high-frequency frictions.

– conditional on the chosen 15-minute frequency controlling appropriately for market microstructure features – since we are pushing closer to continuous sampling, while the drawback is the much shorter time span available for analysis. Since the analysis above points towards the potential presence of short-run dependencies in betas, we focus on the monthly interval to explore this possibility in some depth.

Figures 1H and 2H (H indicates "high-frequency," in reference to the underlying 15-minute sampling) reveal a strong shift upward in both market volatility and the stocks covariances with the market in the latter half of the sample, where the values are both higher and much more volatile. In contrast, there is no visual evidence of any shift in the realized beta series in Figure 3H. The latter figures occasionally display some distinct outliers but these are not concentrated towards the second half of the sample, and the enhanced volatility in the market variance and the individual covariances in the second part of the 1990s is certainly not visible, implying that the dominant dynamic dependencies in the variances and covariances in fact tend to annihilate each other. As before, there is no sign of integer integration in any of the realized variance, covariance or beta measures and the median Ljung-Box statistic for the beta series of 58 in Table 2H is substantially below that of the covariances, 81 and the market variance, 121, confirming the visual impression of relatively less serial dependence in the realized monthly betas.

Further confirmation of the relative strength of the serial dependence across the series is evident in the Figures 4H-8H. The autocorrelations generally appear significant, or nearly so, for the realized market variance and the individual covariances out to an order of fifteen months. This is consistent with the significance of the autocorrelations out to lags four to six in the quarterly measures in Figures 4-6. In contrast, the individual realized betas rarely are significant beyond a lag length of five or six months, consistent with only the first few autocorrelation coefficients being significant in Figure 8. From Figure 7H it is evident that there is a large degree of heterogeneity in the results across the individual stocks, but the dominant pattern is clearly one with significantly positive coefficients at lower lags but also one with a sharply declining correlogram. Because none of the individual realized beta series display significant negative correlation at lower lags, the results point strongly towards the presence of a relatively short-lived but highly significant degree of positive serial dependence in the betas, as also evidenced by the median beta autocorrelation pattern in Figure 8H.

Continuing along the lines above, Figures 9H-13H present sample autocorrelations of the monthly realized measures, each prefiltered by $(1-L)^{.42}$. Since the degree of fractional integration is invariant to the sampling frequency and temporal aggregation, it is natural to compare these figures to those from the quarterly realized volatility measures. The message is in important respects similar to before: the market variance and the covariance measures seem consistent with a pure fractionally integrated process with $d = .42$. In contrast, however, the realized beta series behave differently from before, as there is no direct indication of overdifferencing this time. Of course, at the monthly frequency the serial dependence, as manifest in the correlogram, should appear stronger than for the quarterly series and thus generally blur our ability to distinguish the short run dependent case from the fractionally integrated one. The discrepancy in behavior for the monthly and quarterly series for the realized betas, and not for the realized market variance and the covariances, is nevertheless an indication that the invariance property against sampling frequency and time aggregation is violated for the beta series, as should occur if the series is truly $I(0)$. This lends additional credibility to our earlier assertion that, all told, $I(0)$ dynamics for realized betas seemed likely. Consequently, we again also model the realized betas using a standard autoregression. Table 3H presents results for low order $AR(p)$ model for each individual realized monthly beta series. The modulus of the dominant inverted root of the AR polynomial is now generally higher compared to Table 3, as is to be expected given the shorter estimation horizon for the betas. Figures 14H and 15H document that prefiltering using the estimated $AR(p)$ polynomial produces residuals with white noise characteristics. Overall, there are no signs that a null hypothesis of short memory dependence for the realized betas can be rejected.

In sum, the results across the two sampling frequencies and time horizons are quite consistent. There is strong evidence of a much lower degree of dependence in the realized betas compared to the realized market variance and the realized covariances between the individual stock returns and the market return. Although the realized betas can be approximated by a pure fractionally integrated process with a d of around .20 in the quarterly data – and perhaps an even higher degree of fractional integration in the shorter monthly sample – this may well be an artifact of the short sample size. There is seemingly some heterogeneity across the stock betas and the assumption of a standard short memory autoregressive process with generally significantly positive serial correlation for each of the individual realized betas appears more robust and fully adequate.

5. Assessing Precision: Interval Estimates of Betas

Thus far we have largely abstracted from the presence of estimation error in the realized betas. It is possible to assess the (time-varying) estimation error directly using formal continuous-record asymptotics, and we now do so.

Continuous-Record Asymptotic Standard Errors

We first use the multivariate asymptotic theory recently developed by Barndorff-Nielsen and Shephard (2003) to assess the precision of our realized betas which are, of course, estimates of the underlying integrated betas. This helps us think about separating "news from noise" when examining temporal movements in realized volatility.

Realized beta is simply

$$\hat{\beta}_{it} = \frac{\sum_{j=1}^{N_t} r_{ijt} r_{mjt}}{\sum_{j=1}^{N_t} r_{mjt}^2}, \tag{10}$$

where $\hat{\beta}_{it}$ is the realized beta for stock i in quarter t, r_{ijt} is the return of stock i on day j of quarter t, r_{mjt} is the return of the DJIA on day j of quarter t, and N_t is the number of units (e.g., days) into which quarter t is partitioned.[16] Under appropriate regularity conditions that allow for non-stationarity in the series, Barndorff-Nielsen and Shephard (2003) derive the limiting distribution of realized beta. They show that, as $N \Rightarrow \propto$,

$$\frac{\hat{\beta}_{it} - \beta_{it}}{\sqrt{\left(\sum_{j=1}^{N_t} r_{mjt}^2\right)^{-2} \hat{g}_{it}}} \Rightarrow N(0, 1) , \tag{11}$$

where

$$\hat{g}_{it} = \sum_{j=1}^{N_t} a_{ij}^2 - \sum_{j=1}^{N_t - 1} a_{ij} a_{i\,j+1} \tag{12}$$

and

$$a_{ij} = r_{ijt} r_{mjt} - \hat{\beta}_{it} r_{mjt}^2 . \tag{13}$$

Thus a feasible and asymptotically valid confidence interval for the underlying integrated beta is

$$\beta_{it} \in \hat{\beta}_{it} \pm z_{\alpha/2} \sqrt{\left(\sum_{j=1}^{N_t} r_{mjt}^2\right)^{-2} \hat{g}_{it}} , \tag{14}$$

where $z_{\alpha/2}$ is the appropriate critical value of the standard normal distribution.

In Figure 20 we plot the pointwise ninety-five percent confidence intervals for quarterly beta. They are quite wide, indicating that daily sampling is not adequate to drive out all measurement error. They are, given the width of the bands, moreover, consistent with the conjecture that there is only limited (short range) dependence in the realized beta series. In addition, notice the extremely wide bands associated with the outliers. Hence, potentially influential outliers are generally quite imprecisely estimated and the impact of such incidents on the analysis should be monitored closely.

One advantage of moving to high-frequency data should be that the increased number of observations

---

[16] N has a time subscript because the number of trading days varies slightly across quarters.

facilitates an improved measurement of the quarterly (or monthly) betas. To see if the reduction in measurement error manifests itself in better inference, we reproduce in Figure 21 the pointwise ninety-five percent confidence bands for the quarterly betas over the shorter sample that is covered by our intraday data. These may be compared directly with the corresponding quarterly realized beta standard error bands based on fifteen minute sampling in Figure 22. The improvement is readily visible in the sharp narrowing of the bands. It is evident from Figure 22 that there is quite pronounced positive dependence in the realized quarterly beta measures. In other words, it seems entirely warranted – and perhaps even critical – to complement beta estimates obtained from lower frequency data with ones derived from high-frequency intraday data in order to adequately characterize the dynamic evolution of individual security betas.

Once we pass to the monthly beta estimation frequency, even the intraday sampling provides a somewhat imprecise inference, as can be inferred from Figure 20H. Nonetheless, the standard error bands remain tighter than for the quarterly beta estimates based on daily data, and it is feasible to establish significant time-variation in the monthly betas from Figure 20H.

<u>HAC Asymptotic Standard Errors</u>

As noted previously, the quarterly realized betas are just regression coefficients computed quarter-by-quarter from CAPM regressions using intra-quarter daily data. One could obtain consistent estimates of the standard errors of those quarterly regression-based betas using HAC approaches, such as Newey-West, under the (very stringent) auxiliary assumption that the period-by-period betas are constant. For comparison to the Barndorff-Nielsen-Shephard results above, we compute the HAC standard error bands here.

In Figure 23 we provide the Newey-West ninety-five percent confidence intervals for the quarterly realized betas. Notice that the figure is strikingly similar to Figure 20. Hence, from a practical perspective there is not much difference in the assessment of the estimation uncertainly inherent in the quarterly beta measures obtained from the two alternative procedures. This conclusion is largely confirmed from a comparison of the Newey-West standard error bands for monthly betas based on intraday in Figure 23H with Figure 20H discussed above, where the two series again appear qualitatively similar. However, there is now evidence of significant difference in the width of the band associated with outliers in the distribution, so the differences are worth exploring further in future work.

6. Summary, Concluding Remarks, and Directions for Future Research

We have assessed the dynamics and predictability in realized betas, relative to the dynamics in the underlying market variance and covariances with the market. Key virtues of the approach include the fact that it does not require an assumed volatility model, and that it does not require an assumed model of time variation in beta. We find that, although the realized variances and covariances fluctuate widely and are highly persistent and predictable (as is well-known), the realized betas, which are simple nonlinear functions of those realized variances and covariances, display much less persistence and predictability.

The empirical literature on systematic risk measures, as captured by beta, is much too large to be discussed in a sensible fashion here. Before closing, however, we do want to relate our approach and results to a key literature and two key earlier papers that have important implications for the potential time variation of betas using time series techniques.

First, our results are closely linked to the literature on the latent (single) factor volatility model, as studied by a number of authors, including Diebold and Nerlove (1989), Harvey, Ruiz and Shephard (1994), King, Sentana and Wadhwani (1994), Fiorentini, Sentana and Shephard (1998), and Jacquier and Marcus (2000),

$$r_{it} = \beta_i f_t + v_{it}, \quad f_t \,|\, I_t \sim (0, h_t) \tag{15a}$$

$$v_{it} \overset{iid}{\sim} (0, \omega_i^2), \quad cov(v_{it} v_{jt'}) = 0, \ \forall\, i \neq j,\ t \neq t', \tag{15b}$$

where $i, j = 1, ..., N$, and $t = 1, ..., T$. The i[th] and j[th] time-t conditional variances, and the ij-th conditional covariance, for arbitrary i and j, are

$$h_{it} = \beta_i^2\, h_t + \omega_i^2, \quad h_{jt} = \beta_j^2\, h_t + \omega_j^2, \quad cov_{ijt} = \beta_i \beta_j\, h_t. \tag{16}$$

Assume, as is realistic in financial contexts, that all betas are nonnegative, and consider what happens as h_t moves (as h_t increases, say): All conditional variances increase, and all pairwise conditional covariances increase. Hence the market variance increases, and the covariances of individual equities with the market increase. Two observations are immediate: (1) both the market variance and the covariances of individual equities with the market are time-varying, and (2) because the market variance moves together with the covariances of individual equities with the market, the betas may not vary as much (indeed the betas are in fact constant in the model at hand, by construction!). The upshot is that wide variation in market variance and individual equity covariances with the market, yet no variation in betas, is precisely what one expects to see in a latent (single) factor volatility model. It is also, of course, quite similar to what we found in the data: wide variation and persistence in market variance and individual equity covariances with the market, yet less variation and persistence in betas. Notice also the remarkable similarity in the correlograms for the individual realized covariances in Figure 5. This is another indication of a strong coherence in the dynamic evolution of the individual covariances, consistent with the presence of one dominant underlying factor.

Second, our results nicely complement and expand upon those of Ghysels (1998), who argues that the constant beta CAPM, as bad as it likely is, is nevertheless not as bad as some popular conditional CAPMs. We provide some insight into why allowing for time-varying betas may do more harm than good when estimated from daily data, even if the true underlying betas display quite significant short memory dynamics: it may not be possible to estimate reliably the persistence or predictability in individual realized betas, so good in-sample fits may be spurious artifacts of data mining. We also establish that there should be a real potential for the use of high-frequency intraday data to improve upon this dilemma.

Third, our results also complement and expand upon those of Braun, Nelson and Sunier (1995), who study the discrepancy in the time series behavior of betas relative to the underlying variances and covariances for twelve industry portfolios using bivariate EGARCH models. They also find strong conditional heteroskedasticity in the variances and covariances, while there is a lesser degree of time variation in the conditional betas. Moreover, they find the strong asymmetric relationship between return innovations and future return volatility to be entirely absent in the conditional betas. Hence, at the portfolio level they document similar qualitative behavior between the variances and covariances relative to the betas as do we. However, their analysis is linked directly to a specific parametric representation, it studies industry portfolios, and it never contemplates the hypothesis that the constituent components of beta – variances and covariances – may be of a long memory form. The latter point has been confirmed by numerous studies in recent years. Consequently, our investigation can be seen as a substantive extension of their findings, showing that the main conclusions hold up when beta measurements are performed in a fully nonparametric fashion and fractional integration is explicitly accommodated.

In closing, we sketch an interesting direction for future research, which suggests a different way of using the high-frequency intraday data in the realized beta context, and which exploits the insights from the continuous record asymptotics more directly. We have thus far used monthly realized betas based on underlying high-frequency returns mostly to supplement and illustrate the findings from the quarterly betas constructed from daily data. However, the presence of a valid continuous record asymptotic theory for the size of the measurement errors in the realized beta measures and the indication of only short memory in the betas make an alternative approach to the intraday data based analysis feasible. This is readily seen by casting the system in state space form,

$$\hat{\beta}_{i,t} = \beta_{i,t} + u_{i,t} \tag{17a}$$

$$\beta_{i,t} = a_0 + a_1 \beta_{i,t-1} + v_{i,t} \tag{17b}$$

$$u_{i,t} \sim N\left(0, \left(\sum_{j=1}^{N_t} r_{mjt}^2\right)^{-2} \hat{g}_{it}\right), \qquad v_{i,t} \sim N\left(0, \sigma_{v,i,t}^2\right). \tag{17c}$$

The first relation indicates the measurement equation linking the estimated realized beta to the true underlying beta by explicitly introducing a normally distributed estimation error with the asymptotically valid variance obtained from the BNS procedure outlined earlier. The second relation specifies the transition equation to a standard first-order autoregression with potentially time-varying error. This allows for strictly stationary dependence in the evolution of the true betas. The simplest approach would be to let $v_{i,t}$ have a constant variance, but it is also straightforward to let this variance change in step with the underlying variability in the realized beta estimates so that the beta innovations are allowed to become more volatile as the constituent parts, the market variance and the covariance of the stock return with the market, grow. Finally, we shall assume that the error terms in the measurement and transition equations are uncorrelated. This approach directly utilizes the advantages of our intraday based beta measurements by incorporating estimates of the measurement errors to alleviate the error-in-variables problem and it allows for explicit recognition of the heteroskedasticity in the realized beta series.

References

Andersen, T.G. and Bollerslev, T. (1998), "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, 39, 885-905.

Andersen, T.G., Bollerslev, T. and Diebold, F.X. (2004), "Parametric and Nonparametric Volatility Measurement," in L.P. Hansen and Y. Ait-Sahalia (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland, forthcoming.

Andersen, T., Bollerslev, T., Diebold, F.X. and Ebens, H. (2001), "The Distribution of Realized Stock Return Volatility," *Journal of Financial Economics*, 61, 43-76.

Andersen, T. Bollerslev, T., Diebold, F.X. and Labys, P. (2001), "The Distribution of Realized Exchange Rate Volatility," *Journal of the American Statistical Association*, 96, 42-55.

Andersen, T., Bollerslev, T., Diebold, F.X. and Labys, P. (2003), "Modeling and Forecasting Realized Volatility," *Econometrica*, 71, 579-626.

Andersen, T., Bollerslev, T. and Meddahi, N. (2004), "Analytic Evaluation of Volatility Forecasts," Manuscript, Kellogg School, Northwestern University, *International Economic Review*, forthcoming.

Andrews, D.W.K. and Guggenberger, P. (2003), "A Bias-Reduced Log-Periodogram Regression Estimator of the Long-memory Parameter," *Econometrica*, 71, 675-712.

Ang, A. and Chen, J.C. (2003), "CAPM over the Long Run: 1926-2001," Manuscript, Columbia University and University of Southern California.

Barndorff-Nielsen, O.E. and Shephard, N. (2003), "Econometric Analysis of Realized Covariation: High Frequency Covariance, Regression and Correlation in Financial Economics," Manuscript, revised July 2003.

Baillie, R.T. (1996), "Long Memory Processes and Fractional Integration in Econometrics," *Journal of Econometrics*, 73, 5-59.

Bollerslev, T., Engle, R.F. and Wooldridge, J. (1988), "A Capital Asset Pricing Model with Time-Varying Covariances," *Journal of Political Economy*, 96, 113-131.

Braun, P.A., Nelson D.B. and Sunier A.M. (1995), "Good News, Bad News, Volatility, and Betas," *Journal of Finance*, 50, 1575-1603.

Breidt, F.J., Crato, N. and de Lima, P. (1998), "On the Detection and Estimation of Long Memory in Stochastic Volatility," *Journal of Econometrics*, 73, 325-334.

Campbell, J.Y., Lo, A.W., and MacKinlay, A. Craig (1997), *The Econometrics of Financial Markets*. Princeton: Princeton University Press.

Campbell, J.Y. and Vuolteenaho, T. (2002), "Bad Beta, Good Beta," Manuscript, Harvard University.

Cheung, Y.-W. and Lai K.S. (1993), "A Fractional Cointegration Analysis of Purchasing Power Parity," *Journal of Business and Economic Statistics*, 11, 103-112.

Cohen, R., C. Polk and Vuolteenaho, T. (2002), "Does Risk or Mispricing Explain the Cross-Section of Stock Prices," Manuscript, Kellogg School, Northwestern University.

Comte, F. and Renault, E. (1998), "Long Memory in Continuous Time Stochastic Volatility Models," *Mathematical Finance*, 8, 291-323.

Diebold, F.X. and Kilian, L. (2001), "Measuring Predictability: Theory and Macroeconomic Implications," *Journal of Applied Econometrics*, 16, 657-669.

Diebold, F.X. and Nerlove, M. (1989), "The Dynamics of Exchange Rate Volatility: A Multivariate Latent Factor ARCH Model," *Journal of Applied Econometrics*, 4, 1-22.

Dybvig, P.H. and Ross, S.A. (1985), "Differential Information and Performance Measurement Using a Security Market Line," *Journal of Finance*, 40, 383-400.

Engle, R.F. and Granger, C.W.J. (1987), "Co-Integration and Error Correction: Representation, Estimation and Testing," *Econometrica*, 55, 251-276.

Engle, R.F. and Kozicki, S. (1993), "Testing for Common Features," *Journal of Business and Economic Statistics*, 11, 369-383.

Fama, E.F. (1976)*, Foundations of Finance*. New York: Basic Books.

Fama, E.F. and French, K.R. (1992), "The Cross Section of Expected Stock Returns," *Journal of Finance*, 47, 427-465.

Fama, E.F. and French, K.R. (1993), "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3-56.

Ferson, W.E. and Harvey, C.R. (1991), "The Variation of Economic Risk Premiums," *Journal of Political Economy*, 99, 385-415.

Ferson, W.E., Kandel, S. and Stambaugh, R.F. (1987), "Tests of Asset Pricing with Time-Varying Expected Risk Premiums and Market Betas," *Journal of Finance*, 42, 201-220.

Fiorentini, G., Sentana, E. and Shephard, N. (1998), "Exact Likelihood-Based Estimation of Conditionally Heteroskedastic Factor Models," Working Paper, University of Alcante, CEMFI and Oxford University.

Geweke, J. and Porter-Hudak, S. (1983), "The Estimation and Application of Long Memory Time Series Models," *Journal of Time Series Analysis*, 4, 221-238.

Ghysels, E. (1998), "On Stable Factor Structures in the Pricing of Risk: Do Time-Varying Betas Help or Hurt?," *Journal of Finance*, 53, 549-573.

Granger, C.W.J. (1995), "Modeling Nonlinear Relationships Between Extended-Memory Variables," *Econometrica*, 63, 265-279.

Granger, C.W.J. and Newbold, P. (1986), *Forecasting Economic Time Series* (Second Edition). Orlando: Academic Press.

Hansen, L.P. and Richard, S.F. (1987), “The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models,” *Econometrica*, 55, 587-613.

Harvey, A.C. (1998), “Long Memory in Stochastic Volatility,” in J. Knight and J. Satchell (eds.), *Forecasting Volatility in Financial Markets*. London, UK: Butterworth-Heinemann.

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

Horvath, M.T.K. and Watson, M.W. (1995), “Testing for Cointegration When Some of the Cointegrating Vectors are Prespecified,” *Econometric Theory*, 11, 952-984.

Huang, C. And Litzenberger, R.H. (1988), *Foundations for Financial Economics*. Amsterdam: North-Holland.

Hull, J. and White, A. (1987), “The Pricing of Options on Assets with Stochastic Volatilities,” *Journal of Finance*, 42, 381-400.

Jacquier, E. and Marcus, A.J. (2000), “Market Volatility and Asset Correlation Structure,” Working Paper, Boston College.

Jagannathan, R. and Wang, Z. (1996), “The Conditional CAPM and the Cross Section of Expected Returns,” *Journal of Finance*, 51, 3-53.

Keim, D. and Hawawini, G. (1999), “The Cross Section of Common Stock Returns: A Review of the Evidence and Some New Findings,” Manuscript, Wharton School.

King, M., Sentana, E. and Wadhwani, S. (1994), “Volatility and Links Between National Stock Markets,” *Econometrica*, 62, 901-933.

Lintner, J. (1965a), “Security Prices, Risk and Maximal Gains From Diversification,” *Journal of Finance*, 20, 587-615.

Lintner, J. (1965b), “The Valuation of Risky assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets,” *Review of Economics and Statistics*, 47, 13-37.

Meddahi, N. (2002), “A Theoretical Comparison Between Integrated and Realized Volatility,” *Journal of Applied Econometrics*, 17, 479-508.

Robinson, P.M. (2001), "The Memory of Stochastic Volatility Models," *Journal of Econometrics*, 101, 195-218.

Robinson, P.M. and Marinucci, D. (2001), “Semiparametric Fractional Cointegration Analysis,” *Journal of Econometrics*, 105, 225-247.

Roll, R. (1977), “A Critique of the Asset Pricing Theory's Tests,” *Journal of Financial Economics*, 4, 129-176.

Rosenberg, B. (1973), “Random Coefficient Models: The Analysis of a Cross Section of Time Series by Stochastically Convergent Parameter Regression,” *Annals of Economic and Social Measurement*,” 2, 399-428.

Schaefer, S., Brealey, R., Hodges, S., and Thomas, H. (1975), "Alternative Models of Systematic Risk," in E. Elton and M. Gruber (eds.), *International Capital Markets: An Inter- and Intra-Country Analysis*, 150-161. Amsterdam: North-Holland.

Schwert, G.W. (1989), "Why Does Stock Market Volatility Change Over Time?," *Journal of Finance*, 44, 1115-1153.

Sharpe, W. F. (1963), "A Simplified Model for Portfolio Analysis," *Management Science*, 9, 227-293.

Wang, K.Q. (2003), "Asset Pricing with Conditioning Information: A New Test," *Journal of Finance*, 58, 161-196.

Table 1
The Dow Jones Thirty

Notes: The table summarizes company names and tickers, and the range of the data examined. We use the Dow Jones Thirty as of March 1997. Tickers with asterisks denote stocks with incomplete data, which we exclude from the analysis.

Figure 1a
Time Series Plot of Quarterly Realized Market Variance, Fall 1987 Included

Figure 1b
Time Series Plot of Quarterly Realized Market Variance, Fall 1987 Excluded

Notes: The two subfigures show the time series of quarterly realized market variance, with the 1987:4 outlier included (Figure 1a) and excluded (Figure 1b). The sample covers the period from 1962:3 through 1999:3, for a total of 149 observations. We calculate the realized quarterly market variances from daily returns.

Figure 2
Time Series Plots of Quarterly Realized Covariances

Notes: The Figure shows the time series of quarterly realized covariances, with the 1987:4 outlier excluded. The sample covers the period from 1962:3 through 1999:3, for a total of 148 observations. We calculate the realized quarterly covariances from daily returns.

Figure 3
Time Series Plots of Quarterly Realized Betas

Notes: The Figure shows the time series of quarterly realized betas, with the 1987:4 outlier excluded. The sample covers the period from 1962:3 through 1999:3, for a total of 148 observations. We calculate the realized quarterly betas from daily returns.

Table 2
The Dynamics of Quarterly Realized Market Variance, Covariances and Betas

Notes: The table summarizes aspects of the time-series dependence structure of quarterly realized market variance, covariances and realized betas. Q_{12} denotes the Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation, and ADF^{i} denotes the augmented Dickey-Fuller unit root test with I augmentation lags. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized variance, covariances and betas from daily returns.

Figure 4
Sample Autocorrelations of Quarterly Realized Market Variance

Notes: The figure shows the first 36 sample autocorrelations of quarterly realized market variance. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized variance from daily returns.

Figure 5
Sample Autocorrelations of Quarterly Realized Covariances

Notes: The figure shows the first 36 sample autocorrelations of quarterly realized covariances. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized covariances from daily returns.

Figure 6
Median Sample Autocorrelations of Quarterly Realized Covariances

Notes: The figure shows the medians across individual stocks of first 36 sample autocorrelations of quarterly realized covariances. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the median of Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized covariances from daily returns.

Figure 7
Sample Autocorrelations of Quarterly Realized Betas

Notes: The figure shows the first 36 sample autocorrelations of quarterly realized betas. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized betas from daily returns.

Figure 8
Median Sample Autocorrelations of Quarterly Realized Betas

Notes: The figure shows the medians across individual stocks of first 36 sample autocorrelations of quarterly realized betas. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the median of Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized betas from daily returns.

Figure 9

Sample Autocorrelations of Quarterly Realized Market Variance Prefiltered by $(1-L)^{.42}$

Notes: The figure shows the first 36 sample autocorrelations of quarterly realized market variance prefiltered by $(1-L)^{.42}$. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the median of Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized variance from daily returns.

Figure 10
Sample Autocorrelations of Quarterly Realized Covariances Prefiltered by $(1-L)^{.42}$

Notes: The figure shows the first 36 sample autocorrelations of quarterly realized covariances prefiltered by $(1-L)^{.42}$. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized covariances from daily returns.

Figure 11
Median Sample Autocorrelations of Quarterly Realized Covariances Prefiltered by $(1-L)^{.42}$

Notes: The figure shows the medians across individual stocks of first 36 sample autocorrelations of quarterly realized covariances prefiltered by $(1-L)^{.42}$. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the median of Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized covariances from daily returns.

Figure 12
Sample Autocorrelations of Quarterly Realized Betas Prefiltered by $(1-L)^{.42}$

Notes: The figure shows the first 36 sample autocorrelations of quarterly realized betas prefiltered by $(1-L)^{.42}$. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized betas from daily returns.

Figure 13
Median Sample Autocorrelations of Quarterly Realized Betas Prefiltered by $(1-L)^{.42}$

Notes: The figure shows the medians across individual stocks of first 36 sample autocorrelations of quarterly realized betas prefiltered by $(1-L)^{.42}$. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the median of Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized betas from daily returns.

Table 3
Inverted Roots of $AR(p)$ Models for Quarterly Realized Betas
Notes: The table shows the inverted roots and modulus of the dominant root of autoregressive lag operator polynomials $(1-\hat{\phi}_1 L-\hat{\phi}_2 L^2-...-\hat{\phi}_p L^p)$, where $\hat{\phi}_1,\ \hat{\phi}_2,\ ...,\ \hat{\phi}_p$ are the least squares estimates of the parameters of $AR(p)$ models fit to the realized betas, with p selected by the AIC. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized variance, covariances and betas from daily returns.

Figure 14
Sample Autocorrelations of Quarterly Realized Betas Prefiltered by $(1-\hat{\phi}_1 L-\hat{\phi}_2 L^2-...-\hat{\phi}_p L^p)$

Notes: The figure shows the first 36 sample autocorrelations of quarterly realized betas prefiltered by $(1-\hat{\phi}_1 L-\hat{\phi}_2 L^2-...-\hat{\phi}_p L^p)$, where $\hat{\phi}_1,\ \hat{\phi}_2,\ ...,\ \hat{\phi}_p$ are the least squares estimates of the parameters of $AR(p)$ models fit to the realized betas, with p selected by the AIC. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized variance, covariances and betas from daily returns.

Figure 15
Median Sample Autocorrelations of Quarterly Realized Betas Prefiltered by $(1-\hat{\phi}_1 L-\hat{\phi}_2 L^2-...-\hat{\phi}_p L^p)$

Notes: The figure shows the medians across individual stocks of the first 36 sample autocorrelations of quarterly realized betas prefiltered by $(1-\hat{\phi}_1 L-\hat{\phi}_2 L^2-...-\hat{\phi}_p L^p)$, where $\hat{\phi}_1,\ \hat{\phi}_2,\ ...,\ \hat{\phi}_p$ are the least squares estimates of the parameters of $AR(p)$ models fit to the realized betas, with p selected by the AIC. The dashed lines denote Bartlett's approximate 95 percent confidence band in the white noise case. Q_{12} denotes the Ljung-Box portmanteau statistic for up to twelfth-order autocorrelation. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized variance, covariances and betas from daily returns.

Figure 16
Predictability of Market Volatility and Individual Equity Covariances with the Market
***I*(*d*) Approximation of Dynamics**

Notes: We define predictability as $P_j\ =\ 1-var(e_{t+j,t})/var(e_{t+40,t})$, where $var(e_{t+j,t})\ =\ \sigma^2 \sum_{i=0}^{j-1} b_i^2$, σ_t^2 is the variance of the innovation ε_t, and the b_i's are moving average coefficients; i.e., the Wold representation is $y_t\ =\ (1+b_1 L+b_2 L^2+b_2 L^3+...)\varepsilon_t$. Here we plot P_j for $j=1,...,7$. We approximate the dynamics using a pure long-memory model, $(1-L)^{.42} y_t\ =\ \varepsilon_t$, in which case $b_0\ =\ 1$ and $b_i\ =\ (-1)\, b_{i-1}\,(d-i+2)/(i-1)$. Moreover, because we take d=0.42 for market volatility and for all covariances with the market, all of their predictabilities are the same at all horizons.

Figure 17
Predictability of Betas
***I*(*d*) Approximation of Dynamics**

Notes: We define predictability as $P_j = 1 - var(e_{t+j,t})/var(e_{t+40,t})$, where $var(e_{t+j,t}) = \sigma^2 \sum_{i=0}^{j-1} b_i^2$, σ_t^2 is the variance of the innovation ε_t, and the b_i's are moving average coefficients; i.e., the Wold representation is $y_t = (1 + b_1 L + b_2 L^2 + b_2 L^3 + ...)\varepsilon_t$. Here we plot P_j for $j = 1,...,7$. We approximate the dynamics using a pure long-memory model, $(1 - L)^{.20} y_t = \varepsilon_t$, in which case $b_0 = 1$ and $b_i = (-1)\, b_{i-1}\,(d - i + 2)/(i - 1)$. Moreover, because we take d=0.20 for all betas, all of their predictabilities are the same at all horizons.

Figure 18
Predictability of Betas
***AR*(*p*) Sieve Approximation of Dynamics**

Notes: We define predictability as $P_j = 1 - var(e_{t+j,t})/var(y_t)$, where $var(e_{t+j,t}) = \sigma^2 \sum_{i=0}^{j-1} b_i^2$, $var(y_t) = \sigma^2 \sum_{i=0}^{\infty} b_i^2$, σ_t^2 is the variance of the innovation ε_t, and the b_i's are moving average coefficients; i.e., the Wold representation is $y_t = (1 + b_1 L + b_2 L^2 + b_2 L^3 + ...)\varepsilon_t$. Here we plot P_j for $j = 1,...,7$. We approximate the dynamics using an *AR*(*p*) model, with the autoregressive lag order *p* determined by the AIC. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized covariances from daily returns.

Figure 19
Median Predictability of Betas
***AR*(*p*) Sieve Approximation of Dynamics**

Notes: We define predictability as $P_j = 1 - var(e_{t+j,t})/var(y_t)$, where $var(e_{t+j,t}) = \sigma^2 \sum_{i=0}^{j-1} b_i^2$, $var(y_t) = \sigma^2 \sum_{i=0}^{\infty} b_i^2$, σ_t^2 is the variance of the innovation ε_t, and the b_i's are moving average coefficients; i.e., the Wold representation is $y_t = (1 + b_1 L + b_2 L^2 + b_2 L^3 + ...)\varepsilon_t$. Here we plot P_j for $j = 1,...,7$. We approximate the dynamics using an *AR*(*p*) model, with the autoregressive lag order *p* determined by the AIC. The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded, for a total of 148 observations. We calculate the quarterly realized covariances from daily returns.

Figure 20
Ninety-Five Percent Confidence Intervals for Quarterly Beta, Long Sample, Daily Sampling

Notes: The figure shows the time series of ninety-five percent confidence intervals for the underlying quarterly integrated beta, calculated using the results of Barndorff-Nielsen and Shephard (2001). The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded. We calculate the realized quarterly betas from daily returns.

Figure 21
Ninety-Five Percent Confidence Intervals for Quarterly Beta, Short Sample, Daily Sampling

Notes: The figure shows the time series of ninety-five percent confidence intervals for the underlying quarterly integrated beta, calculated using the results of Barndorff-Nielsen and Shephard (2001). The sample covers the period from 1962:3 through 1999:3. We calculate the realized quarterly betas from daily returns.

Figure 22
Ninety-Five Percent Confidence Intervals for Quarterly Beta, Short Sample, Fifteen-Minute Sampling

Notes: The figure shows the time series of ninety-five percent confidence intervals for the underlying quarterly integrated beta, calculated using the results of Barndorff-Nielsen and Shephard (2001). The sample covers the period from 1993:1 through 1999:3. We calculate the realized quarterly betas from fifteen-minute returns.

Figure 23
Ninety-Five Percent Confidence Intervals for Quarterly Beta, Long Sample, Daily Sampling (Newey-West)

Notes: The figure shows the time series of Newey-West ninety-five percent confidence intervals for the underlying quarterly integrated beta, calculated using the results of Barndorff-Nielsen and Shephard (2001). The sample covers the period from 1962:3 through 1999:3, with the 1987:4 outlier excluded. We calculate the realized quarterly betas from daily returns.