

Johnson, James A.; Medeiros, Marcelo C.; Paye, Bradley S.

Working Paper

Jumps in stock prices: New insights from old data

Texto para discussão, No. 682

Provided in Cooperation with:

Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro

Suggested Citation: Johnson, James A.; Medeiros, Marcelo C.; Paye, Bradley S. (2021) : Jumps in stock prices: New insights from old data, Texto para discussão, No. 682, Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), Departamento de Economia, Rio de Janeiro

This Version is available at:

<https://hdl.handle.net/10419/249730>

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.

TEXTO PARA DISCUSSÃO

No. 682

Jumps in Stock Prices: New Insights
from Old Data

James A. Johnson
Marcelo C. Medeiros
Bradley S. Paye



Jumps in Stock Prices: New Insights from Old Data*

James A. Johnson¹, Marcelo C. Medeiros², and Bradley S. Paye³

¹Terry College of Business, University of Georgia, Athens, GA 30602, USA

²Department of Economics, Pontifical Catholic University of Rio de Janeiro (PUC-Rio), Rio de Janeiro, RJ 22451-900, Brazil

³Pamplin College of Business, Virginia Tech, Blacksburg, VA 24061, USA

March 1, 2021

ABSTRACT

We characterize jump dynamics in stock market returns using a novel series of intraday prices covering over 80 years. Jump dynamics vary substantially over time. Trends in jump activity relate to secular shifts in the nature of news. Unscheduled news often involving major wars drives jump activity in early decades, whereas scheduled news and especially news pertaining to monetary policy drives jump activity in recent decades. Jump variation measures forecast excess stock market returns, consistent with theory. Results support models featuring a separate jump factor such that risk premium dynamics are not fully captured by volatility state variables.

*The authors dedicate this paper to the memory of Harold Mulherin, who kindly provided hand-collected Dow Jones data used in this study. We received helpful comments from Torben Andersen, Tim Bollerslev, Roger Edelen, Robert Engle, Jia-Hau Guo (2019 FMA-Europe discussant), George Jiang, Bryan Kelly, Zhengzi Sophia Li, Julien Penasse (2018 EFA discussant), George Tauchen (discussant at Tim Bollerslev's 60th Birthday Conference), Roberto Reno, Allan Timmermann, Viktor Todorov, and seminar and conference participants at the University of Verona, Washington State University, the 2018 European Financial Association Annual Meetings, the 2019 European Conference of the Financial Management Association, the Third International Workshop in Financial Econometrics, and Tim Bollerslev's 60th Birthday Conference. Finally, we are grateful to an anonymous reviewer for a host of comments and suggestions that greatly improved the paper. James Johnson: jamesaj@uga.edu; Marcelo Medeiros: mcm@econ.puc-rio.br; Bradley Paye: bpaye@vt.edu (corresponding author); 540-231-6523 (phone); 540-231-3155 (fax).

Asset prices exhibit occasional discontinuities or ‘jumps’ with important economic implications. There is evidence of a premium associated with bearing non-diversifiable jump risk in equities (Pan (2002)), for example, and jumps appear to be important in explaining the cross section of expected stock returns (Bollerslev, Li, and Todorov (2016), Cremers, Halling, and Weinbaum (2015)). Jump risk has implications for portfolio choice (Liu, Longstaff, and Pan (2003)) and offers a potential resolution of the credit risk puzzle (Zhou (2001)).

Benchmark models permitting jumps, e.g., Merton (1976), assume that jump dynamics are time-invariant. Although analytically convenient, this assumption is unlikely to reflect reality and several studies document variation in jump activity (Andersen, Bollerslev, and Diebold (2007a), Christoffersen, Jacobs, and Ornthanalai (2012), Santa-Clara and Yan (2010), and Tauchen and Zhou (2011)). Time-varying jump risk theoretically generates associated variation in equity and variance risk premia, and recent research emphasizes the importance of the dynamics of left-tail risk in financial markets.¹ Given the significance of dynamic jump risk, it is important to understand the historical pattern of variation in stock market jump activity and the ultimate sources of such variation. Unfortunately, however, the vast majority of existing evidence is culled from options or high-frequency financial data spanning a relatively modest history.

This paper characterizes the dynamics of jump activity in the US stock market over a long horizon spanning nearly 90 years. To do so, we analyze a novel dataset consisting of intraday stock valuation data beginning in the 1930s. These data involve hourly observations for Dow Jones Averages published in the *Wall Street Journal* (WSJ) beginning in 1933. Data are hand-collected until the mid 1990s, when it becomes possible to construct the Dow Jones Averages from the New York Stock Exchange (NYSE) Trade and Quote (TAQ) database. We apply nonparametric methods that decompose total return variation into jump and diffusive

¹See, e.g., Andersen, Fusari, and Todorov (2015), Andersen, Fusari, and Todorov (2017), Bollerslev and Todorov (2011), Bollerslev, Todorov, and Xu (2015), and Kelly and Jiang (2014). A related literature explores the implications of rare large fundamental shocks or ‘disasters’ on asset prices, e.g., Gabaix (2012) and Wachter (2013).

components and identify intraday return intervals exhibiting jumps.²

Beginning with [Schwert \(1989\)](#), a large literature assesses the extent to which secular variation in stock and bond return volatility can be explained by corresponding variation in leverage, macroeconomic volatility, or other determinants (see also, e.g., [Engle, Ghysels, and Sohn \(2013\)](#), [Engle and Rangel \(2008\)](#), and [Paye \(2012\)](#)). These studies do not distinguish between the diffusive and jump components of return variation and therefore it is unclear whether the two components exhibit similar dynamics, especially over long horizons. To shed light on this question, we construct time series measures that separate total return variation into jump and diffusive components, as well as a relative jump measure capturing the proportion of return variation attributable to jumps. Jump and diffusive volatility co-move positively and both series tend to be countercyclical. The two components of return variation are far from perfectly correlated; however, leading to economically significant cyclical variation in the relative jump measure. Moreover, the long-run dynamics of diffusive and jump volatility are quite distinct. Although both series exhibit roughly ‘U-shaped’ long-run dynamics, jump volatility reaches historical lows approximately 10–15 years later than diffusive volatility and recovers more rapidly. The dynamics of the two return variation components differ markedly during a roughly 25-year period from the 1960s through the mid 1980s, but are more similar in recent years. The starkly differing dynamics of jump and diffusive variation during the 1960s causes a stark downward shift in the relative jump measure. This shift is economically large: jump variation accounts for around 30% (14%) of total variation on average before (after) the break.³

²Our intraday data cover a substantially longer history relative to other common intradaily price series. However, sampling is (for the most part) limited to hourly. The econometric test we apply defines an interval as containing a ‘jump’ when the corresponding return magnitude is very large relative to an estimate of prevailing volatility. The sampling limitation inhibits our ability to distinguish a truly discontinuous price movement from an evanescent burst of volatility that occurs within an hour. We view these two events as similar from the perspective of many investors. Indeed, it can be challenging to identify truly discontinuous price movements even with modern ultra-high-frequency data ([Christensen, Oomen, and Podolskij \(2014\)](#)). Our paper therefore focuses on characterizing the dynamics of ‘jump-like’ returns over a long history.

³The post-break average relative jump magnitude we estimate is broadly consistent with evidence in several existing studies based on shorter, more recent sample periods, including, e.g., [Andersen et al. \(2007a\)](#)

Time-variation in jump volatility could be driven by variation in jump intensity or by time-variation in the average magnitude of jumps. We test for jumps on an interval-by-interval basis and characterize time-variation in both aspects of jump activity. Empirically, variation in jump volatility appears to be driven primarily by variation in jump intensity, which reached historical highs in the 1940s and early 1950s before gradually falling to historically low levels during the late 1970s and early 1980s. Jump intensity then rebounds following these lows and fluctuates over recent decades.

We next address the question of what drives jumps and variation in jump dynamics. Our particular goal is to examine whether there have been notable shifts in the types of news that generate jumps in equity markets. To this end, we analyze historical media coverage following large detected jumps over a long horizon. We categorize each jump as driven by either scheduled or unscheduled news and classify the cause of the jump based on the corresponding media discussion. There are three notable findings. First, nearly all of the jump activity in the first half of our long sample is driven by unscheduled news. In contrast, news associated with scheduled macroeconomic and monetary announcements drives a significant fraction of jumps in recent decades. Second, among the specific sources of jumps, news pertains to World War II and the Korean War drives a large proportion of detected jumps during the 1940s and 50s, a period when jump activity is historically high.⁴ In contrast, economic news, including both scheduled announcements and unscheduled economic news, drives the majority of jumps in recent decades. Finally, the proportion of jumps attributable to monetary announcements has increased markedly in recent decades and dominates announcement-related jump activity. This clustering and evolution of announcement-specific jump activity is important given the documented relevance of scheduled announcement days for the time series and cross-section

and [Tauchen and Zhou \(2011\)](#). The finding of a much higher relative jump variation prior to the mid 1960s is, to our knowledge, new to this paper.

⁴The Second World War was the deadliest conflict in human history, and major participants focused their entire economic and industrial capabilities to the war effort. In addition to over 33,000 US casualties, the Korean War had important macroeconomic effects, including higher taxes and GDP growth, lower consumption and investment, and price and wage controls to curb inflation.

of asset returns (Savor and Wilson (2013, 2014); Lucca and Moench (2015)).

The historical perspective afforded by our data provides new insights concerning the evolution of jump risk. Low-frequency dynamics of equity market jumps appear to be driven by two main factors: time-variation in the international peace and security climate, and the evolution of policies and technology associated with the production and dissemination of information to market participants.⁵ Importantly, the stylized facts we document are inconsistent with most benchmark jump-diffusion specifications in the literature. One class of such models tightly links jump intensity with the level of diffusive volatility (Bates (2000), Pan (2002)). Less restrictive models, such as Santa-Clara and Yan (2010), decouple the stochastic evolution of jump activity from diffusive volatility, but do not incorporate trends or structural shifts that appear to explain a significant portion of variation in stock market jump intensity over the past nine decades.

The canonical principle of a ‘risk-return trade-off’ suggests a linkage between risk premium dynamics and variation in underlying risks. Recent theoretical literature hypothesizes distinct roles for the jump and diffusive components of stock price risk in determining the equity premium (Santa-Clara and Yan (2010), Andersen et al. (2015), etc.). We therefore separately analyze the forecasting power of the jump and diffusive components of return variation. We show that the jump component of return variation contains differential predictive information for returns relative to the diffusive component. In particular, we conduct standard in- and out-of-sample predictive regressions for excess stock return on a broad US equity index. When the predictor variable is the diffusive component of return variation, we find little evidence of predictability, consistent with a number of previous studies that test the intertemporal risk-return relation using measures of realized stock return variation. In stark contrast, we find strong evidence of predictive power for the jump component of return

⁵The evolution of Federal Open Market Committee (FOMC) communications provides a prominent example of the second factor. The FOMC did not typically issue statements accompanying changes in the stance of monetary policy until 1994, and we find a stark shift in jump activity associated with FOMC statements following this change.

variation, as well as for the relative jump measure. The slope coefficients on both predictors are positive as expected by theory and the predictive power is economically significant, especially at longer horizons. At a horizon of 5 years, for example, both in- and out-of-sample research designs indicate R^2 -improvements relative to the historical average benchmark on the order of 10–15%. Our evidence concerning the differential predictive content of jump variance measures supports models featuring a separate jump or tail factor such that risk premium dynamics are not fully captured by volatility state variables driving underlying asset price dynamics.

Previous studies that apply high-frequency data to characterize stock market jump activity include [Andersen et al. \(2007a\)](#), [Barndorff-Nielsen and Shephard \(2006\)](#), and [Huang and Tauchen \(2005\)](#), among others. We also contribute to a stream of literature that explores price discovery associated with macroeconomic news (e.g., [Andersen, Bollerslev, Diebold, and Vega \(2003\)](#), [Andersen, Bollerslev, Diebold, and Vega \(2007b\)](#)) and the causes of jumps including scheduled news ([Andersen et al. \(2007a\)](#), [Evans \(2011\)](#), etc.) as well as unscheduled news (e.g., [Lee and Mykland \(2008\)](#)). Our evidence concerning the stock return forecasting power of jump variation relates to papers emphasizing connections between measures of jump or tail risk extracted from options data and the equity premium (e.g., [Andersen et al. \(2015\)](#), [Bollerslev et al. \(2015\)](#), and [Santa-Clara and Yan \(2010\)](#)), as well as measures of aggregate tail risk constructed from the cross section ([Kelly and Jiang \(2014\)](#)). Relative to this body of work concerning jumps, tail risk, and the equity premium, our paper provides a new long-horizon perspective. Whereas many of the aforementioned studies analyze relatively short periods, we analyze a period spanning over eight decades, including the end of the Great Depression era, the second World War, oil prices shocks and stagflation in the 1970s, as well as the recent financial crisis.⁶ This long-horizon perspective delivers fresh

⁶[Manela and Moreira \(2017\)](#) provide an analogous long-horizon perspective regarding implied volatility by constructing a ‘news implied volatility’ (NVIX) measure spanning over a century. Conceptually, our paper differs by focusing specifically on the jump component of return variation. The jump variation measures we construct are not highly correlated with the NVIX measure and possess significant incremental predictive

insights concerning time-variation in jump activity and the associated implications.

1. Data

The most novel data analyzed in our study are intraday data for Dow Jones price indices covering the period 1933–2019. The data consist of open, close, and hourly observations on the Dow Jones Industrial Average and the Dow Jones Composite Average. We analyze two separate sources of intradaily Dow Jones data to establish robustness of key results. The first of these sources consists of hand-collected data from the *Wall Street Journal* and *Barron's* over the period 1933–1993, and analogous data constructed from the New York Stock Exchange (NYSE) Trade and Quote (TAQ) database for the period 1994–2019. The hand collection process for the Dow Jones data involves recording the open, close, and hourly index levels published by the *Wall Street Journal* and/or *Barron's*. Data for the period 1933–1989 were initially hand-collected by Mason Gerety and Harold Mulherin and analyzed in several empirical studies published in the 1990s, including [Gerety and Mulherin \(1991\)](#), [Gerety and Mulherin \(1992\)](#), and [Gerety and Mulherin \(1994\)](#). Data for 1933–1940 are for the DJIA, with the remaining data pertaining to the broader DJCA. We hand-collected additional DJCA data for 1990–1993 from historical records of the *Wall Street Journal* to extend the original data until the time when the New York Stock Exchange (NYSE)'s Trade and Quote (TAQ) database becomes available (see discussion below). The Supplementary Appendix provides a figure illustrating the hand-collection process for Dow Jones data.

We construct data analogous to the hand-collected Dow Jones Composite Average data for the period 1994–2019 using the NYSE TAQ database. Basing intraday prices on TAQ data avoids potential data entry errors associated with the hand-collection process. To do so, we track constituent firms for the DJCA and the actual DJCA close price over the 1994–2019 period. Hourly component prices are based on TAQ transaction prices (following

power for stock returns. The measures are therefore complementary rather than competing.

standard cleaning procedures discussed in [Barndorff-Nielsen, Hansen, Lunde, and Shephard \(2009\)](#)) with the closest time-stamp to the hour. We compute the cumulative price of all index components at the open, close, and on each hour for each trading day, and record an imputed index divisor for the day based on this total price and the actual DJCA value at close. Hourly prices are then calculated from the intraday price levels and the daily divisor. The year 1993 provides an important check on our TAQ-based data construction process. We verify that the TAQ-based data are very similar to the corresponding hand-collected data from the *Wall Street Journal* for this year.

As a robustness check and alternative data source, we also analyze a long history of intraday Dow Jones observations obtained from Global Financial Data (GFD). These data consist of open, close, and hourly data for the DJIA from January 1933 through December 2019. From May 1933 through December 1940, both our main dataset and the alternative dataset from GFD provide intraday observations on the DJIA at the same points in time. Consequently, a comparison of the data is informative concerning the accuracy of data initially hand-collected by Gerety and Mulherin. From 1941 onward, our main data are based on DJCA observations, whereas the GFD data continue to be based on the narrower DJIA.

Institutional features occasionally change during the long history covered by the data. These changes impact the precise nature of recorded data. For example, from 1933 through early 1952, NYSE trading hours were Monday–Friday 10am–3pm, as well as Saturdays 10am–12pm. During this era, index data for a typical weekday consist of the open, close, and 4 hourly price observations, permitting the construction of 5 intraday returns. Saturday trading was discontinued in mid-1952, and the weekday NYSE close moved from 3pm to 3:30pm until 1974. Trading hours extended to 4pm beginning in 1974 and the open moved to 9:30am in 1985.

Table 1 describes various regimes characterized by institutional shifts or variations in the pattern of collected intraday data. Most regime shift points correspond to (persistent) changes in NYSE trading hours. In one case, however, we classify a regime change due to an

alteration in available data. Specifically, from late September 1952 through October 1963, 3PM prices were not published in the *Wall Street Journal*, so that the final intraday return corresponds to a 90 minute return from 2PM to market close at 3:30PM. The table also documents a number of relatively short periods during 1967–1969 with early closings due to the ongoing ‘paperwork crisis’ at the NYSE.

High-quality intradaily trade and quotation data for stocks in the Dow Jones indices are available in recent decades and many prior studies rely on such data. To complement our more novel long history of hourly prices, we collected 5-minute intradaily price data for the Dow Jones Index over the period 1994–2019. We use these data as a point of comparison in order to contrast the behavior of certain empirical measures of ‘realized’ components of return variation using hourly data with analogs constructed using the 5-minute data. The 5-minute Dow Jones price data were obtained from TradeNavigator.

2. Econometric Methods

This section describes the econometric methods we apply in order to decompose return variation into different components. We first describe a setting in which the fundamental (log) index price evolves in continuous time according to a relatively general stochastic process. Price are observed at a specified sampling frequency, e.g., 5-minutes or hourly. We then describe key econometric results pertaining to nonparametric estimation of the diffusive and jump components of return variation, as well as the detection of intervals containing price jumps.

2.1. The Price Process and Observed Intraday Returns

As is standard in the econometric analysis of high-frequency financial data, we assume the (efficient) log price p (Dow Jones average in our context) follows a jump-diffusion process:

$$p_t = p_0 + \int_0^t b_s ds + \int_0^t \sigma_s dW_s + J_t, \quad (1)$$

where b_s and σ_s represent the instantaneous drift and volatility, respectively, W denotes a standard Brownian motion, and J_t represents a pure jump process of finite activity.⁷ Our paper focuses on the role of the jump component J_t of prices. The price jump at time t is given by $J_t = p_t - p_{t-}$, where $p_{t-} \equiv \lim_{s \uparrow t} p_s$. The price process specified by Eq. (1) is relatively general and permits many realistic features associated with stock price evolution. In addition to jumps, the specification permits general forms of stochastic volatility via the process σ_s , time-variation in the drift term b_s , and the so-called leverage effect induced by negative spot correlation between the process W_s and the volatility process.

We assume that the log price process in Eq. (1) is observed at times t_n for $n = 0, \dots, \tilde{N}$, where $t_i < t_j$ for $i < j$ and the times need not be regularly spaced. Returns based on price observations are computed as $r_n \equiv p_{(t_n)} - p_{(t_{n-1})}$, $n = 1, \dots, \tilde{N}$. The notation t_m^O and t_m^C indicates market open and close times for trading day m for $m = 1, \dots, M$. These times naturally partition the price data into M trading days. We designate the returns $r_m^{ON} \equiv p_{(t_m^O)} - p_{(t_{m-1}^C)}$ as overnight returns for trading days $m = 2, \dots, M$. Removing these returns leaves the set of intraday returns r_n , $n = 1, \dots, N$. It will often be convenient to reference the intraday returns corresponding to a particular trading day. The notation $r_{m,i}$ denotes the i -th intraday return on trading day m , and N_m indicates the number of intraday returns on trading day m .

⁷A more technically precise specification of the price process appears in the Supplementary Appendix, along with certain regularity conditions required for jump tests described below.

2.2. Measuring Jump Variation

We aim to separate total return variation into two components, one attributable to continuous price variation and the other to occasional jumps. To achieve these goals we apply nonparametric, high-frequency econometric methods that have become relatively standard in the empirical literature concerning asset prices jumps. Although we focus on jump activity, the underlying econometric methods for identifying jumps and jump variation depend critically on the ability to accurately measure the *diffusive* component of return variation either over a particular time interval or at a particular moment in time (‘spot’ volatility).

We construct lower frequency (monthly and quarterly) time series measures of the extent and proportion of jump variation in prices. These measures are based on underlying daily measures of total and diffusive return variation computed using intradaily data. We first introduce the *realized variance* for each trading day, computed as

$$RV_m \equiv \sum_{i=1}^{N_m} r_{m,i}^2, \quad (2)$$

By the theory of quadratic variation, realized variance converges uniformly in probability to the increment of the quadratic variation process over the trading day as the sampling frequency increases, i.e.,⁸

$$RV_m \xrightarrow{p} \int_{t_m^O}^{t_m^C} \sigma^2(s) ds + \sum_{t_m^O \leq s < t_m^C} |\Delta p_s|^2 \quad \text{as } N_m \rightarrow \infty. \quad (3)$$

The realized variance captures both diffusive variation (the first term on the right-hand side of Eq. (3)) and jump variation (the second term). In contrast, alternative high-frequency estimators provide measures of diffusive volatility that are robust to jumps. We consider two

⁸See Andersen and Bollerslev (1998), Andersen, Bollerslev, Diebold, and Labys (2001), and Barndorff-Nielsen and Shephard (2002).

such measures: realized bipower variation and the MedRV estimator proposed by [Andersen, Dobrev, and Schaumburg \(2012\)](#). These measures are defined as:

$$RBV_m = \frac{\pi}{2} \left(\frac{N_m}{N_m - 1} \right) \sum_{i=2}^{N_m} |r_{m,i-1}| |r_{m,i}|, \quad (4)$$

$$MedRV_m = \frac{\pi}{6 - 4\sqrt{3} + \pi} \left(\frac{N_m}{N_m - 2} \right) \sum_{i=2}^{N_m-1} \text{med}(|r_{m,i-1}| |r_{m,i}| |r_{m,i+1}|)^2. \quad (5)$$

Both RBV_m and $MedRV_m$ consistently estimates the diffusive component of return variation as the sampling frequency increases, i.e., $RBV_m \xrightarrow{p} \int_{t_0^C}^{t_m^C} \sigma^2(s) ds$, ([Barndorff-Nielsen and Shephard \(2004\)](#), [Barndorff-Nielsen and Shephard \(2006\)](#), [Andersen et al. \(2012\)](#)). However, [Andersen et al. \(2012\)](#) show that $MedRV_m$ has better theoretical efficiency properties and improved finite-sample robustness to jumps and ‘zero’ returns. As these features are quite desirable in our contest, we focus on $MedRV_m$ as the primary diffusive volatility measure, but explore robustness to the use of the more tradition BPV_m measure.

Given the realized variance measure of total return variation and a (consistent) measure of diffusive return variation, a consistent measure of jump variation obtains as the difference. For example, the quantity $RJV_m \equiv RV_m - MedRV_m$ (realized jump variation) provides a natural consistent estimate of the jump variance on trading day m .

Undoubtedly, realized variation measures computed using a relatively small number of hourly returns will be a relatively noisy proxies for true daily variation. Consequently, we do not focus on analyzing the properties of noisy daily jump return variation measures. Instead, we aggregate daily jump variation and related return variation measures to construct less noisy measures at monthly, quarterly, and annual frequencies. The series $MRJV$, $QRJV$, and $ARJV$ denotes the monthly, quarterly, and annual time series, respectively, constructed by summing the daily jump variation measure RJV_m over the corresponding periods. Realized variance and integrated (diffusive) variance are aggregated similarly. Finally, we define a relative jump (RJ) measure capturing the proportional contribution of jumps to total re-

turn variation. For example, the quarterly relative jump measure is defined as the ratio of quarterly jump variance to quarterly realized variance:

$$QRJ \equiv \frac{QJV}{QRV}. \quad (6)$$

2.3. Testing for Specific Jumps

Our second empirical goal is to identify *specific* intradaily time intervals in which price jumps occur and corresponding estimates of the associated jump magnitudes. In order to test for specific jumps, we apply nonparametric ‘threshold’ methods using high-frequency data (Mancini (2001)). The intuition behind the approach involves the fact that diffusive return moments are of order $\sqrt{\Delta}$, where Δ designates the return interval duration. When a jump occurs, the return over a small interval including the jump is dominated by the jump component, and this motivates a test that designates intervals as containing a jump whenever:

$$|r_n| > k\Delta_n^\omega \quad (7)$$

where $\omega \in (0, 0.49)$ and $k > 0$. Under relatively mild assumptions on the price process, it can be shown that such a test consistently identifies the intervals containing jumps as the sampling frequency increases, i.e., as $\Delta_n \rightarrow 0$. (See the Supplementary Appendix for technical details.) However, the asymptotic result leaves important practical choices (ω and k) to the discretion of the researcher. In applications, it is common to refine the general approach by designating intervals as containing a jump whenever

$$|r_n| > \alpha \hat{\sigma}_n \Delta_n^{0.49}, \quad (8)$$

where $\hat{\sigma}_n$ denotes an estimate of the local or ‘spot’ volatility and $\alpha > 0$. Noting that $\Delta_n^{0.49} \approx \sqrt{\Delta}$ and rearranging Eq. (8) gives an interpretation of α as the number of standard

deviations, based on the spot volatility estimate, required for jump designation.

Two key implementation details are the method for estimating the spot variance $\hat{\sigma}_n$ and the choice of α . Estimates of spot volatility are typically based on jump-robust diffusive variation measures computed using returns ‘local’ to the interval in question. We adopt a spot volatility estimator of the form:

$$\hat{\sigma}_n(B) = \max(\tau_n \overline{\text{MedRV}}(n, B), \sigma_{\text{MIN}}), \quad (9)$$

where τ_n represents a time-of-day adjustment factor, $\overline{\text{MedRV}}(n, B)$ denotes the average daily MedRV measure of Eq. (5) computed over a window of B trading days centered around the day containing return interval n , and σ_{MIN} is a minimum volatility threshold.

Eq. (9) involves three key components. First, the ‘local’ volatility component $\overline{\text{MedRV}}(n, B)$ averages MedRV over a window of trading days centered around the day corresponding to return interval in question. When $B = 0$, this simply equals the bipower variation for the corresponding trading day. Because we observe returns at a relatively low hourly frequency, $B = 0$ results in a very noisy estimate of local volatility. For this reason, we set $B > 0$, which has the effect of including the bipower variation from nearby trading days to reduce estimation noise. The empirical section further discusses the choice of B for our main test results, as well as robustness considerations.

The second component is the time of day adjustment factor τ_n , which follows numerous studies that account for diurnal effects in volatility in jump tests similar to ours. It is well-known that volatility tends to be ‘U-shaped’ over the trading day: higher near the open and close and lower during the middle of the trading day (Andersen and Bollerslev (1997)). Our implementation of the adjustment factor, described in detail in the Supplementary Appendix, is relatively simple: we adjust the general spot volatility estimate up or down depending on the ratio of the average volatility for a particular trading hour relative to the sum or the average volatilities across all trading hours. The time of day adjustment factors

are computed separately across the different trading regimes in our data.

The final component that appears in Eq. (9) is a minimum volatility threshold σ_{MIN} . The motivation for incorporating a minimum threshold relates to the noise associated with local volatility estimates. In a repeated testing context, estimation noise will cause some local volatility estimates to be very low. This can lead to spurious and/or economically insignificant jumps during such intervals. We calibrate σ_{MIN} such that hourly jump returns must equal at least 50 basis points in magnitude.⁹

The second crucial implementation detail for the jump test is the choice of α . Increasing the threshold reduces the likelihood of spuriously designating a diffusive price movement as a jump, but also reduces the power to detect true jumps. [Lee and Mykland \(2008\)](#) pose a jump test closely related to Eq. (8) and specify a formula for selecting the threshold value α based on Gaussian extreme value theory, given a specified significance level for the test. To provide some intuition, in the absence of jumps the return over a short interval will be approximately Gaussian. It might therefore seem natural to choose $\alpha = 3$ for a test with significance level around 1%. However, the test is applied to many intraday intervals (over 1,000 per year in our setting). It is thus necessary to select a higher threshold to maintain the desired significance level in this repeated testing context. In the Supplementary Appendix, we show that the recommended threshold derived by [Lee and Mykland \(2008\)](#) implies an α value just over 4 in our empirical setting. Baseline results therefore set $\alpha = 4$, but we explore robustness to alternative choices.

⁹We are grateful to George Tauchen for suggesting this modification, which follows a similar modification in [Li, Todorov, and Tauchen \(2018\)](#). Our main findings concerning time-varying jump activity are robust to modest perturbations in the minimum volatility level, or even to omitting this aspect. In the latter case, we identify even higher rates of jump activity in the 1940s and 1950s, both absolutely and relative to subsequent decades. From this perspective, our imposition of the minimum volatility level is conservative.

2.4. *Finite Sample Performance and Microstructure Effects*

Although daily return variation measure based on hourly data are quite noisy, aggregating to monthly or longer horizon measures mitigates the associated effects. However, the hourly sampling limitation is likely to have a greater impact on our interval-by-interval jump tests. Various studies examine the finite sample properties of jump tests similar to ours in simulation environments (see, e.g., [Lee and Mykland \(2008\)](#), [Huang and Tauchen \(2005\)](#) and [Davies and Tauchen \(2018\)](#)). Results in [Lee and Mykland \(2008\)](#) are particularly notable, as they explicitly include the case of hourly sampling, contrasting this case with both higher frequency (15-minute) and lower frequency (daily) sampling. Simulation results in [Lee and Mykland \(2008\)](#) highlight both the pros and cons of basing a jumps analysis on hourly data. Hourly returns provide much greater power to detect jumps relative to daily returns, with substantially lower rates of spurious jump detection.¹⁰ Indeed, this fact that motivates much of our analysis, as our novel long history of intraday Dow Jones prices affords an opportunity to conduct a meaningful analysis of jump activity over a very long history. On the other hand, as predicted by theory, tests based on hourly returns have reduced power and generate more spurious jumps than tests using returns sampled at higher frequencies.

The reduced power associated with using lower frequency returns is concentrated toward cases where the underlying jump is relatively small in magnitude ([Lee and Mykland \(2008\)](#)). Consequently, our tests using hourly returns will have limited power to detect small jumps. Power will concentrate in the direction of detecting economically large jumps that dominate return variation over a 60-minute interval. In addition, the test may spuriously designate short-lived ‘spikes’ in diffusive volatility as jumps. Although the distinction between truly discontinuous price jumps and short-lived volatility spikes can be important in certain contexts involving, e.g., technical issues of market completeness, both types of price movements can plausibly be viewed as ‘jump-like’ from the perspective of many typical investors. As

¹⁰Our discussion refers specifically to Tables 1 and 2 in [Lee and Mykland \(2008\)](#).

subsequent results reveal, the the ‘jumps’ detected by our test reflect economically large price movements that often receive attention in contemporaneous news media. We provide novel evidence regarding the secular dynamics of these important price movements, recognizing the aforementioned limitations associated with hourly sampling.

A final concern involves the effects of microstructure effects, which can create bias in high-frequency return variation measures. These effects are typically strongest when returns are sampled at very high frequencies. The fact that our returns are hourly and involve only Dow Jones stocks with relatively high trading volume and liquidity should reduce the impact of microstructure biases in our analysis. Still, our analysis spans nearly a century and includes periods when trading volume was relatively limited. Therefore, we explore this issue in the empirical work that follows.

3. Jump Dynamics Over a Long History

This section presents our main empirical results regarding time-variation in jump dynamics.

3.1. Properties of intraday returns and daily return variation measures

Table 2 shows descriptive statistics summarizing properties of intraday returns and key daily return variation measures. To conserve space, we report illustrative results from three 8-year periods, one near the beginning of our sample (Panel A), one near the middle (Panel B), and one toward the end of the sample that includes the financial crisis period (Panel C). The first two columns of Table 2 show, for selected years, the total number of intraday returns during the year and the mean number of returns per trading day. The mean number of returns per day does not always equal an integer (the median) due to occasionally abnormal trading days. This effect is most pronounced during the ‘paperwork’ crisis the late 1960s,

reflected in Panel B results. The next three columns show the mean and median of the intradaily returns for the year, as well as the proportion of these returns that are ‘zeros.’ The latter quantity is partially indicative of the severity of microstructure effects and is discussed in further detail later.

The right-hand portion of Table 2 shows within-year averages of key daily return variation measures RV, MedRV, and BPV for the corresponding year, as well as the within-year average of daily return sample autocorrelation. Theoretically, total return variation must exceed diffusive return variation, with the difference equaling jump variation. Consistent with this pattern, the average annual realized variance is typically greater than average MedRV or BPV. BPV usually exceeds MedRV, with more material differences in high volatility years. Note that, in several years during the 2000s (Panel C), the BPV exceeds even RV, implying a negative estimate of jump variation for the corresponding year. This illustrates why we prefer to use MedRV, which does not exhibit this behavior, as the primary measure of daily diffusive volatility. The average annual autocorrelation is typically negative. This statistic fluctuates from year to year but appears roughly stable over our long sample.

3.2. Jump Return Variation Over Time

We first compare the properties of lower frequency return variation measures computed using the hourly intradaily returns with measures based on 5-minute returns for the period of overlap between the measures (1994–2019). Panel A of Figure 1 contrasts quarterly realized volatility using the hourly and 5-minute returns. Panel B contrasts quarterly diffusive volatility based on the MedRV measure. In both cases, the series are extremely similar. This shows that, following aggregation of noisy daily measures to the quarterly frequency, hourly returns produce realized volatility and diffusive volatility measures that are very similar to more ideal measures computed using 5-minute returns.

Panel C of Figure 1 compares quarterly jump return volatility, computed as the square

root of the quarterly difference between realized volatility and MedRV, for the hourly returns and 5-minute returns. This panel indicates that, despite aggregating noisy daily estimates at the quarterly level, the jump volatility series based on hourly returns exhibits material differences relative to the corresponding series based on 5-minute returns. Specifically, the series based on hourly returns is more erratic and prone to occasional ‘outlier’ quarters. Smoothing in the form of a two-year (8 quarter) backward moving average (shown in Panel D) reveals that, despite noise, jump volatility measures based on hourly data reveal quite similar dynamics at the ‘business cycle’ level to those exhibited by measures computed using 5-minute returns. We conclude that, although quarterly jump volatility measures using hourly data are somewhat erratic, this time series is informative regarding jump dynamics, particularly at the business cycle and lower frequency. Motivated by the contrast of Panels C and D in Figure 1, we will often consider smoothed versions of the jump volatility measures in subsequent analysis.

Having verified the informativeness of jump volatility measures based on hourly returns, we now examine return variation measures over a long history. A first question concerns whether jump and diffusive volatility exhibit similar secular variation. [Schwert \(1989\)](#) and several recent contributions explore the macroeconomic determinants of low frequency variation in stock return volatility (see, e.g., [Engle and Rangel \(2008\)](#), [Engle et al. \(2013\)](#), and [Paye \(2012\)](#)). Total stock return volatility largely reflects the diffusive component. Therefore, it is unclear whether the jump component of return variation exhibits similar dynamics.

Figure 2 addresses this question by plotting time series of jump and diffusive variation measures over the period 1933–2019. Panel A contrasts jump volatility with diffusive volatility. The displayed series are two-year backward moving averages of the underlying quarterly measures in order to reduce the impact of noise and focus on business cycle variation. The plot indicates a nuanced relation between jump volatility (plotted against the left-hand axis) and diffusive volatility (plotted against the right-hand axis). On the one hand, there appears to be a reasonably tight connection between the cyclical dynamics of both series, especially

toward the beginning and end of the long sample period. Both jump and diffusive volatility increase around the financial crisis period, for example. However, the relation between jump and diffusive volatility appears much less tight during a period from the mid 1960s through the early 1980s. For example, diffusive volatility rises substantially during the oil price shocks of the mid 1970s and associated economic contractions, but jump volatility exhibits little response. Finally, there is a noticeable difference in long-run jump and diffusive volatility dynamics during the first half of the sample. Diffusive volatility peaks around 1939, at the end of the Great Depression era and the onset of World War II, and then falls gradually during the 1940s 50s, and early 1960s. In contrast, the level of jump volatility remains fairly steady during this period, until decreasing more abruptly in the mid 1960s.

Panel B of Figure depicts the relative jump measure (RJ). Similar to the volatility series in Panel A, the series plotted is a two-year backward moving average of the underlying quarterly RJ series. The dynamics of the relative jump measure are economically interesting because they shed direct light on the validity of models of time-varying return dynamics that tie together variation in jump dynamics with variation in diffusive volatility. There are several key insights from Panel B. First, the relative jump measure exhibits cyclical fluctuation, indicating time-variation in the *relative* dynamics of the jump and diffusive components of return variation. In contrast to the levels of jump and diffusive volatility depicted in Panel A; however, the relative jump measure does not appear to exhibit a strong relation to macroeconomic conditions. For example, the relative jump measure does not increase substantially during the financial crisis period owing to the fact that diffusive and jump volatility increase by roughly the same proportion at this time (Panel A). A second insight is that there appears to be a significant structural break or shift in the mid 1960s such that jumps account for a relatively lower fraction of return variation. The level shift is economically significant: jump variation comprises approximately 16–17% of total return variation on average prior to the early 1960s, and approximately 6–8% of return variation thereafter. The shift appears to be primarily driven by a reduction in jump volatility at this

time. Although jump volatility recovers in subsequent decades, diffusive volatility increases in tandem, making the reduction in the level of RJ appear to be permanent. Finally, the *volatility* of RJ has decreased in recent decades. This reduction in volatility is consistent with the tight relation between jump and diffusive volatility dynamics over the past two decades depicted in Panel A.¹¹

Table 3 characterizes the properties of quarterly jump volatility, diffusive volatility, and the relative jump measure, and compares these series with other financial risk measures such as bond spreads and option-based risk measures. Jump volatility is persistent, positively skewed and fat-tailed, similar to diffusive volatility and option-implied volatility measures (VIX and NVIX).¹² Panel B shows correlations of jump volatility, RJ , and diffusive volatility with other variables. The correlation between jump and diffusive volatility is approximately 0.69 – strongly positive but also indicating a far from perfect relation between the series. The correlations of jump and diffusive volatility with bond spreads, the tail-risk measure, and option-based risk measures tend to be similar. The relative jump measure is (mechanically) positively correlated with jump volatility and negatively correlated with diffusive volatility, but does not exhibit strong correlation with the other comparison variables. Panel C illustrates business cycle properties of variables by report the slope coefficient, t -statistic, and R^2 from a regression of the corresponding series on the quarterly NBER recession indicator time series. All of the return variation and option-based volatility measures exhibit significant countercyclical relations. The relative jump measure does not exhibit a significant relation with the business cycle as captured by the recession indicator (nor does the term spread or variance risk premium in our samples).

¹¹The Supplementary Appendix provides a four-panel version of Figure 2 that also includes plots of the raw (unsmoothed) quarterly jump and diffusive volatility series and a plot that emphasizes the very low frequency component of jump and diffusive volatility by applying a 16-year moving average filter. The low-frequency comparison shows that, although low-frequency components for both series are roughly ‘U-shaped,’ the long-run component of diffusive volatility reaches historical lows in the 1960s whereas long-run jump volatility reaches lows in the early 1980s.

¹²The sample autocorrelation for jump volatility is significant and positive but relatively low (0.14). However, this value is likely biased downward due to measurement noise associated with quarterly jump variance.

3.3. *Jump test results*

This section discusses results for interval-by-interval tests for jumps. We focus primarily on results for the DJCA hand-collected prices, setting $\alpha = 4$ and $B = 2$, implying that spot volatility is computed as the average over five days (the current trading day ± 2 trading days). This choice is motivated by a comparison of the properties of the spot volatility computed using our procedure with an alternative ‘ideal’ spot volatility computed using 5-minute returns for the current trading day (only). Figure 3 illustrates how spot volatility estimates computed using hourly data compare with the high-frequency benchmark. Panel A compares spot volatility estimates computed only using returns within the trading day in question ($B = 0$). The spot volatility estimates using hourly data roughly track those based on 5-minute returns, but are clearly somewhat noisy. Panel B shows the much better fit for spot volatility estimates using hourly data with $B = 2$. These estimates closely track the more ideal spot volatility based 5-minute returns.¹³

Panel A of Figure 4 shows the detected jumps over our long sample period. There is clearly time-variation in the intensity of jumps, as the density of jumps appears during certain periods, such as in the late 1930s and 1940s, relative to others. The distribution of jump returns appears to be fat-tailed: most jumps are on the order of 50–100 basis points, but there are a number of very large jumps associated with hourly absolute returns in excess than 3%. Panel B plots the thresholds required for jump designation, i.e., the time series of absolute hourly return threshold required for designation as a jump. During certain periods, such as at the peak of the financial crisis the latter portion of 2008 and early 2009, the level of spot volatility is so high that few jumps are detected. However, other periods with low jump activity, such as during the mid 1970s, are not associated with spot volatility spikes. Panel B illustrates that the minimum volatility threshold σ_{\min} binds for significant portions

¹³The choice $B = 2$ provides the best approximation to the high-frequency spot volatility estimates over the period where 5-minute returns are available in the sense of achieving the highest R^2 -value in a regression of daily high-frequency spot volatility estimates on hourly analogs for different choices of B .

of the sample, particularly in the 1940s through the early 1960s, when the level of diffusive volatility is low by historical standards.

Figure 5 plots rolling estimates of key jump activity parameters, including the jump intensity rate, measured as the percentage of jumps per intraday interval, the mean absolute jump return, and the mean jump return. Estimates are based on a rolling window of approximately three years (the window length is $3 \times 250 \times 7 = 5,292$ intraday intervals). Dotted lines show a 95% confidence interval for the corresponding rolling estimates. Dashed, vertical lines indicate regime change points for the underlying Dow Jones intraday returns data (see Table 1). The time-variation in jump intensity rates is statistically as well as economically significant. Peak jump intensity occurs near the beginning of our sample and during US involvement in World War II. The jump rate falls in the aftermath of the war, and then increases again in the late 1940s and early 1950s. Jump intensity then falls gradually over the next 20 years before reaching historical lows in the late 1960s. Jump intensity then cycles over more recent decades, with another peak occurring in the late 1980s. Fluctuations in jump intensity appear to be more modest following the turn of the century.

Rolling jump magnitude estimates (Panel B) are relatively steady over the long period we examine.¹⁴ Comparing Panels A and B reveals that most time-variation in jump volatility is driven by corresponding time-variation in jump intensity rates, and not by time-variation in the volatility of jump returns conditional on jumps occurring. Average jump returns (Panel C) are relatively stable in the early and late portions of our sample, but there is some evidence of dynamics in average jump returns during the 1960s through the 1980s.

¹⁴The upward blip in average jump magnitude in the mid 1960s is in part driven by a jump in May of 1962 of over 6%, which corresponds to the single largest intraday jump in our data. The jump is associated with the ‘flash crash of 1962’ in which market prices plunged on May 28 of that year, only to largely recover the following day. Interestingly, our data illustrate that the recovery was largely concentrated in a single hour toward the *end* of the day on May 29. Although we filter for ‘bounce backs’ (see discussion below), this jump remains due to the time delay between the price decreases on May 29 and the sudden large increase late in the day on May 29.

3.4. *Microstructure Effects and Other Robustness Checks*

Apparent time-variation in jump intensity could derive from variation in different types of market frictions. One potential manifestation of microstructure effects involves the possibility of incomplete overnight price discovery. Specifically, jump activity might appear higher in earlier decades because price discovery related to overnight news bleeds into the following trading day. Reductions in the severity of market frictions over time could explain an apparent drop-off in jump rates, as overnight news is more efficiently incorporated at market open. This conjecture would imply that a disproportionate share of jumps occur during the initial trading interval during early decades. This is not the case in the data, however. A second possibility is that market frictions drive spurious variation in jump activity by imparting time-varying bias on spot volatility estimates and corresponding jump thresholds. Figure 6 shows how two measures of frictions evolve over time. First, we plot for each year the average daily return autocorrelation statistic. Intraday returns exhibit modest negative serial correlation. This correlation is quite stable over time; however, and therefore time-varying bias in spot volatility estimates is unlikely to explain our main results.¹⁵ There is significant time-variation in the proportion of zero returns. Not surprisingly, this exhibits a downward trend with the exception of the first few years of the sample. In general, the secular dynamics of zero return frequencies do not closely track the dynamics of jump intensity that we document.

Recent literature on measuring return variation using high-frequency data emphasizes the role of certain price path anomalies known as ‘drift bursts’, defined as brief, locally explosive trends in asset price paths (Christensen, Oomen, and Renò (2020)) and ‘slow jumps’ in which significant, short-term price movements appear as gradual, trend-like adjustments in very high-frequency data (Barndorff-Nielsen et al. (2009)). Concerning these phenomena, our

¹⁵Another potential concern is that variation in jump activity could be attributable to discrete changes in NYSE trading hours, but Figure 5 illustrates that the most significant shifts in jump activity do not rapidly follow NYSE regime shift times.

lower hourly sampling frequency likely insulates our analysis relative to a setting with ultra high-frequency data. For example, a ‘slow jump’ that occurs over fifteen minutes of trading is likely to be simply deemed a jump in our analysis using hourly data. Similarly, evanescent flash crashes that occur within an hour of trading are unlikely to be deemed jumps. As a specific example, the May 6, 2010 flash crash (see, e.g., [Easley, De Prado, and OHara \(2011\)](#)), in which the DJIA plunged around 9% around 2:30 PM only to recover shortly thereafter, is *not* detected as a jump in our analysis because the corresponding hourly absolute return does not exceed the spot volatility threshold.¹⁶

We conduct a number of additional robustness checks related to aspects of the jump tests. The Supplementary Appendix shows plots of rolling jump activity for alternative jump test thresholds. Mechanically, a higher threshold yields fewer jumps and a lower jump intensity in each decade, with opposite effects for a lower threshold. However, altering the threshold does not change our key insights concerning time-variation in jump intensity. The Supplementary Appendix also shows that we obtain similar results for modest variations in the bandwidth B used to construct the spot volatility estimate.

3.5. *Economic Implications*

We document economically significant low-frequency variation in jump activity and jump characteristics over a period of more than eight decades. What are the economic implications of this evidence? A first implication is that asset pricing models assuming static jump risks, such as the classic Merton jump-diffusion model ([Merton \(1976\)](#)), are inconsistent

¹⁶Misrecorded prices can also generate spurious ‘jumps’ and time-variation in data quality could therefore create the false appearance of time-variation in jump activity. This is unlikely to drive our results for several reasons. First, we audited underlying index levels around jumps followed by significant price movements in the opposite direction using Wall Street Journal records from ProQuest. This uncovered several spurious jumps due to initially misrecorded prices and reported results are based on corrected data. Second, we filter out ‘bounce backs,’ defined as jumps immediately followed by a subsequent jump in the opposite direction. Finally, we obtain similar results concerning variation in jump intensity using two different datasets: the hand-collected DJCA data and the GFD DJIA data.

with the data. We show that the *particular* aspect of constant jump risk most strongly violated in the data is the assumption of a time-invariant jump arrival or intensity rate. Our evidence on this point complements and extends earlier evidence of time-variation in jump activity (e.g., [Huang and Tauchen \(2005\)](#)). Our results also speak to more subtle aspects of models permitting both jumps and stochastic volatility. Many such models closely tie variation in jump activity with variation in the diffusive component of volatility (e.g., [Bates \(2000\)](#), [Pan \(2002\)](#)). Such an assumption increases tractability by reducing the dimensionality of risks faced by investors. But is it empirically plausible? We show that, from a low-frequency perspective, the assumption is not supported by the data and there have been time periods such as the 1960s and 1970s during which the dynamics of jump and diffusive variation behave very differently. In particular, an empirically realistic description of asset return dynamics should allow for long-lived deviations in jump intensity and/or jump variation relative to fluctuations in volatility. However, we also show that, within some sub-periods, and in particular over the past two decades, there is strong co-movement between jump and diffusive volatility. This tighter relation in recent years suggests that, in certain contexts, tying the dynamics of jump intensity with volatility might be reasonable. Of course, an important question is whether jump and volatility dynamics will uncouple once again in the future. We briefly return to this question in the concluding section.

4. What Drives Variation in Jump Activity?

Although there is general consensus regarding the presence of jumps in asset prices (reinforced by our evidence), the causes of jumps are less well-understood. Theory suggests that jumps should be associated with the release of price-relevant news. A number of papers link detected jumps with scheduled macroeconomic and monetary policy news events. For example, [Andersen et al. \(2007a\)](#) and [Evans \(2011\)](#) link detected jumps for equity, bonds, and exchange rates with scheduled macroeconomic news. There is also evidence that *unsched-*

uled news drives asset price jumps (Lee and Mykland (2008), Prokopczuk and Wese Simen (2015)). But it is unclear whether the causes of jumps appear to be stable over time, or whether there have been important shifts over time in the relation between jumps and the associated the flow of financial and macroeconomic information.

In order to analyze potential shifts in the causes driving jumps over a long history, we hand-collect information from the *Wall Street Journal* for publication days following detected jumps. We do this for the largest 100 jumps (measured by absolute jump return) for the subset of intraday jumps identified using a higher threshold of $\alpha = 5$. These screens yield a set of events that are both statistically most likely to be jumps based on our test and also sufficiently large in magnitude such that the corresponding price movements are most likely to receive media attention that will be informative regarding the cause of the jump.¹⁷

For each detected jump, we assess whether contemporaneous media coverage discusses price movements plausibly corresponding to the jump. This turns out to be the case for all investigated jumps, which likely reflects the fact that we focus on a subset of the economically and statistically most significant jumps detected by the test. We classify the jump as attributable to ‘scheduled news,’ if the media coverage refers to a scheduled announcement, and otherwise classify the jump as attributable to unscheduled news. We also break out a subset of jumps attributed to scheduled news associated with Federal Reserve announcements, and any unscheduled jumps attributed to monetary policy news. We split our long sample period into two roughly equal sub-samples: 1933–1975 and 1976–2019 in order to compare with the sources of jumps have changed over the nearly 90-year period of our sample. Panel A of Table 4 shows results.

Our earlier results establish that jump intensity was historically highest during the 1930s-

¹⁷We focus on *Wall Street Journal* coverage in order to provide a consistent method of assessing potential sources of jumps over our long sample period. Numerous media sources provide coverage over the last few decades. However, including these sources, which are not available for earlier periods, might bias our findings regarding time-variation in the source of jumps. That said, in certain cases in which a next-day edition of the *Wall Street Journal* is unavailable, we instead examine coverage in the financial section of the *New York Times*.

1950s, a time period that predates the modern monetary policy era and largely predates the modern practice of regular releases of standardized macroeconomic data. This fact implicitly suggests that jump activity prior to the 1960s is primarily driven by *unscheduled* news. The Table 4 results confirm this. Over the full sample, media coverage associates approximately 17% of jumps as driven by scheduled news. However, virtually none of these jumps occur during the first sub-period, whereas the nearly a third of the researched jumps during the second half of our sample are associated with scheduled announcements. Moreover, monetary policy announcements appear to be particularly important drivers of jumps in recent decades: over half of the researched jumps associated with scheduled announcements involve Fed announcements.

Panel B classifies the media discussion associated with each jump into one of six categories: 1) Fed/monetary policy; 2) economic news; 3) industry/commodities news; 4) political news; 5) war or conflict news; or 6) other.¹⁸ Economic news includes scheduled announcements and data releases outside of Fed announcements, as well as other information focused on macroeconomic conditions other than monetary policy news. Industry/Commodities news reflects cases in which the media attributes a *market* price movement to news concerning commodities or a particular industry. Political news includes election news and other political developments. War news is self-explanatory. The ‘other’ category captures all remaining jumps not assigned to one of the aforementioned categories.¹⁹

War news drives a significant fraction of jumps during the early portion of our sample, but no jumps thereafter. Given that World War II, the Korean War, and the Vietnam War occur during this period, this is perhaps not surprising, but demonstrates that war and conflict events generate significant jump activity. The results therefore suggest that the

¹⁸Jumps in equity prices driven by other types of news, e.g., war news, are presumably ultimately driven by the economic ramifications of such news. Designation of jumps as attributable to ‘economic news’ implies that the underlying media coverage focuses on news regarding a particular measure of economic conditions.

¹⁹The Supplementary Appendix provides an expanded discussion of the criteria for assigning jumps. Detailed notes associated with each jump are available from the authors upon request.

transition to the cold war period following the Korean War and a more peaceful, if tenuous, international atmosphere contributed to a reduction in stock market jump activity.

Monetary policy and economic news drives a much higher proportion of jumps in recent decades. This is likely related to shifts in policies and technology associated with the production and dissemination of economic data. An example of such a shift involves the Federal Reserve’s decisions to follow FOMC meetings with public statements and a related trend towards greater forward guidance concerning the conduct of monetary policy. This shift appears to be particularly important for stock market jumps. Panel B shows that over 25% of researched jumps in the second portion of the sample period are attributed to Fed/monetary policy news.²⁰

A final takeaway pertains to jumps classified in the ‘other’ category in Table 4. Media coverage following such jumps either does not point to a specific piece of news as causing the jump, or explains the price movement as a sudden response to past price pressure rather than new information. In a few cases, media coverage explicitly states that there is no clear cause for the price movement. Almost 30% of researched jumps fall in this category, which is notable given the restrictive criteria for researched jumps (highly significant and large jumps). The fact that a significant fraction of large, sudden price movements occur without apparent associated news echos similar conclusions in earlier studies that focus on much shorter periods (e.g., [Cutler, Poterba, and Summers \(1989\)](#)).

5. Jump Variation and the Equity Premium

This section analyzes time series measures of the jump and diffusive components of return variation and their relation with the equity premium. Benchmark asset pricing models

²⁰In additional unreported results, we find that, since 1990, over 10% of all detected jumps using our baseline test coincide with FOMC statement or minutes release days (detailed tabulation available upon request). This likely understates the actual proportion of jumps driven by monetary news, as it excludes other potential Fed-related jumps associated with speeches and monetary policy news outside of statement or minutes release days.

often imply a linear relation between the conditional market risk premium and conditional market variance, e.g., the intertemporal capital asset pricing model (ICAPM) of [Merton \(1973\)](#). When equity prices are subject to both diffusive and jump variation; however, equity premium dynamics can be more complicated, particularly under models in which jump dynamics are driven by a source of risk distinct from diffusive volatility. Under such a jump-diffusion model, [Santa-Clara and Yan \(2010\)](#) show that equity premium dynamics depend upon the stochastic evolution of both jump intensity and diffusive volatility (and their interaction). [Andersen et al. \(2015\)](#) specify a rich parametric model that incorporates jumps in both prices and volatility, separate dynamics for left and right jump intensities, and cross-excitation in volatility and jump processes. Their model also delivers a spot equity premium that reflects both diffusive and jump risks, where the evolution of the latter is not tied tightly to the former.

Motivated by this theoretical literature, we test whether jump variation measures explain time-variation in the equity premium, conditional on measures of the diffusive component of return variation and other prominent stock return forecasting variables. Our long time series of jump and return measures afford particular advantages in a return-forecasting setting, because they permit us to examine predictability over a range of horizons. This is important because recent literature establishes stronger evidence of the theoretically expected positive relation by focusing on longer-run components of returns and (total) return variance ([Bandi and Perron \(2008\)](#), [Bandi, Bretscher, and Tamoni \(2018\)](#)).

Table 5 contrasts the stock return forecasting power of diffusive volatility with that of jump volatility and the relative jump measure of over a long history. The table shows results for predictive regressions of annual log excess returns on the value-weighted CRSP portfolio on lagged annual return variation measures. We consider forecasting horizons of one to eight years. Following, e.g., [Bandi and Perron \(2008\)](#), the diffusive and jump variation measures used in the regression are log transformed.

Panel A of Table 5 shows results for diffusive return variation (IV), measured using

daily MedRV aggregated over the calendar year. There is little evidence that the diffusive component of volatility forecasts returns. Slope coefficients for IV are generally insignificant and economically small. These results are consistent with a number of prior studies that examine the stock return forecasting power of rolling measures of volatility (e.g., [French, Schwert, and Stambaugh \(1987\)](#)). Panel B shows that, in start contrast to diffusive return variation, there is much stronger evidence of forecasting power for the jump component of return variation. The slope coefficients in these regressions are always positive and are statistically significant at most horizons based on the Newey-West approach with lags equal to twice the forecasting horizon to account for serial correlation associated with overlapping data. The predictive power of jump variation is economically significant, particularly at longer horizons. For example, the R^2 -value at a horizon of 5 years equals almost 10%, and increases further to around 16% at an 8-year horizon. Panel C shows similar strong evidence of predictive power for the relative jump measure RJ . Again, slope estimates are positive at all horizons and generally statistically significant. The R^2 -values for the relative jump measure, e.g., just over 10% at a horizon of 5 years, indicate economically significant forecasting power.

Before proceeding to an out-of-sample research design, we note that the Supplementary Appendix includes additional results as robustness checks. For example, we find similar results using quarterly data (forecasting at the same horizons) in which we employ smoothed jump variation measures using two years of backward aggregation, as depicted in [Figure 2](#). We also obtain qualitatively similar results using return variation measures based on the alternative GFD DJIA intradaily returns. Finally, we show results for a set of bivariate predictive regressions that include various alternative popular stock return forecasting measures from prior literature in order to show that our jump return variation measures do not simply reflect information in, e.g., the default or term spread or other conventional predictors.

[Table 6](#) presents stock return forecasting results based on an alternative out-of-sample (OOS) design. The forecast horizon H ranges from 1 year (4 quarters) to 8 years (32 quarters)

based on underlying overlapping quarterly data. Each return variation measure is smoothed using a backward moving average of two years (8 quarters) using underlying quarterly data. A log-transform is applied to the smoothed diffusive and jump variance measures. The benchmark forecast is the historical average return forecast. Forecasts are computed using a recursive estimation scheme with an initial estimation window of 80 quarters (20 years) for results in Panel A, and an initial estimation window of 120 quarters (30 years) for results in Panel B. The OOS sample for results in Panel A therefore begins in 1953Q2 and ends with the last H -period OOS return that may be computed ending in 2019Q4, whereas for results in Panel B the sample begins in 1963Q2 and ends similarly. The table reports ΔR^2_{OOS} defined as the increase in OOS R^2 relative to the benchmark forecast, expressed as a percentage. The statistic denoted CW p -val gives the p -value associated with the test for superior predictive ability, proposed by [Clark and West \(2006\)](#). The null hypothesis in the Clark-West test is that the benchmark historical average forecast has a mean squared prediction error that is less than or equal to that of the competing model that includes the indicated return variation measure.

Similar to in-sample regression designs, the out-of-sample (OOS) tests indicate little evidence that diffusive volatility predicts stock returns. In contrast, forecasts based on the jump variation measure (JV) typically outperform the benchmark, consistent with earlier in-sample inference results. OOS forecasting gains produced by the jump variation measure relative to the benchmark are economically significant, particularly at longer horizons. For example, ΔR^2 values for jump variation at the 6 year horizon are approximately 10%. The relative jump measure also yields statistically and economically significant OOS forecast improvements relative to the historical average benchmark. The corresponding improvements in forecast accuracy are particularly impressive at longer horizons, with ΔR^2 statistics in the 15–20% range at horizons of 5–8 years.

In order to shed further light on relative OOS forecast performance, [Figure 7](#) plots the cumulative difference in the sum of OOS squared forecast errors for diffusive variance, jump

variance, and the relative jump measure, in the spirit of [Goyal and Welch \(2008\)](#). We show results for forecast horizons of 2 and 6 years. The red dashed line represents the threshold for (cumulatively) outperforming of the benchmark historical average return forecast. The plots for diffusive variation (IV) in both panels are generally flat or downward sloping, indicating no benefit relative to the benchmark. The relative jump measure, in contrast, delivers large gains relative to the historical average during a period from the mid 1960s through the mid 1970s. Although the performance of this predictor is mixed in the remaining part of the OOS period, it again performs well in the early 2000s and delivers significant forecast gains over the full OOS period. The jump variance measure also performs well from the mid 1960s through the early 1970s, and in recent years following the financial crisis. However, it does not perform well in the early part of the OOS period (1950s). This is in part due to the fact that predictive regression slope estimates during this period, based on relatively short estimation samples, are sometimes negative rather than positive as expected. This also explains why the OOS results for the JV are stronger in Panel B using the 30-year initial estimation window.

Broadly, we our tests provide evidence that measures of jump variation capture predictable variation in the equity premium. The predictive power of the jump measures appears particularly strong in the 1960s and early 1970s, a period in which the dynamics of jump variation do not closely match those of diffusive variation (see Panel A of Figure 2).

6. Conclusion

This paper analyzes the long-horizon dynamics of jump activity in stock returns using a novel long history of intradaily price data. Jump volatility and jump incidence rates vary considerably over time. Jump intensity is greatest during the 1940s and 50s, falls to historical lows in the 1970s and 1980s, and partially recovers thereafter. Unscheduled news drives most equity jump activity historically, although the proportion of jumps attributable

to scheduled announcements – especially monetary announcements – increases markedly in recent decades. Finally, we connect variation in jump dynamics with time-variation in the equity premium. Jump variation measures positively forecast aggregate excess stock returns, consistent with theoretical models featuring time-varying jump risk. Importantly, jump and diffusive variation contain different predictive signals for stock returns. The predictive signal of jump variation is strongest at short horizons, in contrast to the long-run predictive signal of the diffusive component of return variation. These results support recently proposed jump-diffusion models that decouple the dynamics of jump risk from those of diffusive return variation.

References

- Andersen, T. G., Bollerslev, T., 1997. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance* 4, 115–158.
- Andersen, T. G., Bollerslev, T., 1998. Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review* pp. 885–905.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., 2007a. Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *The Review of Economics and Statistics* 89, 701–720.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., Labys, P., 2001. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association* 96, 42–55.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., Vega, C., 2003. Micro effects of macro announcements: Real-time price discovery in foreign exchange. *The American Economic Review* 93, 38–62.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., Vega, C., 2007b. Real-time price discovery in global stock, bond and foreign exchange markets. *Journal of International Economics* 73, 251–277.
- Andersen, T. G., Dobrev, D., Schaumburg, E., 2012. Jump-robust volatility estimation using nearest neighbor truncation. *Journal of Econometrics* 169, 75–93.
- Andersen, T. G., Fusari, N., Todorov, V., 2015. The risk premia embedded in index options. *Journal of Financial Economics* 117, 558–584.
- Andersen, T. G., Fusari, N., Todorov, V., 2017. Short-term market risks implied by weekly options. *The Journal of Finance* 72, 1335–1386.

- Bandi, F., Bretaer, L., Tamoni, A., 2018. Long-run economic uncertainty, unpublished Manuscript.
- Bandi, F. M., Perron, B., 2008. Long-run risk-return trade-offs. *Journal of Econometrics* 143, 349–374.
- Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., Shephard, N., 2009. Realized kernels in practice: Trades and quotes. *The Econometrics Journal* 12, C1–C32.
- Barndorff-Nielsen, O. E., Shephard, N., 2002. Estimating quadratic variation using realized variance. *Journal of Applied Econometrics* 17, 457–477.
- Barndorff-Nielsen, O. E., Shephard, N., 2004. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* 2, 1–37.
- Barndorff-Nielsen, O. E., Shephard, N., 2006. Econometrics of testing for jumps in financial economics using bipower variation. *Journal of financial Econometrics* 4, 1–30.
- Bates, D. S., 2000. Post-'87 crash fears in the s&p 500 futures option market. *Journal of Econometrics* 94, 181–238.
- Bollerslev, T., Li, S. Z., Todorov, V., 2016. Roughing up beta: Continuous versus discontinuous betas and the cross section of expected stock returns. *Journal of Financial Economics* 120, 464–490.
- Bollerslev, T., Todorov, V., 2011. Tails, fears, and risk premia. *The Journal of Finance* 66, 2165–2211.
- Bollerslev, T., Todorov, V., Xu, L., 2015. Tail risk premia and return predictability. *Journal of Financial Economics* 118, 113–134.
- Christensen, K., Oomen, R., Renò, R., 2020. The drift burst hypothesis. *Journal of Econometrics* .

- Christensen, K., Oomen, R. C., Podolskij, M., 2014. Fact or friction: Jumps at ultra high frequency. *Journal of Financial Economics* 114, 576–599.
- Christoffersen, P., Jacobs, K., Ornathanalai, C., 2012. Dynamic jump intensities and risk premiums: Evidence from S&P 500 returns and options. *Journal of Financial Economics* 106, 447–472.
- Clark, T. E., West, K. E., 2006. Using out-of-sample mean squared prediction errors to test the martingale difference hypothesis. *Journal of Econometrics* 135, 155–186.
- Cremers, M., Halling, M., Weinbaum, D., 2015. Aggregate jump and volatility risk in the cross-section of stock returns. *The Journal of Finance* 70, 577–614.
- Cutler, D. M., Poterba, J. M., Summers, L. H., 1989. What moves the stock market. *Journal of Portfolio Management* 15, 4–11.
- Davies, R., Tauchen, G., 2018. Data-driven jump detection thresholds for application in jump regressions. *Econometrics* 6, 16.
- Easley, D., De Prado, M. M. L., OHara, M., 2011. The microstructure of the flash crash: flow toxicity, liquidity crashes, and the probability of informed trading. *The Journal of Portfolio Management* 37, 118–128.
- Engle, R. F., Ghysels, E., Sohn, B., 2013. Stock market volatility and macroeconomic fundamentals. *Review of Economics and Statistics* 95, 776–797.
- Engle, R. F., Rangel, J. G., 2008. The spline-garch model for low-frequency volatility and its global macroeconomic causes. *The Review of Financial Studies* 21, 1187–1222.
- Evans, K. P., 2011. Intraday jumps and us macroeconomic news announcements. *Journal of Banking & Finance* 35, 2511–2527.

- French, K. R., Schwert, G. W., Stambaugh, R. F., 1987. Expected stock returns and volatility. *Journal of Financial Economics* 19, 3–29.
- Gabaix, X., 2012. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics* 127, 645–700.
- Gerety, M. S., Mulherin, J. H., 1991. Patterns in intraday stock market volatility, past and present. *Financial Analysts Journal* 47, 71–79.
- Gerety, M. S., Mulherin, J. H., 1992. Trading halts and market activity: An analysis of volume at the open and the close. *Journal of Finance* 47, 1765–1784.
- Gerety, M. S., Mulherin, J. H., 1994. Price formation on stock exchanges: The evolution of trading within the day. *Review of Financial Studies* 7, 609–629.
- Goyal, A., Welch, I., 2008. A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies* 21, 1455–1508.
- Huang, X., Tauchen, G., 2005. The relative contribution of jumps to total price variance. *Journal of Financial Econometrics* 3, 456–499.
- Kelly, B., Jiang, H., 2014. Tail risk and asset prices. *The Review of Financial Studies* 27, 2841–2871.
- Lee, S. S., Mykland, P. A., 2008. Jumps in financial markets: A new nonparametric test and jump dynamics. *Review of Financial Studies* 21, 2535–2563.
- Li, J., Todorov, V., Tauchen, G., 2018. Jump factor models in large cross-sections, working paper.
- Liu, J., Longstaff, F. A., Pan, J., 2003. Dynamic asset allocation with event risk. *The Journal of Finance* 58, 231–259.

- Lucca, D. O., Moench, E., 2015. The pre-FOMC announcement drift. *The Journal of Finance* 70, 329–371.
- Mancini, C., 2001. Disentangling the jumps of the diffusion in a geometric jumping brownian motion. *Giornale dell'Istituto Italiano degli Attuari* 64, 19–47.
- Manela, A., Moreira, A., 2017. News implied volatility and disaster concerns. *Journal of Financial Economics* 123, 137–162.
- Merton, R. C., 1973. An intertemporal capital asset pricing model. *Econometrica* 41, 867–887.
- Merton, R. C., 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.
- Pan, J., 2002. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics* 63, 3–50.
- Paye, B. S., 2012. déjà vol: Predictive regressions for aggregate stock market volatility using macroeconomic variables. *Journal of Financial Economics* 106, 527–546.
- Prokopczuk, M., Wese Simen, C., 2015. What makes the S&P 500 jump? Working paper .
- Santa-Clara, P., Yan, S., 2010. Crashes, volatility, and the equity premium: Lessons from s&p 500 options. *The Review of Economics and Statistics* 92, 435–451.
- Savor, P., Wilson, M., 2013. How much do investors care about macroeconomic risk? evidence from scheduled economic announcements. *Journal of Financial and Quantitative Analysis* 48, 343–375.
- Savor, P., Wilson, M., 2014. Asset pricing: A tale of two days. *Journal of Financial Economics* 113, 171–201.

Schwert, G. W., 1989. Why does stock market volatility change over time? *The Journal of Finance* 44, 1115–1153.

Tauchen, G., Zhou, H., 2011. Realized jumps on financial markets and predicting credit spreads. *Journal of Econometrics* 160, 102–118.

Wachter, J. A., 2013. Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance* 68, 987–1035.

Zhou, C., 2001. The term structure of credit spreads with jump risk. *Journal of Banking & Finance* 25, 2015–2040.

Fig. 1: Quarterly Return Variation Measures: Hourly Versus 5-Minute Data

The figure compares a variety of realized return volatility measures computed using hourly DJCA returns with analogs computed using 5-minute DJIA returns. Panel A contrasts quarterly realized volatility using the hourly and 5-minute returns. Panel B contrasts quarterly diffusive volatility based on the MedRV measure. Panel C compares quarterly jump return volatility, computed as the square root of the quarterly difference between realized volatility and MedRV. Finally, Panel D compares jump volatility that is smoothed by applying a two-year (8 quarter) backward moving average. All series are plotted for the overlapping sample period of 1994–2019.

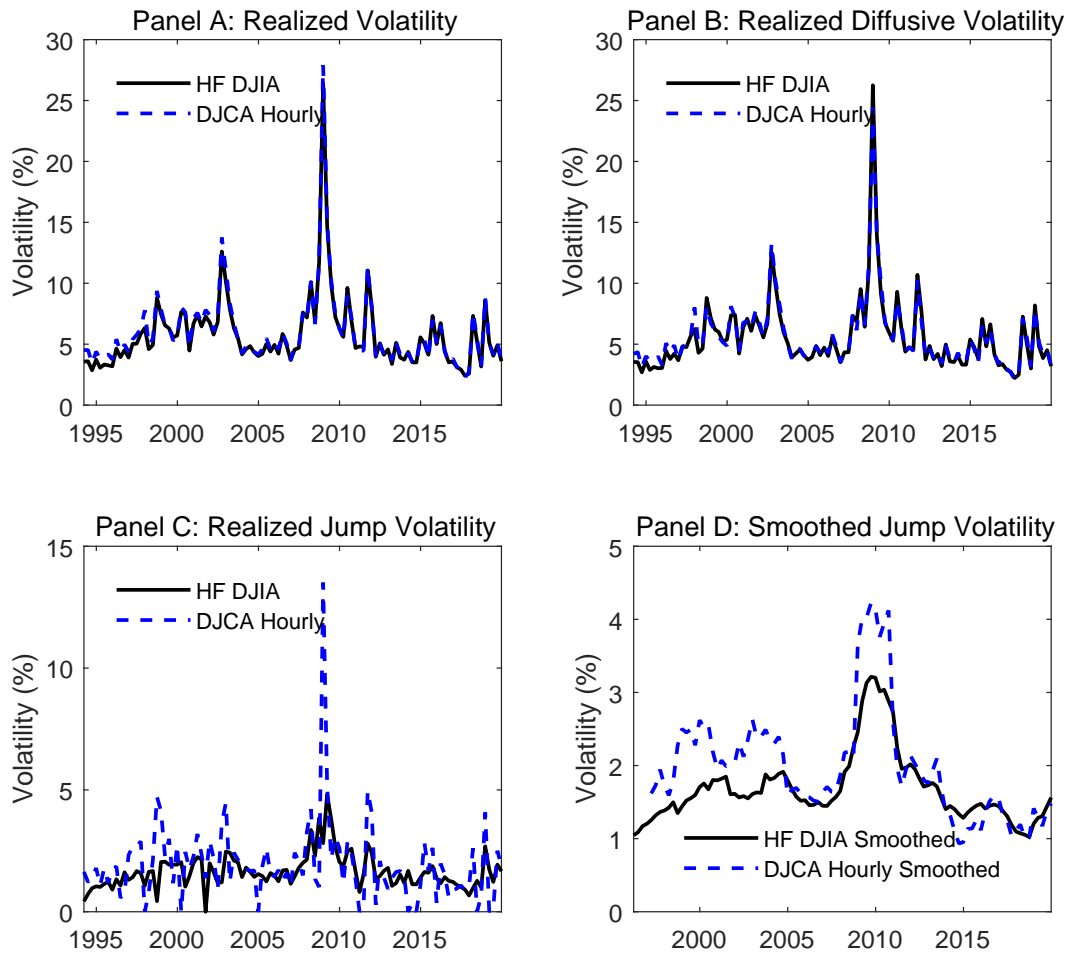


Fig. 2: Return Variation Dynamics Over a Long History

The figure plots measures of jump and diffusive return variation. Panel A plots contrasts jump and diffusive volatility dynamics. The plotted series are computed by applying a two-year (backward) moving average filter to the raw quarterly diffusive volatility (MedRV) series and jump volatility series, respectively. Panel B plots the quarterly relative jump measure RJ , expressed as a percentage. The raw series is smoothed via a two-year backward moving average filter prior to plotting.




Fig. 3: Bandwidth Choice: Hourly Data Versus 5-Minute Benchmark

This figure compares daily spot volatility estimates computed using hourly Dow Jones data with benchmark spot volatility estimates based on 5-minute DJIA returns. All estimates apply the MedRV measure of diffusive volatility. Panel A plots the daily spot volatility estimate using 5-minute returns (black line) against the daily spot volatility estimate using hourly returns (dashed blue line). Panel B provides a similar plot, except the dashed blue line corresponds to the spot volatility estimate computed using $B = 2$, which averages daily MedRV using hourly data over five trading days (current day ± 2 trading days). In order to facilitate detailed comparison, the plot time period is the 4-year period of 2002–2005.




Fig. 4: Jump Activity: Dow Jones Hourly Returns

The figure summarizes jump activity from 1933–2019 based on hourly Dow Jones returns. Jumps are identified based on the jump test described in Section 3 of the paper, with a threshold choice of $\alpha = 4$. Panel A plots the time series of jump returns for each jump identified. Panel B plots the absolute jump returns (left-hand axis) and compares with the time series of truncation levels (right-hand axis). The truncation level is equal to $v_n = \alpha \hat{\sigma}_n$ and is equal to four times the spot volatility estimate for these results.




Fig. 5: Rolling Jump Test Attributes: Dow Jones Hourly Data (1935–2019)

The figure plots rolling measures summarizing jump activity for intraday Dow Jones returns over the period 1935–2019. All rolling estimates are computed over a window of $5,292 = 3 \times 252 \times 7$ intraday returns (approximately three years in calendar time). Panel A plots rolling estimates of jump intensity, defined as the percentage of intraday intervals within the estimation window containing detected jumps. Panel B plots rolling estimates of the mean of the absolute value jump return as a percentage based on those jumps detected within the window. Panel C plots the mean percentage jump return based on those jumps detected within the window. Gray dotted lines depict 95% confidence intervals associated with these rolling estimates. Vertical, dashed, red lines illustrate change points in operating hours for the NYSE (see Table 1).




Fig. 6: Market Frictions Over Time

The figure plots two annual measures of the extent of market frictions associated with the Dow Jones Composite Average series. The first series is the annual average of daily sample return autocorrelation statistics. The second is the annual fraction of total returns that are “zeros”, i.e., a static index value and a corresponding return of exactly zero. The sample period is 1933–2019.




Fig. 7: Out-of-Sample Return Forecasting Performance

The figure illustrates the out-of-sample (OOS) stock return forecasting performance of alternative return variation measures. For each return variation measure, the figure displays the cumulative difference in the sum of squared errors (CDSSE) at each point in time during the OOS evaluation period. The three alternative return forecasting measures are diffusive variance (IV), jump variance (JV) and the relative jump measure (RJ). The return variation measures are smoothed using a two-year backward moving average of the raw quarterly measures. A log transformation is applied to the diffusive and jump variance measures. All forecasts are computed using underlying quarterly data with forecasting regression parameters estimated using an expanding window with an initial estimation sample of 80 quarters (20 years).

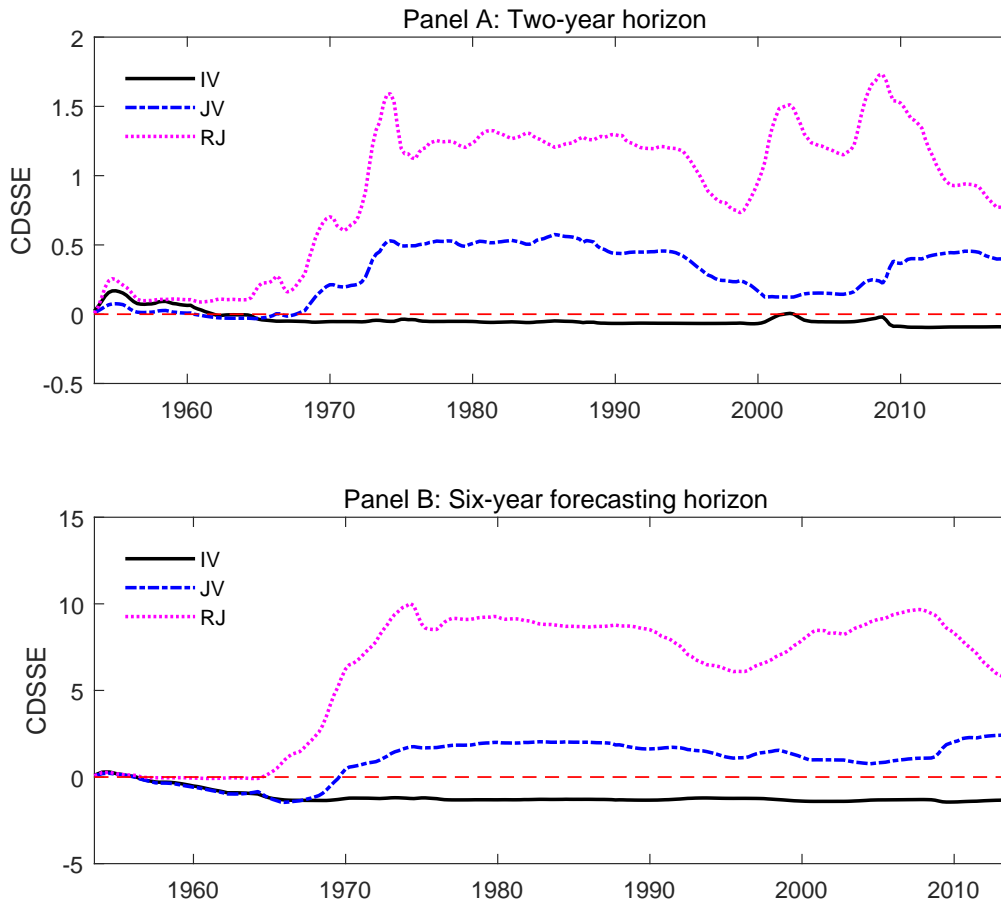


Table 1: Dow Jones Intraday Data: Regime Classification

This table shows the hours over which the New York Stock Exchange operated during our sample period, 5/18/1933–12/29/2017. There are four main “regimes” during which the NYSE changed its open and/or closing times; within each regime, there are occasional subperiods in which other changes took place. The number of observations shown alongside each schedule shows the total number of intraday prices obtained from the *Wall Street Journal*. This table describes only significant, persistent institutional shifts, omitting holidays and unscheduled emergency closings.

<i>Regime</i>	<i>Dates</i>	<i>Schedule</i>	<i>Observations</i>
1	5/18/1933 - 9/26/1952 <i>Exceptions:</i> *	10am - 3pm M-F, 10am-12pm Sat	6 (3 on Saturdays)
	7/2/1945 - 9/7/1945	10am - 3pm M-F	6
	5/20/1946 - 9/27/1946	10am - 3pm M-F	6
	5/26/1947 - 9/26/1947	10am - 3pm M-F	6
	5/24/1948 - 9/24/1948	10am - 3pm M-F	6
	5/23/1949 - 9/23/1949	10am - 3pm M-F	6
	5/29/1950 - 9/29/1950	10am - 3pm M-F	6
	5/28/1951 - 9/28/1951	10am - 3pm M-F	6
	5/26/1952 - 9/26/1952	10am - 3pm M-F	6
2a	9/29/1952 - 11/1/1963	10am - 3:30pm, M-F	6 [†]
2b	11/2/1963 - 9/30/1974 <i>Exceptions:</i>	10am - 3:30pm, M-F	7
	8/8/1967 - 8/18/1967	10am - 2pm, M-F [‡]	5
	1/22/1968 -3/1/1968	10am - 2pm, M-F	5
	1/2/1969 - 7/3/1969	10am - 2pm, M-F	5
	7/7/1969 - 9/26/1969	10am - 2:30pm, M-F	6
	9/29/1969 - 5/1/1970	10am - 3pm, M-F	6
3	10/1/1974 - 9/26/1985	10am - 4pm, M-F	7
4	9/30/1985 - 12/29/2017	9:30am - 4pm, M-F	8 [§]

* The summer NYSE calendar did not include Saturday hours from 1945 through 1952.

[†] In the first subsection of regime 2, observations are given hourly from 10am to 2pm, then the close at 3:30pm, omitting any 3pm observations.

[‡] The “paperwork crisis” forced scheduled closings intermittently throughout 1967–1970.

[§] Through the end of 1993, the data are still collected manually from the *Wall Street Journal*; afterward, hourly index values are reconstructed from TAQ data.

Table 2: **Descriptive statistics for selected years**

This table shows descriptive statistics for intraday returns and daily return variation and correlation measures. *Count* equals the number of intraday intervals during the year. *#/Day* shows the mean number of returns per day. The following two statistics are the mean and median of annual intraday returns. The statistic ‘% Zero’ indicates the percentage of return intervals with zero returns. The final four columns report the daily averages of realized variance (*RV*), bipower variation (*BPV*), the alternative *MedRV* estimator, and the daily sample autocorrelation (*AC*). Panel A covers the first 8 years of our sample (excluding a partial year in 1933). Panel B covers the the 8 year period 1963–1970 and Panel C covers the 8 year period 2003–2010. See the Supplementary Appendix for a longer table reporting statistics for all years from 1933–2019.

Year	Intraday Return Statistics					Daily Averages			
	<i>Count</i>	<i>#/Day</i>	<i>Mean</i>	<i>Median</i>	% Zero	<i>RV</i>	<i>BPV</i>	<i>MedRV</i>	<i>AC</i>
Panel A: 1934–1941									
1934	1244	5.00	-0.010	-0.010	1.287	1.000	0.832	0.750	-0.110
1935	1250	5.00	0.006	0.019	1.760	0.477	0.340	0.318	-0.155
1936	1264	5.00	-0.008	0.007	1.108	0.368	0.299	0.267	-0.141
1937	1250	5.00	-0.029	-0.005	0.960	1.354	1.212	1.157	-0.103
1938	1250	5.00	-0.006	0.000	1.280	1.185	0.955	0.868	-0.109
1939	1245	5.00	-0.012	-0.007	2.008	0.702	0.530	0.487	-0.097
1940	1255	5.00	-0.022	0.000	21.514	0.573	0.485	0.439	-0.118
1941	1250	5.00	-0.015	-0.023	7.760	0.261	0.229	0.198	-0.100
Panel B: 1963–1970									
1963	1292	5.17	0.004	0.004	2.185	0.127	0.094	0.081	-0.087
1964	1516	5.99	0.002	0.003	1.188	0.066	0.060	0.066	-0.063
1965	1511	6.00	-0.001	0.006	0.927	0.100	0.095	0.101	-0.053
1966	1498	5.96	-0.016	-0.014	0.873	0.248	0.220	0.233	-0.006
1967	1486	5.95	0.005	-0.003	1.290	0.154	0.118	0.113	-0.038
1968	1300	5.83	-0.001	0.003	1.373	0.114	0.105	0.105	-0.089
1969	1116	4.56	-0.019	-0.022	1.416	0.150	0.143	0.133	-0.108
1970	1440	5.71	0.005	-0.004	2.178	0.384	0.329	0.299	-0.010
Panel C: 2003–2010									
2003	1757	7.00	0.008	0.011	0.455	0.648	0.681	0.587	-0.136
2004	1757	7.00	0.005	0.007	0.000	0.335	0.328	0.302	-0.159
2005	1750	7.00	-0.004	0.002	0.171	0.395	0.381	0.336	-0.126
2006	1757	7.00	0.003	0.007	0.000	0.379	0.389	0.347	-0.110
2007	1757	7.00	-0.001	0.014	0.000	0.606	0.591	0.514	-0.133
2008	1771	7.00	-0.008	-0.012	0.000	4.193	3.701	3.390	-0.159
2009	1764	7.00	0.006	0.020	0.000	1.685	1.711	1.495	-0.152
2010	1764	7.00	0.004	0.017	1.644	0.728	0.729	0.661	-0.113

Table 3: **Characteristics of Return Variation Components and Comparison Financial Variables**

This table presents summary statistics for quarterly return variation measures and selected financial variables. Panel A presents the sample mean, standard deviation, skewness and kurtosis for each variable as well as the sample autocorrelation. Panel B shows pairwise sample correlation estimates for jump volatility (\sqrt{JV}), diffusive volatility (\sqrt{IV}) and the relative jump measure (RJ). \sqrt{IV} is based on quarterly diffusive variance measured as the average daily MedRV statistic across trading days with the corresponding quarter. \sqrt{JV} is computed as the square root of the difference between the average daily realized variance and diffusive variance. Panel C shows the slope estimate, t -statistic based on Newey-West error with 2 lags, and R^2 -value for a univariate regression of the corresponding variable on the contemporaneous value of the quarterly NBER recession indicator series. DEF is the default spread (%). TS is the term spread (%). KJ is the tail risk measure of Kelly and Jiang (2014). $NVIX$ is the news implied volatility measure of Manela and Moreira (2017). VRP is the variance risk premium. VIX is the CBOE volatility index. The sample size, and sample start and end points are displayed at the bottom of the table. The units of measurement are quarterly percentages for \sqrt{JV} , \sqrt{IV} , DEF , $TERM$, $NVIX$, and VIX .

	\sqrt{JV}	\sqrt{IV}	RJ	DEF	$TERM$	KJ	$NVIX$	VRP	VIX
Panel A: Descriptive Statistics									
Mean	1.97	4.28	21.11	1.04	1.72	0.43	24.27	15.51	35.15
Std. Deviation	1.40	2.66	13.62	0.52	1.28	0.04	4.25	14.73	30.33
Skeness.	2.47	2.52	0.39	1.54	-0.27	-1.09	1.04	-0.74	3.30
Kurtosis	16.81	15.36	2.40	5.61	3.51	3.81	7.45	15.34	18.50
AC(1)	0.14	0.39	0.41	0.92	0.85	0.92	0.81	0.27	0.64
Panel B: Correlations									
\sqrt{JV}	1.00	0.69	0.38	0.33	0.17	0.08	0.26	0.13	0.40
\sqrt{IV}	0.69	1.00	-0.29	0.35	0.25	0.07	0.30	0.12	0.57
RJ	0.38	-0.29	1.00	-0.04	-0.08	0.09	0.05	0.07	-0.06
Panel C: Regression on Quarterly NBER Recession Variable									
β	0.74	1.93	0.70	0.47	0.06	-0.01	2.37	2.96	52.65
t -stat	2.07	2.17	0.28	2.69	0.23	-0.82	1.79	0.33	2.49
R^2 -value	3.00	5.69	0.03	8.88	0.02	0.59	3.49	0.33	24.87

Table 4: **Jump Attribution by Time Period**

We investigate media reactions to large market movements identified as jumps in every year of our sample, and then break the sample into ‘Early’ (1935–1975) and ‘Late’ (1976–2019) periods. Jumps are the top 100 jumps, as measured by jump magnitude, identified under the baseline method with an increased threshold of $\alpha = 5$. This table summarizes the media attribution of those jumps. In Panel A, news events that happen according to a predetermined schedule are categorized as *Scheduled* news. Examples of scheduled news include scheduled economic announcements such as those by the Federal Reserve or the Bureau of Labor Statistics, as well as political events such as elections. Jump attributions of all other news events are categorized as *Unscheduled*. In Panel B, news is categorized according to one of five broad classes. *Economic* news includes both scheduled economic announcements, such as those regarding the Fed funds rate or inflation, as well as media stories regarding economic factors. *Industry/Commodity* news stories are those in which large market swings are attributed to a group of firms, an industry, or commodity prices, and can even be regarding single firms in instances where a bellwether firm changes dramatically enough to impact the industry and, in turn, the market as a whole. *Political* news includes elections as well as other political developments. *War/Conflict* news involves any event regarding US military conflicts or similar news – World War II, the Korean war, the Vietnam war, the Gulf war, and the attacks on September 11, 2001.

Panel A: Scheduled vs. Unscheduled News						
<i>Period</i>	<i>All Scheduled</i>	<i>Scheduled Fed</i>		<i>All Unscheduled</i>	<i>Unscheduled Fed</i>	<i>Total</i>
Early	0 (0.0%)	0 (0.0%)		48 (100.0%)	1 (2.1%)	48
Late	17 (32.7%)	9 (17.3%)		35 (67.3%)	5 (9.6%)	52
Total	17 (17.0%)	9 (9.0%)		83 (83.0%)	6 (6.0%)	100

Panel B: News Categories							
<i>Year</i>	<i>Economic</i>	<i>Fed</i>	<i>Indus./Comm.</i>	<i>Political</i>	<i>War/Conflict</i>	<i>Other</i>	<i>Total</i>
Early	8 (16.7%)	1 (2.1%)	8 (16.7%)	6 (12.5%)	7 (14.6%)	18 (37.5%)	48
Late	16 (30.8%)	14 (26.9%)	10 (19.2%)	2 (3.8%)	0 (0.0%)	10 (19.2%)	52
Total	24 (24.0%)	15 (15.0%)	18 (18.0%)	8 (8.0%)	7 (7.0%)	28 (28.0%)	100

Table 5: **Stock Return Forecasting Power**

The table presents results for predictive regressions for stock returns. The model is

$$RET_{t+1:t+H} = \alpha + \beta \text{VAR}_t + \epsilon_{t+H},$$

in which the dependent variable $RET_{t+1:t+H} \equiv \sum_{h=1}^H R_{t+h} - RF_{t+h}$ consists of the cumulative log excess return on the market portfolio over H periods, where R_{t+h} and RF_{t+h} denote the quarterly log return on the CRSP value-weighted portfolio and the log risk-free rate, respectively. VAR_t denotes a particular (univariate) return variation measure which is specified in the ‘Variable’ column of the table. The three alternative predictors considered are diffusive variance (IV), jump variance (JV), and the relative jump measure (RJ). Each return variation measure is smoothed using a backward moving average of two years (8 quarters) using underlying quarterly data. A log-transform is applied to the smoothed diffusive and jump variance measures. The forecast horizon H ranges from 1 year ($H = 4$ quarters) to 8 years ($H = 32$ quarters) based on underlying overlapping quarterly data. The sample period begins in 1933Q2 with the final observation occurring as the final H -period return culminating in 2019Q4. For each predictor, the table shows slope estimates with corresponding t -statistics and the forecasting regression R^2 -value. Reported t -statistics are based on Newey-West standard errors with $2 \times H - 1$ lags. The total sample observations N for each horizon appears at the bottom of the table.

Horizon (years)	1	2	3	4	5	6	7	8
Panel A: Diffusive Variance (IV)								
$\hat{\beta}$	0.00	-0.01	-0.02	0.00	0.03	0.02	0.03	0.06
t -stat	0.14	-0.31	-0.50	0.03	0.31	0.25	0.33	0.57
R^2	0.02	0.14	0.53	0.00	0.36	0.29	0.50	1.46
Panel B: Jump Variance (JV)								
$\hat{\beta}$	0.03	0.03	0.03	0.09	0.16	0.17	0.20	0.26
t -stat	1.81	1.11	0.75	1.68	2.54	2.78	3.46	4.34
R^2	1.92	1.06	0.84	5.14	11.64	12.26	14.40	21.56
Panel C: Relative Jump Measure (RJ)								
$\hat{\beta}$	0.04	0.07	0.09	0.12	0.16	0.19	0.21	0.25
t -stat	1.70	1.90	1.94	2.26	2.51	2.55	2.67	3.30
R^2	2.38	4.50	6.15	9.06	12.72	15.04	16.75	20.74
N	336	332	328	324	320	316	312	308

Table 6: **Out-of-Sample Predictive Regressions for Stock Returns**

The table shows out-of-sample (OOS) results for predictive regressions. The model is

$$RET_{t+1:t+H} = \alpha + \beta VAR_t + \epsilon_{t+H},$$

in which the dependent variable $RET_{t+1:t+H} \equiv \sum_{h=1}^H R_{t+h} - RF_{t+h}$ is the H -period cumulative log excess return on the market portfolio, where R_{t+h} and RF_{t+h} denote the log return on the CRSP value-weighted portfolio and the log risk-free rate, respectively. VAR_t denotes a particular (univariate) return variation measure which is specified in the ‘Variable’ column of the table. The three alternative predictors considered are diffusive variance (IV), jump variance (JV), and the relative jump measure (RJ). Each return variation measure is smoothed using a backward moving average of two years (8 quarters) using underlying quarterly data. A log-transform is applied to the smoothed diffusive and jump variance measures. The forecast horizon H ranges from 1 year (4 quarters) to 8 years (32 quarters) based on underlying overlapping quarterly data. The benchmark forecast is the historical average return forecast. Forecasts are computed using a recursive estimation scheme with an initial estimation window of 80 quarters (20 years) for results in Panel A, and an initial estimation window of 120 quarters (30 years) for results in Panel B. The OOS sample for results in Panel A begins in 1953Q2 and ends with the last H -period OOS return that may be computed ending in 2019Q4, whereas the sample for results in Panel B begins in 1963Q2 and ends at the same point. The table reports ΔR_{OOS}^2 defined as the increase in OOS R^2 relative to the benchmark forecast, expressed as a percentage. The statistic denoted CW p -val gives the p -value associated with the test for superior predictive ability, proposed by [Clark and West \(2006\)](#). The null hypothesis in the Clark-West test is that the model without VAR_t , i.e., the historical average benchmark forecast has a mean squared prediction error that is less than or equal to that of the model that includes VAR_t .

Variable (X_t)	Statistic	Horizon (Years)							
		1	2	3	4	5	6	7	8
Panel A: Initial 20-year estimation window									
IV	ΔR_{OOS}^2	-0.74	-0.67	-1.70	-3.36	-2.14	-4.19	-5.97	-5.30
	CW p -val	0.58	0.56	0.65	0.89	0.75	0.93	0.98	0.93
JV	ΔR_{OOS}^2	0.90	2.82	2.43	3.27	9.53	7.33	5.13	8.91
	CW p -val	0.10	0.03	0.07	0.07	0.01	0.03	0.05	0.02
RJ	ΔR_{OOS}^2	0.88	5.53	7.09	10.97	15.44	17.43	18.33	22.20
	CW p -val	0.08	0.01	0.01	0.02	0.03	0.03	0.03	0.03
Panel B: Initial 30-year estimation window									
IV	ΔR_{OOS}^2	-1.01	-0.72	-1.05	-1.79	-0.67	-1.37	-1.54	-0.16
	CW p -val	0.90	0.80	0.68	0.86	0.68	0.84	0.88	0.45
JV	ΔR_{OOS}^2	1.52	3.39	3.17	5.00	11.93	11.01	10.25	15.75
	CW p -val	0.06	0.02	0.04	0.04	0.00	0.01	0.01	0.00
RJ	ΔR_{OOS}^2	2.24	5.29	7.52	11.39	16.12	18.68	20.28	25.11
	CW p -val	0.06	0.02	0.02	0.02	0.01	0.03	0.03	0.02