

Andersen, Torben G.; Bollerslev, Tim

Article

Realized return volatility, asset pricing, and risk management

NBER Reporter Online

Provided in Cooperation with:

National Bureau of Economic Research (NBER), Cambridge, Mass.

Suggested Citation: Andersen, Torben G.; Bollerslev, Tim (2006) : Realized return volatility, asset pricing, and risk management, NBER Reporter Online, National Bureau of Economic Research (NBER), Cambridge, MA, Iss. Fall 2006, pp. 7-10

This Version is available at:

<https://hdl.handle.net/10419/61896>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Realized Return Volatility, Asset Pricing, and Risk Management

Torben G. Andersen and Tim Bollerslev*

It is now widely accepted that expected returns, volatility, and broader financial risk measures all vary over time. In particular, there is a pronounced clustering in return volatility; occasional extreme return outliers—especially on the negative for equities; and an increase in return correlations during market downturns. This makes it more complicated for academics, regulators, and practitioners seeking to understand, monitor, act, and react to financial market dynamics to assess market conditions in real time. Textbook prescriptions for portfolio choice, asset pricing, and risk management typically are based on a static setting with known and invariant return distributions. These approaches are ill-suited for practical decision making: market agents know neither the parameters nor the parametric family of the return distribution, and the shape of the distribution is likely to change over time. Depending on the horizon, the challenges differ, with the notable exception that accurate assessment of the current volatility level remains pivotal. At daily or shorter intervals, it is critical also to understand the likely reaction of markets to impending news releases and to control for the intraday pattern in the market activity and return dynamics. For weekly and monthly frequencies, the persistence of volatility and the extent of asymmetry

between return and volatility innovations both figure importantly in determining return distributions. For even longer quarterly and annual horizons, the main issues again relate to the temporal persistence of volatility, but good estimates of the non-negligible longer-run expected returns now also become critical.

The increased availability of tick-by-tick financial trade records and real-time news reports, coupled with our enhanced capacity to store and process vast amounts of data, have led to important new insights in regards to the issues discussed above. Specifically, over the last few years a very active research agenda into the direct (model-free) measurement of the realized return variation and covariation of financial assets at daily or even higher intraday frequencies has developed.

The intuition behind the realized volatility measure has been recognized for a while, albeit within a simplified setting. In a frictionless market with an unlimited set of price observations available over any interval, it is, quite generally, feasible to perfectly estimate instantaneous volatility if the process is not subject to jumps. However, given the discreteness of the price grid and other market microstructure effects, as well as the limited number of price observations available over short time intervals, even for liquid securities, instantaneous volatility cannot be measured with reasonable precision without (excessively) strong identifying assumptions. In the face of these practical limitations, we have focused a large part of our recent work on developing robust, yet accurate, volatility measures over non-trivial daily, or longer, time intervals that exploit the information available from intraday data.

In so doing, it is important to recognize the main qualitative features that

affect the intraday return process but are absent at daily and lower frequency levels. Most importantly, the intraday volatility pattern and the presence of outliers (jumps) render standard ARCH-type volatility models inadequate unless they are explicitly extended to accommodate such features. We show that the original studies applying standard modeling and inference techniques to the newly available intraday data were seriously misspecified; they produced badly downward biased estimates of the degree of volatility persistence.¹ Meanwhile, by controlling for specific intraday features, we got much closer to the type of volatility dynamics obtained from daily data, although our model specification is still not entirely adequate. In short, direct estimation of the high-frequency volatility process is difficult and very sensitive to market microstructure effects and news.²

We instead advocate daily (or longer-horizon) volatility and covariability measures obtained by aggregating intraday squared returns and absolute return cross-products. Focusing on a non-negligible time interval enables us to exploit many return observations, ensuring that the estimated measure is reasonably precise. Moreover, by restricting the measurement to (a multiple of) a trading day and relying on equally-spaced returns sampled, say, every five or ten minutes, we can largely eliminate the intraday volatility pattern and other market microstructure effects. Formally, as the number of returns observed over the period grows toward infinity, the realized volatility provides a consistent measure of the ex-post return variation. Intuitively, the impact of the mean return is removed by the shrinking of the intraday time intervals, as the expected price movements become negligible relative to the return innovations.

**Andersen is a Research Associate in the NBER's Program on Asset Pricing. He is also the Nathan S. and Mary P. Sharp Distinguished Professor of Finance at the Kellogg School of Management, Northwestern University. Bollerslev is a Research Associate in the NBER's Program on Asset Pricing. He is also the Juanita and Clifton Kreps Professor of Economics, Duke University, and Professor of Finance at the Fuqua School of Business. Their profiles appear later in this issue.*

Importantly, these measures are conceptually distinct from model-based volatility estimates and/or forecasts from traditional models such as GARCH. They represent actual realized return variability assessed from ex-post data rather than ex-ante (conditional) return variances implied by a parametric model. Because volatility is genuinely stochastic, the realized variability inevitably differs from the ex-ante expectations, even if these are based on the true model. In other words, realized volatility represents the (true) expected volatility plus an unpredictable volatility innovation. In contrast, even if the daily squared return is almost unbiased for the underlying volatility, it is an extremely noisy estimator.

In this regard, when checking the adequacy of specific volatility models, we document the extraordinary improvement in the signal-to-noise ratio obtained by using the realized volatility estimators relative to the common practice of using the ex-post squared returns.³ We find that, in certain realistic scenarios, the one-day-ahead volatility forecasts from the true model may explain up to half of the subsequent variation in the realized volatility; the same forecasts only “explain” about 5 percent of the variation in the future squared daily returns. We pursue the topic in detail—using simulation techniques and more elegant analytical means—in joint work with Steve Lange⁴ and Nour Meddahi⁵ respectively, emphasizing the impact of the forecast horizon and sampling frequency.

Given the direct construction of realized volatility from intraday returns, volatility in effect may be treated as observable, albeit with a limited measurement error. This sets the stage for standard time-series analysis of (logarithmic) volatility, a theme pursued jointly with Frank Diebold and Paul Labys in analyzing the volatility and covariability of foreign exchange returns,⁶ and with Diebold and Heiko Ebens for individual stock returns.⁷ This integrated approach to volatility measurement and modeling is pursued further with Diebold⁸ and Nour Meddahi⁹, respectively. In that work, we directly demonstrate the effectiveness of the approach for volatility forecasting. Moreover, with an

accurate volatility proxy in-hand, we can study the properties of daily returns standardized by (realized) volatility. We find these to be much closer to Gaussian than is the case for standardized return residuals from stochastic volatility models, underscoring the potential gains from adapting the more precise volatility measures.¹⁰

A related contentious issue concerns the nature of the longer-run dependencies in return volatility. Recent work using daily returns has produced evidence of so-called long memory, implying a slow hyperbolic decay in the absolute and squared return auto-correlation patterns, rather than the faster geometric decay associated with traditional volatility models. This, of course, has important implications for longer-run conditional volatility and return distribution forecasts. Meanwhile, it has been suggested that this apparent long-memory is (spuriously) induced by infrequent structural changes in the volatility. Thus, it may be better captured by regime-shifting type models. The sharply enhanced inferential power obtained through the realized volatility measures allows for much stronger tests of the long-memory property over much shorter (calendar) samples than is possible with only daily or lower frequency data. Our original study along these lines strongly supported the long-memory hypothesis.¹¹ That finding has been confirmed by numerous later studies, even if this remains an active research topic.

The intraday return data also facilitate the study of market reactions to economic news. We find that a complete account of the foreign exchange return dynamics must include controls for the jumps that occur in response to scheduled U.S. macroeconomic news releases, such as the employment report and CPI inflation. Such news induce an immediate price revision along with an intensive and more refined price discovery process, associated with sharply enhanced volatility, lasting up to about two hours.¹² On the days of these releases, the induced jump and volatility contribute very significantly to the overall daily return variability. In work with Diebold and Clara Vega, we study more detailed issues, such as the impact of the expected announcement figure versus

the surprise component and the sequence of releases relating to economic developments over a given month.¹³

In addition, expanding our perspective to include equity and bond markets, we document important linkages between the state of the business cycle and the financial market reaction to real and inflationary economic news. For example, we find that interest rates and equity market returns are negatively correlated in the expansion phase but positively correlated during recessions.¹⁴ This approach has the potential to elicit direct evidence on the structural linkages across macro markets and thus enable us to study their time variation over both business cycles and distinct policy regimes.

Another avenue for exploring asset pricing issues using the intraday returns is to relate asset-specific realized volatility to the evolution of systematic macroeconomic factors in order to gauge the potential risk exposure of the security. Our joint work with Diebold and Ginger Wu provides one step in this direction. We find interesting systematic shifts over the business cycle in the size of the market betas of so-called value stocks relative to growth stocks, suggesting that the former are systematically perceived as more risky than the latter, which may help to explain the puzzling “value premium.” Nonetheless, a more complete study, explicitly accounting for additional risk factors over longer time spans, is needed to validate the asset pricing implications of the documented features.¹⁵

The many useful applications of realized volatility have motivated a recent, somewhat technical, literature that seeks to minimize the aforementioned measurement errors induced by the presence of market microstructure frictions. The alternative realized volatility measures developed in this literature may also be used for robust inference concerning a variety of features in the underlying price process. In work with Diebold, we provide an overview of some of the developments in this rapidly progressing literature.¹⁶ In further work with Diebold, we have focused on the application of the realized volatility measures along with some new related concepts termed power and bipower vari-

ation measures — obtained by summing properly scaled functions of the intraday absolute returns — to identify the timing and size of discontinuities, or jumps, in the prices for broad stock, bond, and foreign exchange markets. We find that the jumps are less persistent than the smooth, or diffusive, volatility component. We go on to show how this may be used in the construction of more accurate return variability forecasts by decomposing the realized volatility into its diffusive and jump components.¹⁷

Our recent work with Dobrislav Dobrev,¹⁸ and Per Frederiksen and Morten Nielsen¹⁹, provides a more systematic study of the applicability of the realized volatility tools in analysis of equity return distributions. On extracting the significant jumps and transforming the daily return process into a financial time-scale, with each “financial day” representing an equal amount of realized volatility, we find that the returns are indistinguishable from i.i.d. Gaussian. Importantly, these results directly confirm the theoretical underpinnings for the general continuous time jump-diffusive price representation commonly used in asset pricing and financial economics. More broadly, the findings confirm the practical reliability of the new realized volatility tools and the associated theory, and pave the way for further progress in characterizing and forecasting the full conditional return distributions. More research is needed, in particular in terms of the corresponding tools for the multivariate setting.

Numerous other useful applications of the realized volatility concept still await. For instance, specifying and directly estimating more realistic parametric, continuous-time asset pricing models may be made easier by matching the implications from the models with the directly observable realized volatility measures.²⁰ Also, finance theory often implies specific conditional volatility distributions, and/or conditional correlations, between the asset volatilities and the volatility of the systematic risk factors. One example is the volatility risk premium inherent in financial derivatives prices.²¹ Another example is the affine term structure models, which imply that the yield volatility of zero-cou-

pon bonds at any maturity is spanned by the level of contemporaneous yields across the risk-free term structure.²²

In light of the rising prominence of the realized volatility concept for a variety of applications, it occupies a key position in our recent surveys, written jointly with Peter Christoffersen and Diebold, on risk management²³ and volatility forecasting.²⁴ We are currently working on a variety of additional aspects and applications of realized volatility. These include a more detailed investigation of the frequency and dynamic dependencies in the jump dynamics and direct studies of the presence and time-variation in volatility risk premiums. We expect to report on our findings from these projects in the near future.

¹ T.G. Andersen and T. Bollerslev, “Intraday Periodicity and Volatility Persistence in Financial Markets,” *Journal of Empirical Finance*, Vol. 4 (1997), pp.115–58, and reprinted in *Foreign Exchange Markets*, Chapter 5 (2005), pp.89–132, edited by R.J. Sweeney, *International Library of Critical Writings in Financial Economics*, Series Editor, R. Roll, Elgar Publishing Ltd, Cheltenham Glos, United Kingdom.

² T.G. Andersen and T. Bollerslev, “Deutschemark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies,” *NBER Working Paper No. 5783*, October 1996, and *Journal of Finance*, Vol. 53 (1998), pp.219–65, also reprinted in *Foreign Exchange Markets*, Chapter 6 (2005), pp.133–79.

³ T.G. Andersen and T. Bollerslev, “Answering the Skeptics: Yes, ARCH Models Do Provide Good Volatility Forecasts,” *NBER Working Paper No. 6023*, April 1997, and *International Economic Review*, Vol. 39 (1998), pp.885–905, also reprinted in *Forecasting Financial Markets*, (2002), edited by T.C. Mills, *International Library of Critical Writings in Economics*, Series Editor, M. Blaug, Elgar Publishing Ltd, Cheltenham Glos, United Kingdom.

⁴ T.G. Andersen, T. Bollerslev, and S. Lange, “Forecasting Financial Market Volatility: Sampling Frequency vis-a-vis

Forecast Horizon,” *Journal of Empirical Finance*, Vol. 6 (1999), pp.457–77.

⁵ T.G. Andersen, T. Bollerslev, and N. Meddahi, “Analytical Evaluation of Volatility Forecasts,” *International Economic Review*, Vol. 45 (2004), pp.1079–110.

⁶ T.G. Andersen, T. Bollerslev, F.X. Diebold, and P. Labys, “The Distribution of Exchange Rate Volatility,” *NBER Working Paper No. 6961*, February 1999, and *Journal of American Statistical Association*, Vol. 96 (2001), pp.42–55, also reprinted in *Stochastic Volatility: Selected Readings*, Chapter 15 (2005), pp.451–79, N. Shephard, ed.; *Advanced Texts in Economics*, Series Editors, M. Arellano, G. Imbens, G.E. Mizon, A. Pagan, and M. Watson, Oxford University Press, Oxford, U.K.

⁷ T.G. Andersen, T. Bollerslev, F.X. Diebold, and H. Ebens, “The Distribution of Stock Return Volatility,” *NBER Working Paper No. 7933*, October 2000, and *Journal of Financial Economics*, Vol. 61 (2001), pp.43–76.

⁸ T.G. Andersen, T. Bollerslev, and F.X. Diebold, “Modeling and Forecasting Realized Volatility,” *NBER Working Paper No. 8160*, March 2001, and *Econometrica*, Vol. 71 (2003), pp.579–625.

⁹ T.G. Andersen, T. Bollerslev, and N. Meddahi, “Correcting the Errors: On Forecast Evaluation Using High-Frequency Data and Realized Volatilities,” *Econometrica*, Vol. 73 (2005), pp.279–96.

¹⁰ T.G. Andersen, T. Bollerslev, F.X. Diebold, and P. Labys, “Exchange Rate Returns Standardized by Realized Volatility Are (Nearly) Gaussian,” *NBER Working Paper No. 7488*, January 2000, and *Multinational Finance Journal*, Vol. 4 (2000), pp.159–79.

¹¹ T.G. Andersen and T. Bollerslev, “Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns,” *NBER Working Paper No. 5752*, September 1996, and *Journal of Finance*, Vol. 52 (1997), pp.975–1005. Related simulation-based evidence is also reported in T. Bollerslev and Jonathan H. Wright, “Semiparametric Estimation of

Long-Memory Volatility Dependencies: The Role of High-Frequency Data, Journal of Econometrics, Vol. 98 (2000), pp. 81–106.

¹² T.G. Andersen and T. Bollerslev, “Deutschemark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies,” NBER Working Paper No. 5783, October 1996, and Journal of Finance, Vol. 53 (1998), pp. 219–265, also reprinted in Foreign Exchange Markets, Chapter 6 (2005), pp. 133–79.

¹³ T.G. Andersen, T. Bollerslev, F.X. Diebold, and C. Vega, “Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange,” NBER Working Paper No. 8959, May 2002, and American Economic Review, Vol. 93 (2003), pp. 38–62.

¹⁴ T.G. Andersen, T. Bollerslev, F.X. Diebold, and C. Vega, “Real-Time Price Discovery in Stock, Bond and Foreign Exchange Markets,” NBER Working Paper No. 11312, May 2005.

¹⁵ T.G. Andersen, T. Bollerslev, F.X. Diebold, and J. Wu, “A Framework for Exploring the Macroeconomic Determinants of Systematic Risk,” NBER Working Paper No. 11134, February 2005, and American Economic Review, Vol. 95 (2005), pp. 398–404. For a more detailed analysis, see also T.G.

Andersen, T. Bollerslev, F.X. Diebold, and J. Wu, “Realized Beta: Persistence and Predictability,” in T. Fomby (editor), Advances in Econometrics: Econometric Analysis of Economic and Financial Time Series, Volume B (2006), pp. 1–40.

¹⁶ T.G. Andersen, T. Bollerslev, and F.X. Diebold, “Parametric and Nonparametric Volatility Measurement,” NBER Technical Working Paper No. 279, August 2002, forthcoming in Y. Ait-Sahalia and L.P. Hansen (editors), Handbook of Financial Econometrics, North Holland.

¹⁷ T.G. Andersen, T. Bollerslev, and F.X. Diebold, “Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility,” NBER Working Paper No. 11775, November 2005, and forthcoming in Review of Economics and Statistics.

¹⁸ T.G. Andersen, T. Bollerslev, and D. Dobrev, “No-Arbitrage Semi-Martingale Restrictions for Continuous-Time Volatility Models subject to Leverage Effects, Jumps and i.i.d. Noise: Theory and Testable Distributional Assumptions,” forthcoming as an NBER Working Paper and in the Journal of Econometrics, 2006.

¹⁹ T.G. Andersen, T. Bollerslev, Per H. Frederiksen, and Morten Ø. Nielsen, “Continuous-Time Models, Realized Volatilities, and Testable Distributional Implications for Daily Stock Returns,”

forthcoming as an NBER Working Paper.

²⁰ T. Bollerslev and H. Zhou, “Estimating Stochastic Volatility Diffusions Using Conditional Moments of Integrated Volatility,” Journal of Econometrics, Vol. 109 (2002), pp. 33–65.

²¹ T. Bollerslev, M. Gibson, and H. Zhou, “Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities,” forthcoming as an NBER Working Paper.

²² T.G. Andersen and L. Benzoni, “Can Bonds Hedge Volatility Risk in the U.S. Treasury Market? A Specification Test for Affine Term Structure Models,” forthcoming as an NBER Working Paper.

²³ T.G. Andersen, T. Bollerslev, P. Christoffersen, and F.X. Diebold, “Practical Volatility and Correlation Modeling for Financial Market Risk Management,” NBER Working Paper No. 11069, January 2005, and forthcoming in M. Carey and R. Stulz, eds., Risk for Financial Institutions, Chicago: University of Chicago Press.

²⁴ T.G. Andersen, T. Bollerslev, P. Christoffersen, and F.X. Diebold, “Volatility Forecasting,” NBER Working Paper No. 11188, March 2005, and in G. Elliott, C. W.J. Granger, and A. Timmermann, eds., Handbook of Economic Forecasting, Chapter 15 (2006), pp. 777–878, North Holland.