Measuring Tail Risks at High Frequency

Brian Weller*

Duke University

October 25, 2018

Abstract

I exploit information in the cross section of bid-ask spreads to develop a new measure of extreme event risk. Spreads embed tail risk information because liquidity providers require compensation for the possibility of sharp changes in asset values. I show that simple regressions relating spreads and trading volume to factor betas recover this information and deliver high-frequency tail risk estimates for common factors in stock returns. My methodology disentangles financial and aggregate market risks during the 2007--2008 Financial Crisis; quantifies jump risks associated with Federal Open Market Committee announcements; and anticipates an extreme liquidity shock before the 2010 Flash Crash.

JEL: C58, G01, G12, G14, G17

Keywords: Tail Risks, High-Frequency Market Making, Bid-Ask Spreads

*Thanks for many helpful comments to the Editor, Andrew Karolyi, anonymous referees, Torben Andersen, Snehal Banerjee, Eric Budish, Anna Cieslak, Nicolas Crouzet, Ian Dew-Becker, Thierry Foucault, Ravi Jagannathan, Bryan Kelly, Robert Korajczyk, Xu Lu, Konstantin Milbradt, Marina Niessner, Dimitris Papanikolaou, John Shim, Viktor Todorov, Haoxiang Zhu, my discussants, Mikhail Chernov, Dobrislav Dobrev, Arie Gozluklu, Eric Jondeau, Thomas Katzschner, and Allan Timmermann, and seminar participants at Northwestern Kellogg, HEC Paris, University of Lausanne and EPFL, Duke University, the NBER Asset Pricing Meeting, the 4th Industrial Organisation of Securities and Derivatives Markets Conference, the 11th Annual Central Bank Conference on the Microstructure of Financial Markets, the 12th Annual Meeting of the Financial Research Association, and the Paris December International Finance Meeting. I am grateful to the Kellogg Data Center and especially Tomasz Wisniewski for large-scale computing support. Send correspondence to Brian Weller, Duke University, 213 Social Sciences Building, Box 90097, Durham, NC 27708. Telephone: (919) 660-1720. E-mail: brian.weller@duke.edu. Website: https://sites.google.com/site/brianmweller/.

I. Introduction

Fear of extreme events lies at the heart of financial markets. Although crashes are readily identified after the fact, anticipating extreme events remains a pressing challenge. The rarity of extreme events makes forecasting them difficult using standard dynamic techniques on historical data. For this reason, existing approaches to risk assessment instead typically rely on options data to extract market forecasts of extreme events. However, the rarity of liquid, deep out-of-the-money options limits the estimation frequency and potential scope of these procedures. The objective of this paper is to introduce a complementary methodology to overcome these limitations.

In this paper, I use the cross section of bid-ask spreads to develop a new, real-time measure of extreme event risk. By drawing on high-frequency quote data for thousands of U.S. stocks, I improve the resolution of tail-risk estimates from months to minutes and the set of potential factors from those with liquid options to any factors that explain the cross section of realized stock returns. More generally, I demonstrate that the behavior of market intermediaries offers a rich new resource for understanding aggregate economic shocks and potential systemic threats.

Extreme market events take many forms, and the tail risks detected by my approach consist primarily of sharp and sudden factor crashes or jumps. This set of anticipated risks encompasses sharp factor price movements of several basis points, on the order of the median half-spread, to extreme price jumps, realized, for example, during the Black Monday crash of October 19, 1987, or during any of several market crashes at the height of the 2007–2008 Financial Crisis. These extreme market movements jeopardize years of investment returns in minutes or hours, and they are endemic to equity, currency, and commodity markets alike.

The market-making sector provides a natural setting for recovering high-frequency estimates of anticipated risks.[1] Liquidity providers set quotes to balance expected gains from intermediation against potential losses from trading against better-informed market participants. One important source of adverse selection derives from failing to adjust quotes immediately in response to informa-

[1] I use the term "market maker" to encompass all market liquidity providers rather than designated market makers (DMMs) alone. Equity market makers differ from traditional securities dealers in operating competitively in limit-order markets, but the name continues to be used to describe modern liquidity providers. For example, Virtu Financial, one of the world's largest liquidity providers, refers to itself as a market maker even for its non-DMM roles (https://www.virtu.com/market-making).

tion before another party trades against them. Large factor innovations are especially damaging to market makers because such "picking-off" losses may be realized in many securities simultaneously. High-frequency liquidity providers harness extensive market data and advanced algorithms to manage this source of risk, and they continually revise their quotes in response to anticipated sudden price changes. Prevailing quotes thus reflect market-maker risk assessments on a near-instantaneous basis for every exchange-traded security.[2]

My approach extracts tail risk estimates using this quote information for a large cross section of stocks. The technique consists of a two-step regression approach in the style of Fama and MacBeth (1973) regressions for estimating factor prices. The first step estimates risk exposures (betas) using time-series regressions of stock returns on factor realizations. The second step estimates cross-sectional regressions of a liquidity composite on factor betas from the first step. The coefficients from these cross-sectional regressions are tail risk estimates ξ_{kt} for factors $k = 1, \ldots, K$ at date t.

The interpretation of ξ_{kt} derives from a simple model of competitive liquidity provision. All else equal, liquidity providers quote larger spreads and fewer shares for high-beta securities to offset greater expected losses from being adversely selected on extreme factor moves. The difference in spreads across securities with different factor exposures increases with the anticipated size and arrival rate of factor jumps, but only for jumps large enough to cause picking-off losses by pushing asset values outside the quoted best bid and offer. At the same time, uninformed volume also varies across stocks, and spreads can be high because this volume is low or because risk is high. Combining these features, it follows that the cross-sectional slope of Vh/d—a liquidity composite of volume times spreads over depth—with respect to factor loadings relates to expected jump risk for each factor at date t.[3] I formalize this intuition in a model of high-frequency market making based on Budish, Cramton, and Shim (2015).

To facilitate exposition, I assume in this example that picking-off risk is the sole source of the

[2]Partly as a result of continual market-maker quote revisions, the ratio of order volume to trading volume typically exceeds 30 for U.S. stocks and 500 for U.S. exchange-traded products (Market Information Data Analytics System, http://www.sec.gov/marketstructure/datavis/ma_overview.html). Exchange-traded products are defined as CRSP securities with share code 73 and primarily consist of exchange-traded funds.

[3]Spreads and volume are both essential ingredients of the liquidity composite. The Online Appendix considers spreads and volume as separate dependent variables and shows that cross-sectional slopes are often negative, i.e., higher beta stocks have lower spreads.

bid-ask spread. Inventory risk or non-jump adverse selection risk violates this assumption. For individual assets, including factor-mimicking ETFs, separating out the component of the bid-ask spread associated with tail risk at each date is difficult or impossible. By contrast, cross-sectional tail risk estimates are only contaminated by these other sources of the bid-ask spread if they line up with factor betas. I verify empirically that such contamination does not occur. To rule out contamination by inventory risk, I rerun my primary analysis using the adverse selection component of the bid-ask spread in place of the effective spread under the identifying assumption that this component of the spread is unrelated to inventory risk. To rule out contamination by non-jump adverse selection, I include controls for a commonly used measure of "slow" adverse selection, the probability of informed trading, or PIN (Easley and O'Hara (1992); Easley, Kiefer, O'Hara, and Paperman (1996)). Implied market tail risks are nearly identical to the baseline specification in both analyses.[4]

Importantly, my procedure recovers an *expected* cost of extreme realizations for each time interval because bid-ask spreads reflect liquidity providers' forward-looking information on tail risks. Hence slopes estimated using the liquidity composite represent conditional factor risks directly—they do not require time-series averaging as do risk premia estimates in standard Fama-MacBeth regressions with realized returns as the dependent variable in the second stage. The recovered high-frequency market tail risk series aligns well with measures of anticipated and realized jump tails. The correlation with weekly left jump tail estimates from options data (Bollerslev and Todorov (2014)) exceeds 75%, and a one standard deviation increase in the jump tail measure is associated with 4.4 more realized medium-scale jumps per hour (t-statistic of 7.4). Importantly, my measure predicts near-term tail realizations controlling for other tail risk and volatility measures, including the volume-synchronized probability of informed trading (VPIN) measure of Easley, López de Prado, and O'Hara (2012).

The tail-risk measure serves as a real-time barometer of market risks across diverse and challenging economic environments. As a leading application, I analyze the May 6, 2010 Flash Crash as a prototypical large and plausibly unexpected systematic jump. Existing tail risk estimation techniques do not have sufficient resolution to anticipate the Flash Crash or to reliably detect changes

[4] The Online Appendix also confirms that results do not change when evaluating the model on subsets of stocks less subject to these risks.

in perceived tail risk after the event. Consequently, regulators possessed little ability to intervene until long after the event concluded because they lacked reliable tools for evaluating and responding to near-term catastrophe risks.

My measure helps to fill this gap by providing intraday assessments of tail risks. My tail risk measure is a natural leading indicator for liquidity crashes because it draws on liquidity providers' perceptions of extreme event risk. For the Flash Crash, market tail risk increases by several standard deviations (relative to its value two days before) an hour before the Flash Crash begins, by 20 standard deviations in the quarter hour before the Flash Crash begins, and by a remarkable 96 standard deviations at the height of the event. By contrast, realized volatility and idiosyncratic tail risk (corresponding with level changes in spreads) increase sharply only as the crash develops, suggesting that market makers correctly anticipated a liquidity crisis in the market factor and only later adjusted spreads to accommodate liquidity spillovers uncorrelated with the market-factor liquidity shock.

Notably, the extreme tail risks detected around the Flash Crash are not the result of a high false-positive rate. The only other occasions that register more than a 10-standard deviation increase in tail risk also correspond with extreme market distress: unprecedented market interventions in the wake of the collapse of Lehman Brothers (September 18, 2008); global stock market crashes in Asia, Europe, and the U.S., including the worst-ever weekly drop in the S&P 500 (October 10, 2008); and a large negative employment shock associated with the most new jobless claims since September 11, 2001 (November 13, 2008).

The Flash Crash also serves as an example of the dual uses of my measure. In addition to utilizing it as a forward-looking indicator, I apply the tail risk measure retrospectively to assess whether market makers register persistently elevated crash fears after the event. Both market and idiosyncratic anticipated jump risks quickly revert to pre-Crash levels and are statistically indistinguishable from the pre-Crash period by the following week. Despite regulatory agency concerns that the Flash Crash might undermine market confidence, the meltdown has only a short-lived impact on market perceptions of the risk of extreme market liquidity events.

I next exploit the methodology's new intraday resolution to document the evolution of tail risks around major scheduled macroeconomic news. I show that anticipated jumps vary throughout Federal Open Market Committee (FOMC) announcement days in regular patterns of decreased tail risk (relative to non-announcement days) prior to the announcement, heightened tail risk in the quarter hours before and containing the announcement, and slightly elevated tail risk after the announcement. This finding suggests that the pre-FOMC announcement drift documented by Lucca and Moench (2015) and the anomalous performance of the CAPM documented by Savor and Wilson (2013, 2014) cannot be rationalized by unobserved market jump risk without concurrent time variation in risk premia.

Finally, I demonstrate that the methodology separately identifies tail risks in a multifactor setting, even when candidate factors are very highly correlated. For this purpose, I study the coevolution of aggregate market and financial sector risks during the 2007–2008 Financial Crisis.[5] Despite the Financial Select Sector SPDR ETF (XLF) having an cross-year average daily correlation of 89% with the SPY over this period, cross-sectional differences in risk exposures are nonetheless large enough to recover precise estimates for anticipated shocks specific to the financial sector. The most extreme changes in the time series of financial sector tail risks often differ from those of the market jump series and correspond to major uncertainty innovations specific to financial firms, e.g., bank nationalization rumors and congressional votes on Fannie Mae and Freddie Mac rescue packages. The methodology thus offers a unique and useful tool for understanding the 2007–2008 Financial Crisis and assessing ongoing sectoral risks. This application provides a novel link between Brunnermeier and Pedersen (2009)'s concepts of market and funding liquidity by learning about fears of extreme disruptions to banks and the financial sector from the behavior of market intermediaries.

II. Related Literature

A. Tail-Risk Measurement

The primary objective of this study is to develop a forward-looking measure of instantaneous tail risk for a variety of return factors. Like Kelly and Jiang (2014), this paper takes a cross-

[5] I verify that the methodology also applies for traditional asset pricing factors such as "value" (HML) in the Online Appendix.

sectional approach to obtain conditional tail risk estimates. Kelly and Jiang (2014) show that the aggregate market tail inherits individual asset tail dynamics if asset return tails follow a power law. If tail realizations are not too infrequent, this cross-sectional approach can detect physical market-factor tail shapes with short panels on the order of one month. My approach differs in two key respects. First, my estimation strategy relies on bid-ask spreads rather than on tail return realizations. Because every spread is informative at all times rather than only in the "rare event" states associated with extreme returns, I significantly increase the conditioning frequency at which tail risk estimates can be constructed. Second, my measure recovers tail expectations, which jointly summarize factors' ex ante tail position and shape, rather than the realized tail shape beyond a time-varying threshold value. As an example of this distinction, Kelly and Jiang (2014)'s time-varying tail threshold increases sharply during the 2007–2008 Financial Crisis, and the implied tail shape looks no more extreme than during the preceding years as a result.

The two prevailing alternatives for tail risk measurement take advantage of options panels or of high-frequency time series for individual securities. The most closely related work in this literature is Bollerslev and Todorov (2014), which uses a cross section of S&P 500 index options to recover *time-varying* jump tails. In so doing, the authors exploit the fact that differential exposure to jump risk is the key source of variation in prices of close-to-maturity deep out-of-the-money options. Likewise, this paper makes use of the insight that differential exposures to jump risk drive variation in the size of (volume-adjusted) bid-ask spreads—themselves interpretable as prices of very short-dated options (Copeland and Galai (1983))—across stocks with different factor betas.

Bakshi, Kapadia, and Madan (2003) consider skewness and kurtosis for systematic and idiosyncratic risks as implied by differential pricing of individual equity options. Bollerslev, Tauchen, and Zhou (2009) estimate the variance risk premium in a model-free setup. Backus, Chernov, and Martin (2011) recover the distribution of implied consumption disasters from options data. Options provide richer moneyness cross sections than bid-ask spreads, and for this reason, such options-based analyses can identify a *distribution* of potential risks whereas my approach cannot. Bollerslev and Todorov (2011a,b) supplement options analyses with high-frequency data and extreme value in-fill arguments to estimate jump tails for the aggregate market.

My methodology complements these approaches. This paper adds the ability to estimate tail risks (1) in the very near term, (2) for a broad set of factors, (3) with high-frequency conditioning, and (4) under alternative sets of assumptions. Options-based approaches have difficulty assessing near-term risks because option maturities are long relative to intraday or daily events,[6] and many options on individual names are too illiquid to be used for recovering non-market factor information. Likewise, combining realized jumps with extreme value theory can recover only very slow-moving variation in jump tails, and it is not yet applicable to candidate factors not directly traded in liquid factor-mimicking securities (e.g., size, value, and momentum). Conversely, my approach is limited in not being able to describe the full distribution of potential jump events or to gauge the persistence of negative shocks in a forward-looking way.

B. Market Microstructure

The key relation between spreads and tail risks emerges from Budish, Cramton, and Shim (2015)'s model of high-frequency market making in the presence of picking-off risk. I augment their model by imposing a factor structure on the jump process and by considering the resulting cross section of spreads across multiple assets. Forerunners in developing this source of risk include Copeland and Galai (1983), Harris and Schultz (1997), and Foucault, Röell, and Sandås (2003), among others. Indeed, the equilibrium condition of Budish, Cramton, and Shim (2015) and this study can also be motivated using the quotes-as-options framework of Copeland and Galai (1983). Viewed from this perspective, bid-ask spreads naturally extend short-dated options to the near-instantaneous expiration horizons that are especially well-suited for isolating jump risks (Bollerslev and Todorov (2011b)).

Many other works investigate the information content of the limit order book. Of this set, Foucault, Moinas, and Theissen (2007) is closest to this paper in showing that limit order books contain volatility information in addition to directional information for returns, albeit for individual assets.

This paper also shares the spirit of Nagel (2012) in relating returns to intermediation to forward-

[6]Even Carr and Wu (2003) filter out options with time to maturity less than one week, and their important study explicitly focuses on option price dynamics as time to maturity goes to zero.

looking *market* volatility. Specifically, Nagel (2012) shows that short-term reversal returns are highly correlated with the VIX, and he interprets this relation as evidence that financially constrained intermediaries are less able to provide liquidity in bad economic times. I derive a similar relationship between intermediary behavior and the VIX with picking-off risk as the central driver. When anticipated jump risks—a component of the VIX—are greater, the required return to intermediation is higher and the cross-sectional slope of bid-ask spreads with respect to factor exposures is steeper. By contrast with time-varying tail risk, Nagel (2012)'s constrained-intermediary mechanism has less bite at high frequencies or for non-market factors considered in this paper.

A key contribution of this work is to extend this intuition to a broad set of factor risks and show that bid-ask spreads embed rich information about the underlying structure of asset returns. In this sense, my paper also relates to the broad literature on common factors in liquidity and trading volume (e.g., Chordia, Roll, and Subrahmanyam (2000), Lo and Wang (2000), Hasbrouck and Seppi (2001), and Korajczyk and Sadka (2008)). Prior work such as Lo and Wang (2000), Hasbrouck and Seppi (2001), and Cremers and Mei (2007) also recover common factors in volume and liquidity measures, but these factors differ from realized or excess return factors typically studied in the asset pricing literature.

III. Spreads and Asset-Pricing Risks

A. Picking-Off Risks and Return Tails

To motivate my measure of factor tail risk, I build on the models of Copeland and Galai (1983) and Budish, Cramton, and Shim (2015) (BCS). Liquidity consumers or "fundamental traders" (FT) arrive at rate λ_{FT}, and in the aggregate, they pay liquidity providers the half-spread h multiplied by their arrival rate per unit time. Information events arrive at rate λ_{jump}, and they shift the fundamental value of a security by a stochastic value, J. When prices jump, fast arbitrageurs trade at the old or "stale" quoted price if the size of the jump exceeds the half bid-ask spread, costing the liquidity provider $J - h$. Liquidity providers' expected costs per unit time due to "fast" adverse selection are thus the arrival rate of information events λ_{jump} multiplied by the expected price jump conditional on the jump size exceeding the half spread h. Competitive intermediation drives expected profits per unit time to zero, delivering the equilibrium condition of Budish, Cramton, and Shim (2015):

$$\underbrace{\lambda_{FT} \times h}_{E[\text{benefit/time}]} = \underbrace{\lambda_{jump} \times \Pr(J > h) \times E[J - h | J > h]}_{E[\text{cost/time}]}. \tag{1}$$

Equation (1) relates anticipated price movements J to the half-spread h at each date t (both in percent). Higher arrival rates of liquidity consumers drive spreads toward zero, whereas faster information arrivals or larger jumps conditional on information arrivals increase h. Other determinants of the spread such as inventory costs and non-jump adverse selection can be added on the right of Equation (1), and I consider in detail the potential effects of these omitted terms in later sections.

I modify the BCS setup in two ways. First, I allow liquidity demand and liquidity supply to exceed one share or futures contract. Denote the (stochastic) quantity demanded as q and the quoted depth as d. For simplicity, I assume that queue positions are random from the perspective of liquidity providers so that all units of liquidity offered have the same expected revenues and costs per unit time.[7] With this assumption, the equilibrium condition of Equation (1) generalizes to

$$\begin{aligned} \lambda_{FT} \times h \times \overbrace{\{E[q | q \le d] \times \Pr(q \le d) + d \times \Pr(q > d)\}}^{\equiv q^*} = \\ \lambda_{jump} \times \Pr(J > h) \times E[J - h | J > h] \times d. \end{aligned} \tag{2}$$

Liquidity providers at the best bid or offer only intermediate for liquidity demands up to $q = d$. Beyond that point, larger liquidity demands instead convert into resting limit orders or "walk the book" to consume liquidity at higher prices (for which a different equilibrium condition applies). q^* is then the expected liquidity consumed by fundamental traders at the best bid or offer given a trade and quoted depth d, and expected intermediation revenues scale with $q^* \le d$. By contrast, the expected cost of quoting more depth increases with d because well-capitalized arbitrageurs take the entire offered depth when given the (arbitrage) opportunity.

[7] In high-frequency settings, continual churn in the limit order book makes position order difficult to track (Yueshen (2014)).

Second, I relate market maker liquidity provision to the factor structure in jump returns. Factor structures in jumps have strong empirical support and are the basis for an active literature on jump regressions (e.g., Todorov and Bollerslev (2010), Li, Todorov, and Tauchen (2017), and Bollerslev, Li, and Todorov (2016)). I decompose factor and idiosyncratic jumps in an asset's fundamental value as

$$r^d = \sum_k \beta_k r_k^d + \tilde{r}^d, \tag{3}$$

for a set of return factors k and idiosyncratic jump return $\tilde{r}^d$. I make three simplifying assumptions on the jump processes to facilitate taking the adapted model to the data:

Assumption 1. *Jump arrivals are independent both among factors and between factors and idiosyncratic discontinuous returns.*

Assumption 2. *Idiosyncratic jumps are distributed i.i.d. across assets.*

Assumption 3. *The distribution of jumps for each factor is symmetric.*

These assumptions streamline the construction of the tail risk measure but are otherwise inessential. Assumption 1 excludes co-jumps and more complex jump dependencies among factors. Whether excluding co-jumps is reasonable depends on whether the considered return factors are plausibly orthogonal to one another. Relaxing this assumption to allow for co-jumps is readily accommodated by adding cross terms to each cross-sectional estimation, as described in Appendix A. Assumption 2 excludes heterogeneity in the rate of idiosyncratic information arrival among assets. This assumption, too, may be relaxed (as in the Online Appendix), for example by estimating security loadings on the common factor in idiosyncratic volatility—"CIV" of Herskovic, Kelly, Lustig, and Nieuwerburgh (2016)—under the assumption that such factor loadings extend also to idiosyncratic jumps. In empirical work, I assume symmetry of jumps (assumption 3) because the distribution of *realized jumps* for individual stocks and the SPY and XLF ETFs is very close to symmetric.[8] In the Online Appendix, I estimate up- and down-jump risks separately and confirm that *jump risks* also

[8] Jumps detected using Lee and Mykland (2008) and Bollerslev, Todorov, and Li (2013) methodologies share this symmetry property, and Table III of Bollerslev and Todorov (2011b) also finds a one-factor structure for the objective probabilities of left- and right-jump tails.

are symmetric in normal times, but that symmetry between up- and down-jump risk deteriorates around the most extreme events such as the 2010 Flash Crash.

Under these assumptions, I combine Equations (2) and (3) to obtain

$$\underbrace{(\lambda_{FT} q^* + \lambda_{jump} d)}_{\equiv \bar{V}} \frac{h}{d} = \sum_k \underbrace{\lambda_k E\left[r_k^d | r_k^d > \bar{h}_k\right]}_{\equiv \xi_k} |\beta_k| + \underbrace{\tilde{\lambda} E\left[\tilde{r}^d | \tilde{r}^d > \bar{h}\right]}_{\equiv \tilde{\xi}}. \tag{4}$$

Appendix A derives this expression. The left-hand side of Equation (4) represents a liquidity composite of normalized half-spreads h/d and expected volume, the trade-size weighted arrival rate of both trader types $\bar{V} = \lambda_{FT} q^* + \lambda_{jump} d$. The coefficient on each β_k, $\lambda_k E\left[r_k | r_k > \bar{h}_k\right]$ or ξ_k, represents the tail risk for factor k. Faster event arrival rates or more damaging potential events must be compensated in equilibrium by higher anticipated trading revenues $\bar{V} h$.

Equation (4) holds for each asset in isolation. It also holds for individual securities as part of a multi-asset liquidity provision strategy because picking-off risks are additive and do not interact across securities. To see why, consider two securities with the same factor exposures and idiosyncratic tail risk. Losses to the market maker in the first security are mirrored by losses in the second security, and sufficiently large jump events damage the market maker in each market independently. Now suppose that the risk loadings of the second security were flipped, that is, $\beta_{1k} = -\beta_{2k}$ for all k. Holding depth equal, the same events that trigger arbitrageur buys (sales) in the first security also trigger arbitrageur sales (buys) in the second security. Despite offsetting factor exposures, potential stale-quote costs double. Intermediation costs for multiple securities are the sum of costs for individual securities (setting aside fixed costs).[9] Because cross-subsidies cannot occur with free entry into liquidity provision, Equation (4) must hold for each asset independently.

B. Cross-Sectional Recovery of Tail Risks

Estimating tail risks from Equation (4) takes place in two stages. First, the econometrician computes betas with respect to candidate realized return factors. I estimate backward-looking,

[9] Equivalently we can view liquidity provision in a single asset as being short a strangle with strike prices $+h$ and $-h$. Liquidity provision in multiple assets entails holding a portfolio of short strangles in different assets. These option payoffs are additive, and hence, the costs of liquidity provision also add.

rolling annual betas using daily returns r_{it} on candidate factor realizations f_{kt} for each stock in the filtered sample i:

$$r_{it} = \alpha_i + \sum_k \beta_{ik}^{(t)} f_{kt} + \epsilon_{it},\ \forall i. \tag{5}$$

Armed with these betas, I estimate (symmetric) tail risks cross-sectionally across stock-level observations via least absolute deviations regression:

$$\left(\frac{Vh}{d}\right)_{it} = \tilde{\xi}_t + \sum_k \xi_{tk} \left|\beta_{ik}\right| + \delta_{it}. \tag{6}$$

For much of the ensuing analysis, all variables (and products of variables, where appropriate) are hourly averages over the respective time interval. d is the bid and offer depth summed across exchanges in shares. V is realized volume divided by two to account for half of volume being buys or sells, on average. I treat realized volume as an unbiased, but noisy proxy for the volume expected by liquidity providers. h is the effective half-spread. δ_{it} is a stock-specific error term for date t. ξ_{tk} represents the average anticipated jump risk over the interval for factor k. The time fixed effect $\tilde{\xi}_t$ controls for common movements in asset-level tail risk not associated with the market factor or other return factors. In applications I assume that factor loadings estimated from daily data over the prior year apply to high-frequency data for the subsequent trading day; to the extent that this assumption is not met, ξs will be estimated with error, and I am less likely to find significant relations between tail realizations and my tail-risk measure in the subsequent analyses.

Equations (5) and (6) resemble Fama-MacBeth regressions for determining prices of factor risk. They differ in that all ξ_t estimates are of independent interest rather than only inputs into a single time-series average value. Equation (6) recovers market maker expectations directly, whereas the second stage of traditional Fama-MacBeth regressions recovers only very noisy realizations around a factor's date-t expected value. For this reason the precision of tail-risk estimates depends on the size of the cross-section rather than the length of the time series.

Equation (6) also clarifies the conditions under which we can interpret ξ_{tk} as the tail risk at time t. From the perspective of omitted-variable bias, a variable correlated with the liquidity composite and factor betas has the potential to contaminate ξ_{tk}. Candidate omitted variables can take the form of traditional microstructure sources of the bid-ask spread, such as inventory risk and non-jump adverse selection, as well as omitted high-frequency return factors. As in standard asset-pricing applications, the recovered coefficients ξ are a price of exposure to the projection of the full set of risks on the factors specified in the time-series regression. I consider and rule out contamination by inventory risk and non-jump adverse selection in Section V.B, and I speak to other potential omitted factors in Section V.C.

For estimation, I use median regression rather than OLS in the second-stage regressions because occasional the liquidity composite Vh/d occasionally takes extreme values caused by data recording errors and idiosyncratic departures from the model. OLS regression places too much weight on fitting this small number of influential points with extreme volumes or spreads, whereas median regression is more robust to such outliers. A drawback to this robust-regression approach is that median regression complicates the problem of using estimated rather than known betas from the first-stage regressions. Direct corrections to standard errors (as in Shanken (1992)) are unavailable for this hybrid approach, and the discontinuous moment conditions associated with median regression make GMM numerically intractable.

For this reason, I compute standard errors via pair bootstrap to resolve the generated-regressors problem. The bootstrap enables straightforward construction of standard errors in settings with complicated dependence across regression stages and calendar time, including quantile-regression settings (see, e.g., Section 3.9 of Koenker (2005)). This procedure first uses the full sample of stocks and 252 rolling trading days to obtain point estimates for ξ_{tk}. Then, for each replication $r = 1, \ldots, R$, I construct resampled coefficient estimates $\xi_{tk}^{(r)}$ as follows:

1. For each stock $i = 1, \ldots, N$, draw with replacement 12 strings of 21 trading days during the 252 trading days preceding date t;
2. Estimate $\beta_{ik}^{(r)}$ for all stocks i using Equation (5);
3. For each time interval on date t, draw N stocks with replacement from the set of N stocks;
4. Estimate $\xi_{tk}^{(r)}$ using Equation (6) on the N randomly chosen stocks.

Confidence intervals for ξ_{tk} are constructed directly from quantiles of $\xi_{tk}^{(1)}, \ldots, \xi_{tk}^{(R)}$ for $R = 1000$

resamples. I use strings of 21 trading days as bootstrap blocks to account for potential serial correlation in errors in Equation (5). The bootstrap hence simultaneously accounts for uncertainty in betas by resampling over time and in tail risk estimates by resampling over stocks and their generated betas.

IV. Data

The primary data sources for this study are the Center for Research in Security Prices (CRSP) U.S. Stock Database and the New York Stock Exchange Trade and Quote (TAQ) data. CRSP provides security attribute data (e.g., share codes), unique ticker-entity mappings, and daily shares outstanding for each security. The TAQ data aggregate orders from all Consolidated Tape Association exchanges and are timestamped to the second. Traded volume over each interval is directly observed. I follow Holden and Jacobsen (2014) to obtain cleaned effective spreads and market depths from the underlying TAQ data. I add average liquidity rebates to effective half-spreads to obtain the gross-of-fees benefit of liquidity provision that accrues to market makers. I assume that rebates are roughly constant across stocks (i.e., that Tape A vs. B vs. C differences are small) and equal to 22 cents per 100 shares. This average rebate size is found in the present-day NYSE price list for the most active liquidity providers (Tier 1; https://www.nyse.com/markets/nyse/trading-info), in Table 1 of the maker-taker analysis of Foucault, Kadan, and Kandel (2013) (who in turn reference a 2009 industry publication), as well as in a recent comprehensive study of maker-taker fees as the average value for active liquidity providers from January 2008 through December 2010 (Cardella, Hao, and Kalcheva (2017)).[10] In addition to CRSP and TAQ, I obtain historical intraday Chicago Board Options Exchange Volatility Index (VIX) data from Pi Trading.

The main data sample consists of all common stocks (CRSP share code = 10 or 11) in the TAQ database from January 2005 to December 2013. Although the TAQ database starts in 1993, three features of this early data complicate analysis. First, the assumptions of continually updated spreads

[10] Reasonable alternatives for the level of the rebate have minimal effect on results. For example, reducing the rebate adjustment from \$0.22 per hundred shares to \$0.00 maintains the same shape in the time series of jumps while shifting the recovered idiosyncratic tail risk down slightly. Because rebates are not set as a function of risk attributes, rebate measurement errors should be uncorrelated with asset betas, and in turn, rebate mismeasurement should have little effect on factor tail risk estimates.

and minimal order processing costs are not plausible until algorithmic-trading and connectivity innovations in the mid-2000s. Second, the key zero-expected profit condition of Equation (1) likely does not apply to specialist or collusive market making characterizing earlier years (Christie and Schultz (1994)). By contrast, free entry is a critical dimension by which algorithmic market making differs from specialist market making (Hendershott, Jones, and Menkveld (2011)). Free entry not only makes the break-even condition more likely to hold through competition, but it also encourages the collection of new signals by liquidity providers. More sophisticated market makers sharpen the adverse selection faced by less-informed intermediaries because the latter group becomes the marginal liquidity providers exactly when order flow is most unfavorable (Han, Khapko, and Kyle (2014)). Third, the influx of opportunistic liquidity providers in the 2000s likewise alleviates the collective capital constraints faced by market intermediaries. Consequently, inventory risk is less likely to be a key factor in determining liquidity prices than for the NYSE specialists studied by Comerton-Forde, Hendershott, Jones, Moulton, and Seasholes (2010). Notwithstanding these concerns, I also consider 2004 in summary plots to illustrate the potential breakdowns of my measure for pre-HFT liquidity provision.

I restrict the sample to exclude the 15 minutes after market open and before market close. These periods are characterized by unusual trader composition and informational events, such as elevated informed trading activity at market open in response to overnight events. For much of my analysis, I split the remainder of the trading day into six hourly bins for 9:45–10:45am through 2:45–3:45pm. The filtered sample consists of roughly 2,800 stocks for each hour of each trading day from 2004 to 2013. Additional data cleaning and filtering details are provided in Appendix B.

V. Results

My empirical analysis proceeds in two parts. In the first part, I recover hourly tail risk estimates for a one-factor market model and compare these estimates to market tail realizations and to alternative near-term forecasts such as the VIX. In the second part (Section VI), I apply the tail risk extraction methodology to verify the performance of the jump tail measure for the 2010 Flash Crash, major macroeconomic news events, and the 2007–2008 Financial Crisis. This last example

Figure I: **Hourly Market Tail Risks, 2004–2013**

This figure plots rolling 10-day means of hourly cross-sectional slope estimates ξ_{tm} (solid blue line) for Equation (6) for each trading date in 2004–2013. Dotted bands depict the corresponding 95% confidence intervals estimated by pairs bootstrap. Realized volatility (dashed red series) is estimated using minutely squared returns on the SPY and scaled to the hourly frequency. The VIX is plotted using the right axis and 'x'-markers for comparison. NBER recession dates are shaded gray.

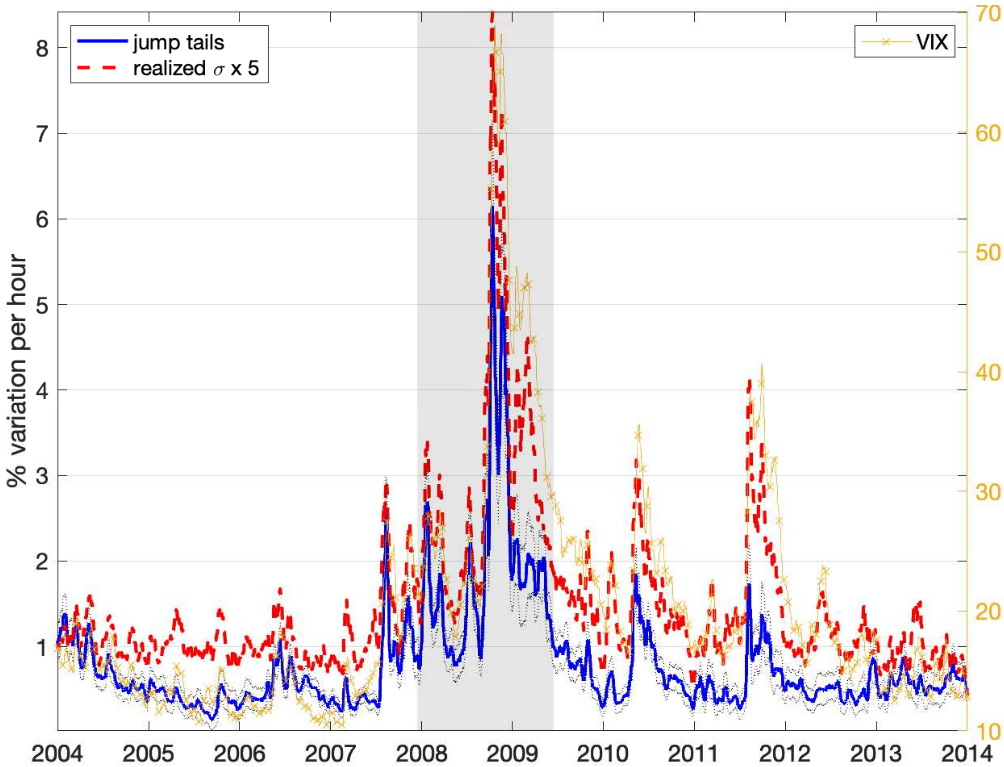

demonstrates the power of the two-stage procedure for separately identifying factor tail risks in a multifactor setting.

Figures I and II plot recovered market tail risks by hour over the 2004 to 2013 sample period. Figure I plots 10-day rolling means of ξ_{MKT} to emphasize low-frequency variation in tail risk. Reassuringly, spread-implied jump risk spikes relative to realized volatility during the financial crisis and global recession, a time of heightened extreme-event risk. Confidence intervals are tight enough to ensure that market tail risk is statistically distinguishable from zero throughout the sample. This feature reflects the value of exploiting large cross sections for identifying conditional risks.

Figure II illustrates my measure's ability to capture market news in real time. I mark the fifteen largest changes in implied market risks over the preceding 24 hours with black Xs, separating peaks

Figure II: **Largest Increases in Tail Risk, 2004–2013**

This figure plots hourly estimated market tail risks in a one-factor market model. The fifteen largest increases in tail risks within a one-month window are overlaid with a black 'X.' Changes are measured as the tail risk at date t and hour h less the tail risk at date $t-1$ and hour h. The table below offers a brief description of coincident events on tail-risk news days. Standardized values divide by the full time-series standard deviation of changes. Bolded events coincide with the most extreme increases in the VIX within 24 hours. Implied market tails on October 10, 2008 are truncated for visual clarity.

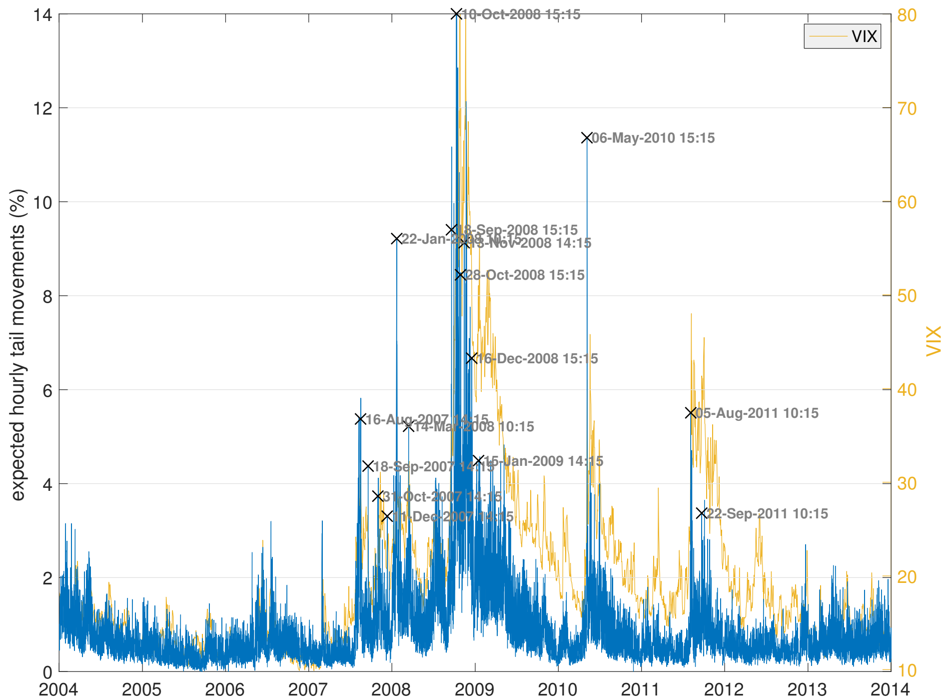

Date	Value	Event
16-Aug-07 14:15	7.24	**Fed approves changes to its primary credit discount window facility**
18-Sep-07 14:15	6.65	FOMC lowers Fed funds target 50bps to 4.75%
31-Oct-07 14:15	5.65	FOMC lowers Fed funds target 25bps to 4.5%
11-Dec-07 14:15	4.75	FOMC lowers Fed funds target 25bps to 4.25%
22-Jan-08 10:15	9.63	**FOMC lowers Fed funds target 75bps to 3.5%**
14-Mar-08 10:15	5.61	New York Fed drops deal to save Bear Stearns
18-Sep-08 15:15	11.96	**SEC short-selling ban; global campaign by central banks; Paulson briefs Congress on TARP**
10-Oct-08 15:15	24.80	**Stock market crashes in Asia, Europe, and the U.S.**
28-Oct-08 15:15	8.92	First round of TARP bank bailouts ($115B)
13-Nov-08 14:15	11.79	Large negative jobless claims surprise; most new claims since September 11, 2001
16-Dec-08 15:15	6.24	FOMC lowers Fed funds target to 0-0.25%
15-Jan-09 14:15	5.28	Senate approves release of $350 billion of TARP funds
6-May-10 15:15	17.44	**2010 Flash Crash**
5-Aug-11 10:15	6.19	S&P downgrades U.S. credit rating to AA+
22-Sep-11 10:15	4.17	**Stock markets plunge after Fed warns of "significant downside risks" to growth**

by a minimum distance of 10 trading days to isolate distinct events. The largest shocks to market tail risk consist primarily of major macroeconomic news and Federal Reserve policy changes. Scheduled and surprise events are captured “in progress”—for example, the 1:45–2:45pm window captures the typical timing of FOMC announcements, and the 2:45–3:45pm window on May 6, 2010 captures the Flash Crash.

Figure III indicates that the tail risk measure also captures well-known intraday patterns in volatility and jump risks (e.g., Andersen and Bollerslev (1997), Bollerslev and Todorov (2011b)). Both plots rank 2008–2009 as the most extreme years and 2005 as the least extreme year for expected jumps (left) and realized jumps (right). As with volatility, the pronounced skewed-U pattern manifests in tail risks for each year of the sample.

Although the tail risk measure matches the shape of intraday patterns, the magnitude disagrees with previous work. Comparing panels of Figure III, the anticipated tail risk measure implies roughly 100 times as much discontinuous variation as is realized in medium-scale jumps. For example, the opening hours of 2008 and 2009 average 3.15 and 2.13 percent anticipated jump variation per side against 2.92 and 2.11 basis points of realized jump variation per side.

This difference in magnitude has several potential causes. First, the tail risk measure captures the entire distribution of anticipated jumps larger than a few basis points. Frequent small jumps inflate the measure relative to realized medium-scale jump variation. Second, market intermediaries may be risk averse. In that case, the risk-neutral, zero-expected profit condition of Equation (1) omits a potentially large scaling term on picking-off costs, which in equilibrium, equates to greater bid-ask spreads or smaller quoted depth scaled in proportion to jump exposures. Third, other sources of the bid-ask spread may contaminate my measure. Section V.B addresses this possibility in depth. For all these reasons, subsequent analyses only consider the tail-risk measure up to scale.

Relation to Other Tail Risk and Volatility Measures

Table I relates my spread-implied tail risk measure to four other tail risk and volatility measures proposed in the literature: options-implied tail risks from Bollerslev and Todorov (2014); the VIX; hourly realized volatility; and the volume-synchronized probability of informed trading (VPIN) for

Figure III: **Intraday Jumps by Hour, 2004–2013**

This figure plots hourly means of market tail risks by hour and year (top) and of weighted realized basis point jumps by hour and year (bottom). Realized basis point jumps are a weighted sum of the number of events in which the minutely return exceeds 5, 10, 25, and 100 basis points, with respective weights of 5, 10, 25, and 100. The total jump variation is then divided by two to reflect average positive- and negative-jump variation.

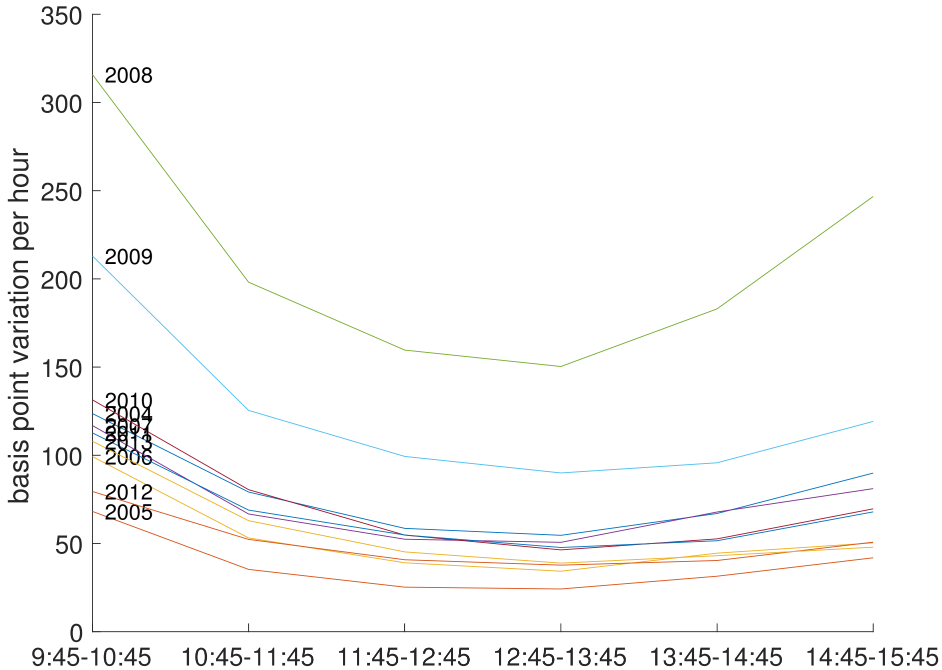

Anticipated Jump Risks

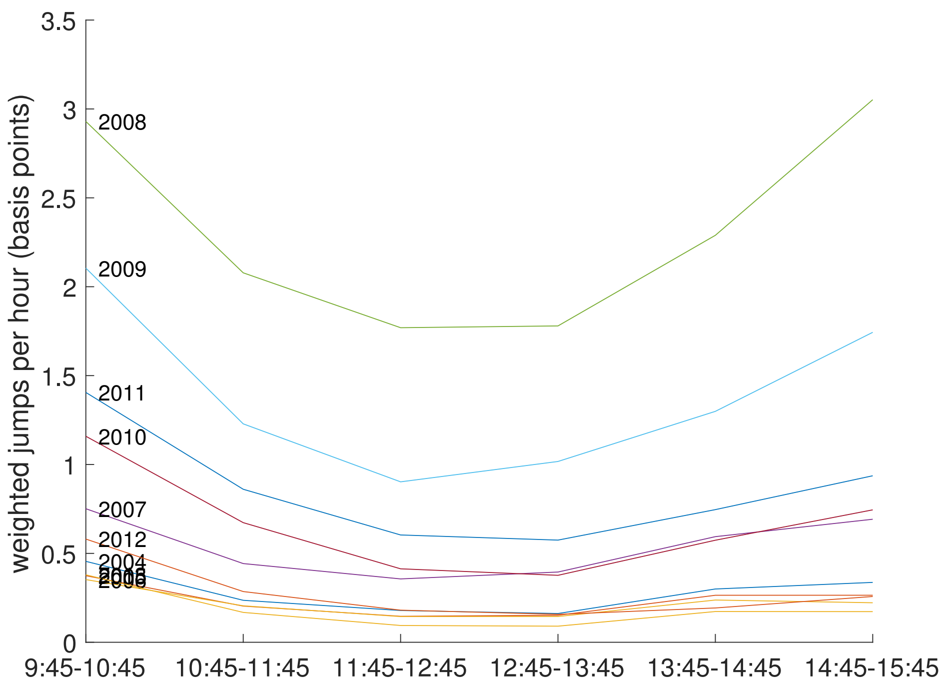

Realized Basis Point Jumps

Table I: **Correlations of Tail Measure with Other Volatility and Tail Measures**

This table reports correlations of tail and volatility measures over the 2004–2013 sample period. The spread-implied measure ξ_{MKT} uses Equation (6) to compute hourly market tail risk estimates. VIX is the (30-day) CBOE Volatility Index. Realized volatility is the square root of the average squared one-minute SPY returns within each hour. Options-implied tails (O-I Tail) are the weekly parametric left-tail risk estimates from Figure 7 of Bollerslev and Todorov (2014). VPIN uses bulk-volume classification and volume bars (10 buckets) for the front-month E-mini S&P 500 futures contract. These two series are available through 2011 only. Daily and weekly values are equal-weighted hourly values within the respective time bin.

Weekly Correlations

	ξ_{MKT}	O-I Tail	VIX	Realized Vol.	VPIN
ξ_{MKT}	–	0.75	0.83	0.83	0.82
O-I Tail	0.75	–	0.89	0.75	0.65
VIX	0.83	0.89	–	0.86	0.79
Realized Vol.	0.83	0.75	0.86	–	0.74
VPIN	0.82	0.65	0.79	0.74	–

Daily Correlations

	ξ_{MKT}	VIX	Realized Vol.	VPIN
ξ_{MKT}	–	0.76	0.62	0.75
VIX	0.76	–	0.66	0.75
Realized Vol.	0.62	0.66	–	0.57
VPIN	0.75	0.75	0.57	–

Hourly Correlations

	ξ_{MKT}	VIX	Realized Vol.	VPIN
ξ_{MKT}	–	0.65	0.43	0.65
VIX	0.65	–	0.52	0.74
Realized Vol.	0.43	0.52	–	0.45
VPIN	0.65	0.74	0.45	–

E-mini S&P 500 futures.[11] Daily and weekly values are equal-weighted averaged hourly values within each respective time bin.

All volatility and tail risk measures are highly correlated at the daily and weekly frequencies. Relative to the options-implied tails, my spread-implied tails are more similar to realized volatility and VPIN and less similar to the VIX. This relationship is expected in that the spread-based measure draws on short-horizon market maker expectations of extreme-event risks, whereas the options-implied measure reflects the long-horizon information embedded in options with more than a week to expiration. Focusing on a competing high-frequency extreme-event risk measure, VPIN is highly correlated with my measure on the weekly horizon, and it becomes less correlated as frequency increases (in tandem with the VIX). I show in Section V.A that this imperfect correlation is associated with the different measures picking up distinct signals for explaining and forecasting jump realizations.

A. *Empirical Tests*

The key verification regression takes the following form for the market factor (and is later repeated for a financial sector proxy in Section VI.C):

$$\begin{aligned} tail\_realization_t \quad = \quad & \alpha + \beta \xi_{t-\Delta,MKT} + \gamma VIX_{t-\Delta} + \delta CV_{t-\Delta} + \\ & \zeta VPIN_{t-\Delta} + \alpha_{-1} tail\_realization_{t-1} + \\ & \beta_{-1} \xi_{t-1-\Delta,MKT} + \gamma_{-1} VIX_{t-1-\Delta} + \\ & \delta_{-1} CV_{t-1-\Delta} + \zeta_{-1} VPIN_{t-1-\Delta} + \epsilon_t. \end{aligned} \tag{7}$$

I measure tail realizations in units of spreads and basis points and in event counts and event sums (weighting by event size). Given spread-implied tail risk's strong comovement with other forward-looking variation measures, e.g., the VIX, I include the VIX as a control to ensure that the tail measure indeed has additional explanatory power for tail events. I also include total continuous

[11] VPIN has been updated since its inception in Easley, López de Prado, and O'Hara (2012). I use bulk volume classification with wide volume bars (10 buckets) as suggested by the classification accuracy tests of Easley, de Prado, and O'Hara (2016). See Andersen and Bondarenko (2015) for additional discussion of several variants of VPIN.

variation CV to isolate the contribution of the tail risk measure to explaining jumps rather than continuous variation. CV is defined as the sum of squared minutely price movements smaller than 2.5 standard deviations of minutely price movements adjusted for time-of-day effects within each year (following the continuous and jump variation decomposition of Mancini (2009), among others, and Bollerslev, Todorov, and Li (2013)'s methodology specifically).[12] Realized continuous variation comoves very strongly with jump variation, so including it as a control presents a particularly strong test of the interpretation of the recovered coefficients as an estimate of anticipated jump tails. Finally, I include the latest bulk-volume classified VPIN measure to distinguish the information content of my tail risk measure from a leading alternative.

The tail realization measures used in the regression are as follows. Basis-point jumps count the number of events in which the minutely midpoint return exceeds 10 basis points, my implicit "large jump" threshold. The jump sum is a weighted count of the number of events in which the minutely return exceeds 5, 10, 25, and 100 basis points, with respective weights of 5, 10, 25, and 100. Spread jumps count the number of occasions in which the minutely return exceeds 5 quoted half-spreads. Because bid-ask spreads widen with volatility, this count measure mechanically provides a partial control for time-varying volatility. The corresponding jump sum measure is a weighted count of the number of events in which the minutely return exceeds 1, 5, 10, and 25 half-spreads, with concomitant weights of 1, 5, 10, and 25.

I use relatively simple measures of realized tail events for two reasons. First, large price movements generate picking-off opportunities regardless of the underlying volatility environment so long as they move the asset's "latent price" outside of the spread within a very short time period. The model does not distinguish between rare, truly discontinuous price movements and extremely rapid continuous ones associated with high local volatility.[13] Both harm liquidity providers as other market participants pick off stale quotes. Second, extreme market movements are quite rare, and including moderately large market movements dramatically increases the number of non-zero obser-

[12] I consider two alternative measures of continuous variation in the Online Appendix. Coefficient estimates on ξ_{MKT} are robust to the way CV is constructed.

[13] Large basis point and spread movements typically include the infrequent jumps captured by formal jump detection techniques (e.g., Lee and Mykland (2008) and Bollerslev, Todorov, and Li (2013)), but they also include more frequent medium-scale price movements that may not register as jumps in high volatility environments.

vations for the dependent variable. The weighted count measures in particular strike a compromise by including moderately extreme events and upweighting truly extreme factor realizations.

I run the regression of Equation (7) for two horizons. First, I run the contemporaneous regression with $\Delta = 0$. This regression should have high explanatory power if the model is true because market makers adjust their quotes nearly every instant to reflect anticipated risks. In this specification, the recovered tail-risk measure $\xi_{t,MKT}$ has the interpretation of the *within-hour average anticipated jump risk*. Second, I implement a true forecasting regression with $\Delta = 1$. This specification tests whether the lagged tail risk measure predicts future tail realizations over the next hour. Recognizing that jump intensities and volatility are persistent, I add lagged tail realizations and explanatory variables in both specifications. These additions challenge the tail-risk measure because the sizable persistence in volatility and jump risks is differenced out. Throughout I drop observations with lags using information from the last hour of the preceding trading day, though results are robust to this choice.

The first panel of Table II presents results from the baseline test of contemporaneous forecast jump tails on realized jumps. I normalize ξ_{MKT} and VPIN by dividing by their standard deviations to facilitate interpretation of coefficients (for comparison, the standard deviation of the VIX in this period is 10.3). For all tail realization measures, an elevated market tail measure coincides with an increase in the number of realized jumps within the hour, and the coefficients are statistically significant and economically large.[14] For example, a one standard deviation increase in the market jump tail risk is associated with 4.4 additional realized basis-point jumps per trading hour and 70.9 additional weighted jumps (the measure captures both intensity and size). The coefficient on the jump-tail measure is only slightly reduced when adding the VIX as a control, and they are both associated with near-term jumps. The coefficient on continuous variation is inconsistent or driven out by the jump tail risk measure; by contrast, the jump-tail estimates perform well in explaining the residual variation in realized jumps, which supports its interpretation as a measure of *extreme event risk* rather than of contemporaneous or anticipated *volatility*. In most specifications, including

[14] Standard errors are robust to heteroskedasticity (both panels) and serial correlation of up to 126 trading hours. Adding hour-year fixed effects has minimal impact on point estimates or their statistical significance. Results are also not driven by extreme observations by recomputing both tables by taking logs of one plus the dependent variable.

Table II: **Contemporaneous and Predictive Regressions for Jump Realizations**

This table presents results from regressions of realized jumps against estimated tail risks,

$$\begin{aligned} tail\_realization_t &= \alpha + \beta\xi_{t-\Delta,MKT} + \gamma VIX_{t-\Delta} + \delta CV_{t-\Delta} + \zeta VPIN_{t-\Delta} + \\ &\quad \alpha_{-1} tail\_realization_{t-1} + \beta_{-1}\xi_{t-1-\Delta,MKT} + \\ &\quad \gamma_{-1} VIX_{t-1-\Delta} + \delta_{-1} CV_{t-1-\Delta} + \zeta_{-1} VPIN_{t-1-\Delta} + \epsilon_t, \end{aligned}$$

where $\Delta = 0$ and $\Delta = 1$ correspond to the first and second subtables, respectively. Tail realizations are measured in counts of minutely returns exceeding basis point or spread thresholds. The count variable sums jumps exceeding 10 basis points or 5 half-spreads, and the sum variables are a weighted sum of jump sizes exceeding 5, 10, 25, and 100 basis points or 1, 5, 10, or 25 half-spreads. Continuous variation is estimated by hour with a 2.5 standard deviation threshold on minutely price movements. VPIN is constructed using bulk-volume classification with volume bins on front-month E-mini S&P 500 futures. Regressions consist of hourly observations for the 2005–2013 sample in the baseline specification and for 2005-2011 where VPIN is included. Standard errors are HAC with monthly (126 observation) bandwidth. $\xi_{t,MKT}$ and VPIN are normalized by their standard deviations in both panels. t-statistics are in parentheses.

SPY Basis-Point Jumps

	Jump Count			Jump Sum		
ξ_{MKT}	4.41***	4.03***	3.91***	70.89***	64.78***	57.30***
	(7.39)	(7.32)	(11.56)	(9.82)	(10.24)	(13.44)
VIX		0.84***	0.82***		14.27***	13.72***
		(5.37)	(5.11)		(9.89)	(8.41)
CV			-0.57			10.56
			(-0.36)			(0.58)
$VPIN$			0.97***			25.31***
			(2.69)			(8.46)
Obs.	11280	11280	6398	11280	11280	6398
R^2	0.86	0.87	0.88	0.89	0.91	0.91

SPY Spread Jumps

	Jump Count			Jump Sum		
ξ_{MKT}	6.29***	5.42***	5.74***	59.43***	54.83***	51.74***
	(4.72)	(4.24)	(6.54)	(7.57)	(7.16)	(10.24)
VIX		1.64***	1.32***		11.95***	10.01***
		(7.36)	(4.84)		(7.77)	(5.88)
CV			-4.21***			-7.47
			(-3.63)			(-0.52)
$VPIN$			4.63***			27.70***
			(6.32)			(6.32)
Obs.	11280	11280	6398	11280	11280	6398
R^2	0.77	0.78	0.82	0.84	0.85	0.87

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table II: **Contemporaneous and Predictive Regressions for Jump Realizations (Cont.)**

SPY Basis-Point Jumps

	Jump Count			Jump Sum		
ξ_{MKT}	6.39***	5.14***	3.25***	100.34***	79.52***	48.76***
	(11.75)	(11.54)	(6.27)	(12.36)	(12.28)	(5.92)
VIX		1.01***	0.92***		17.54***	15.85***
		(6.62)	(4.73)		(8.21)	(5.55)
CV			7.53**			110.02***
			(2.35)			(3.41)
$VPIN$			0.63			22.90***
			(0.99)			(2.99)
Obs.	9016	9016	5114	9016	9016	5114
R^2	0.73	0.76	0.80	0.77	0.80	0.83

SPY Spread Jumps

	Jump Count			Jump Sum		
ξ_{MKT}	10.41***	7.84***	6.42***	88.34***	70.98***	51.04***
	(6.91)	(5.35)	(5.20)	(10.24)	(8.31)	(5.60)
VIX		1.70***	1.36***		15.21***	12.59***
		(7.21)	(4.66)		(9.54)	(6.05)
CV			-0.70			50.83**
			(-0.27)			(2.18)
$VPIN$			5.01***			26.08***
			(8.18)			(4.91)
Obs.	9016	9016	5114	9016	9016	5114
R^2	0.62	0.66	0.72	0.68	0.71	0.76

*** $p < 0.01$, ** $p < 0.05$, $^{*}p < 0.1$

volume-bar VPIN attenuates the explanatory power of ξ_{MKT} only slightly, although both variables explain substantial variation in realized jumps. The second panel of Table II presents analogous results for forecasting using lagged tail risk measures. All specifications indicate comparably strong predictability of market jump realizations for the hour ahead.

These tests confirm the suggestive evidence of Figures I–III. The market tail risk measure is associated with both low- and high-frequency realized jump risks. It contains some of the same information as the forward-looking VIX, but its dynamics are intermediate between those of the VIX and of near-term realized volatility. Importantly, the tail risk measure is not spanned by measures of continuous variation, and in fact it drives out these measures for realized spread jumps. These features accord with the design of the measure as a tool for assessing instantaneous extreme event risks. In addition, the tail risk measure forecasts well near-term market jump events.

The spread-implied tail risk measure substantially boosts predictive power over lagged tail realizations and other covariates. Lags alone deliver a period-ahead forecast R^2 of 61% (not tabulated) because tail risk is highly persistent. ξ_{MKT} explains more than 30% of the remaining variation to increase R^2 to 73%. By contrast, VIX alone increments R^2 only to 69% (not tabulated) for realized basis-point jumps.

To provide a back-of-the-envelope estimate of the economic importance of the increase in explanatory power from using my measure, suppose that jumps are independent with average size 15 basis points and that an investor constructs ξ_{MKT} in real time. With an average of 2.9 basis point jumps per hour and approximately 1,600 trading hours per year, annualized "medium-scale" jump variation is 7.0%. An investor equipped to avoid an additional 12% of this variation avoids 0.84% market volatility, and using the VIX and lagged covariates as a benchmark predictor, adding my measure allows the investor to avoid an additional 4% of this variation and avoid 0.28% market volatility. Relative to annual SPY volatility of 19.0% during this period, these values represent moderate 4.4% and 1.5% reductions in exposure to market volatility, respectively.

B. Other Sources of the Spread

Inventory Risk

Risk-averse market makers must be compensated for exposure to price variation of assets in inventory. If market makers do not cheaply hedge inventory risks, e.g., using liquid factor-mimicking indexes, the component of the bid-ask spread associated with inventory risk may contaminate tail-risk estimates. To address this possibility, I isolate the component of spreads contributed by adverse selection risks—including picking-off risk—rather than by inventory risks.

I follow Glosten (1987) and decompose the effective (bid-ask) spread into "adverse selection" and "realized spread" components. The adverse selection component captures losses to the market maker when she sells (buys) at t and the price subsequently increases (decreases), and the realized spread component captures the difference between the total spread collected—including compensation for inventory risk—and this "permanent" price impact. Glosten (1987)'s now-standard proxy for the adverse selection component of spreads is given by the scaled change in midpoint prices between time t and $t+5$ minutes,

$$h_{it}^{*} = q_{it}\left(\frac{m_{i,t+5m} - m_{it}}{m_{it}}\right), \tag{8}$$

where m_{it} is the prevailing quote midpoint in security i at time t, $q_{it} = +1$ for market-maker sells, and $q_{it} = -1$ for market-maker buys. Buys and sells are determined by the Lee and Ready (1991) algorithm. I average this adverse selection value by stock and minute to obtain a continuous proxy for adverse selection costs, and I substitute h^{*} for effective half spreads in constructing a modified liquidity composite, Vh^{*}/d.

Table III replicates Table II using the modified liquidity composite as the dependent variable. All coefficients retain comparable levels of economic and statistical significance, but point estimates are typically smaller using the adverse selection component of the spread. The similarity of the economic relations and explained variation in realized jumps with and without compensation for inventory risks in the liquidity composite indicates that the tail risk measure is not materially undermined by variation in inventory risk throughout the sample. As Figure IV suggests, the first-order impact of using the modified liquidity composite is a simple rescaling of recovered market tail

Table III: **Contemporaneous and Predictive Regressions for Jump Realizations with an Adverse Selection Proxy**

This table presents results from regressions of realized jumps against estimated tail risks,

$$\begin{aligned} tail\_realization_t &= \alpha + \beta \xi_{t-\Delta,MKT} + \gamma VIX_{t-\Delta} + \delta CV_{t-\Delta} + \zeta VPIN_{t-\Delta} + \\ & \alpha_{-1} tail\_realization_{t-1} + \beta_{-1}\xi_{t-1-\Delta,MKT} + \\ & \gamma_{-1} VIX_{t-1-\Delta} + \delta_{-1} CV_{t-1-\Delta} + \zeta_{-1} VPIN_{t-1-\Delta} + \epsilon_t, \end{aligned}$$

where $\Delta = 0$ and $\Delta = 1$ correspond to the upper and lower panels, respectively. Tail estimates are constructed using the adverse selection component of the spread rather than effective spreads in the liquidity composite. Tail realizations are measured in counts of minutely returns exceeding basis point or spread thresholds. The count variable sums jumps exceeding 10 basis points or 5 half-spreads, and the sum variables are a weighted sum of jump sizes exceeding 5, 10, 25, and 100 basis points or 1, 5, 10, or 25 half-spreads. Continuous variation is estimated by hour with a 2.5 standard deviation threshold on minutely price movements. VPIN is constructed using bulk-volume classification with volume bins on front-month E-mini S&P 500 futures. Regressions consist of hourly observations for the 2005–2013 sample in the baseline specification and for 2005-2011 where VPIN is included. Standard errors are HAC with monthly (126 observation) bandwidth. $\xi_{t,MKT}$ and VPIN are normalized by their standard deviations in both panels. t-statistics are in parentheses.

		SPY Basis-Point Jumps				
		Jump Count			Jump Sum	
ξ_{MKT}	2.19***	2.04***	1.34***	34.96***	32.58***	19.01***
	(6.03)	(6.28)	(6.18)	(8.49)	(9.13)	(6.70)
VIX		1.02***	0.97***		17.18***	15.86***
		(8.29)	(6.48)		(10.31)	(7.24)
CV			2.82***			60.94***
			(1.57)			(2.68)
$VPIN$			2.14***			43.31***
			(5.13)			(8.42)
Obs.	11280	11280	6398	11280	11280	6398
R^2	0.82	0.84	0.86	0.85	0.88	0.89

		SPY Spread Jumps				
		Jump Count			Jump Sum	
ξ_{MKT}	3.04***	2.55***	1.98***	30.24***	27.40***	18.29***
	(4.13)	(3.75)	(4.46)	(6.37)	(6.28)	(5.66)
VIX		1.88***	1.55***		14.37***	12.00***
		(6.77)	(5.27)		(9.04)	(6.25)
CV			0.60			35.36**
			(0.39)			(1.98)
$VPIN$			6.31***			43.24***
			(10.11)			(10.23)
Obs.	11280	11280	6398	11280	11280	6398
R^2	0.74	0.76	0.80	0.80	0.82	0.85

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table III: **Contemporaneous and Predictive Regressions for Jump Realizations with an Adverse Selection Proxy (Continued)**

	SPY Basis-Point Jumps					
	Jump Count			Jump Sum		
ξ_{MKT}	3.23***	2.61***	1.16***	50.87***	40.83***	17.74***
	(11.31)	(9.95)	(4.14)	(11.51)	(10.00)	(4.10)
VIX		1.17***	0.96***		20.17***	16.73***
		(9.46)	(5.49)		(11.27)	(6.29)
CV			12.10***			175.23***
			(3.61)			(5.01)
$VPIN$			1.47**			35.94***
			(2.34)			(5.03)
Obs.	9016	9016	5114	9016	9016	5114
R^2	0.68	0.74	0.79	0.71	0.77	0.82

	SPY Spread Jumps					
	Jump Count			Jump Sum		
ξ_{MKT}	5.38***	3.93***	2.30***	46.57***	36.50***	19.25***
	(7.11)	(4.99)	(4.11)	(10.73)	(7.92)	(4.77)
VIX		1.92***	1.50***		17.29***	13.55***
		(5.50)	(4.29)		(10.85)	(6.39)
CV			7.51**			117.14***
			(2.27)			(4.29)
$VPIN$			6.72***			39.95***
			(11.19)			(8.47)
Obs.	9016	9016	5114	9016	9016	5114
R^2	0.56	0.63	0.70	0.61	0.67	0.74

*** $p < 0.01$, ** $p < 0.05$, $^{*}p < 0.1$

Figure IV: **Comparison of Market Jump Risks Net of Alternative Sources of the Spread**
Figures plot rolling ten-day means of hourly estimated market tail risks for each trading date in 2004–2013. The solid blue line is the baseline estimation of Equation (6). The dashed red line replaces the effective half-spread with the adverse selection component of the spread. The dotted yellow line adds a stock-quarter control for the probability of informed trading. The bottom plot aligns the series by subtracting series means and dividing by series standard deviations.

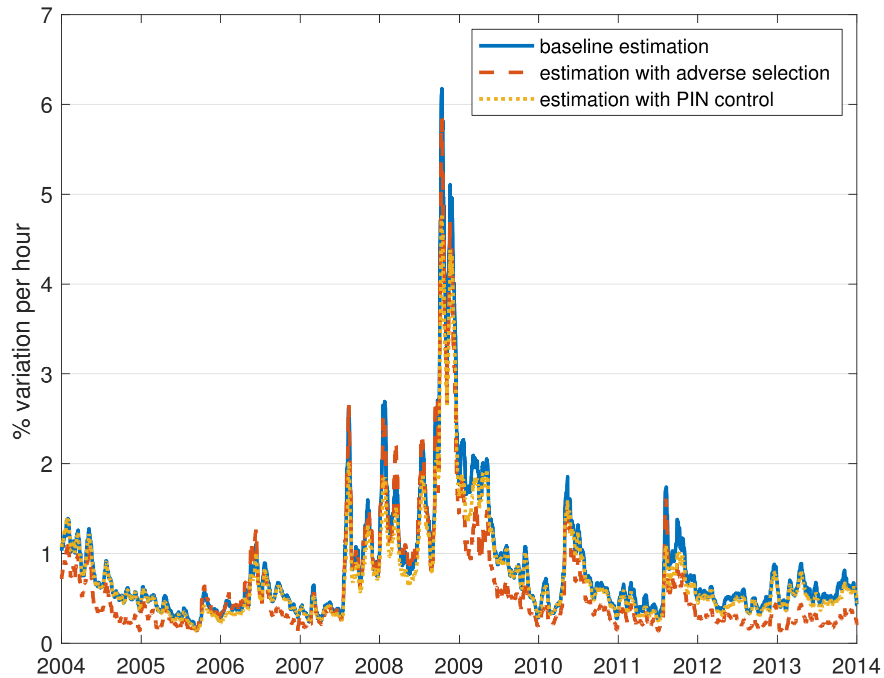

Forecast Market Tail Risks

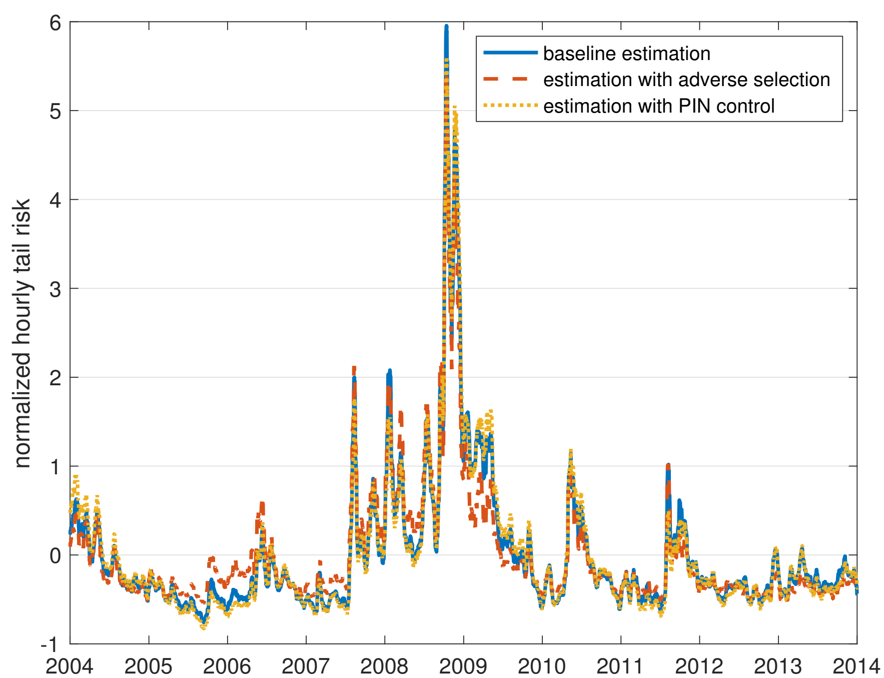

Forecast Market Tail Risks (Normalized)

risks, and the correlation between the modified and unmodified tail risk series exceeds 92%.

At the same time, the smaller coefficients obtained using h^* rather than h indicate either that liquidity providers require greater compensation for their services when market jump risk is high or that the lagged tail-risk measure has greater auxiliary forecasting power than before, likely due to measurement error in the modified series. Two facts point to the latter explanation rather than an omitted-variable bias: (1) correlations between the two series are very high, but (2) using the modified liquidity composite reduces the hourly autocorrelation of the tail-risk series from 79% to 66%. That the VIX, continuous variation, and VPIN appear to absorb some of the baseline tail-risk measure's explanatory power also points to a measurement-error explanation for the smaller coefficients.

Non-Jump Adverse Selection

Adverse selection imposes costs on market makers through two qualitatively different modes:

1. Intermediation against informed traders with long-lived information ("slow");
2. Picking off by stale-quote snipers ("fast").

Cross-sectional variation in spreads reflects differences in exposures to long-lived informational risk in addition to picking-off risk. The tail-risk measure suffers from omitted-variable bias if assets with high betas are more exposed to non-jump or "slow" adverse selection risk.

Such alignment is tantamount to market participants having private information on the underlying factor. Gorton and Pennacchi (1993) and others suggest that private informational advantages are unlikely for systematic factors. They argue that a key advantage of composite or factor-mimicking products such as the SPY is their low risk of slow adverse selection because insider information is typically known at the security level rather than at the aggregate level.

I support this argument by showing that controlling for stock-level slow adverse selection does not meaningfully affect recovered tail risks for the market factor. Specifically, I include the probability of informed trading (PIN) measure of Easley and O'Hara (1992) and Easley, Kiefer, O'Hara, and Paperman (1996) to control for the arrival rate of informed traders. The PIN measure is constructed under the assumption that order flow tilts in the direction of information that persists throughout the trading day unbeknownst to the market maker. Such information is long-lived with respect to the horizon of HFT market makers, and as such, PIN should primarily capture costs of slow rather than fast adverse selection. I compute stock-level PIN estimates quarterly for the 2005–2013 sample period using the methodology of Yan and Zhang (2014).[15]

Table IV reports the results of the contemporaneous regression of realized jumps on implied market tail risk, where the implied tail risk estimation equation adds the stock-date PIN characteristic to the right-hand side of Equation (6). Like the tail-risk measure net of potential inventory risk, the tail risk measure net of slow adverse selection risk performs similarly to the baseline specification in matching time variation in market tail realizations. Indeed, no tail risk coefficient in Table IV is appreciably different from its counterpart in Table II, and the time-series correlation of baseline and PIN-adjusted series is 98%. Omitted stock-level "slow" adverse selection does not bias my measure because it exhibits low cross-sectional correlation with factor betas; although slow adverse selection may be important at the security level, it "washes out" in my estimate of *factor* risks.

C. Other Sources of Omitted Variable Bias

The discussion of inventory risk and non-jump adverse selection in Section V.B confronts the most likely empirical threats to the interpretation of ξ_{kt} as tail risk for factor k at date t.[16] Short of conducting an instrumental variables analysis, however, it is unlikely that omitted-variable bias can be completely excluded in interpreting ξ_{kt}. This said, there are three reasons to believe that my tail risk measure is not seriously contaminated by omitted factors.

First, several recent papers document a low-dimensional high-frequency factor structure, by contrast with the higher-dimensional factor structure present in low-frequency data and mimicked by models such as the Barra U.S. Equity Model (USE4). Pelger (2017) finds a market factor and three industry factors (finance, oil, and electricity). Aït-Sahalia and Xiu (2018) recover a market

[15]An active literature describes issues encountered in estimating PIN in fast-moving equity markets. Results are robust across alternate PIN estimation methodologies.

[16]Following the recommendation of a referee, I also control for potential bias associated with idiosyncratic risks using idiosyncratic volatility and idiosyncratic volatility factor loadings in the Online Appendix. My tail risk estimates are virtually unchanged in both cases.

Table IV: **Contemporaneous and Predictive Regressions for Jump Realizations Controlling for PIN**

This table presents results from regressions of realized jumps against estimated tail risks,

$$\begin{aligned} tail\_realization_t &= \alpha + \beta \xi_{t-\Delta,MKT} + \gamma VIX_{t-\Delta} + \delta CV_{t-\Delta} + \zeta VPIN_{t-\Delta} + \\ & \alpha_{-1} tail\_realization_{t-1} + \beta_{-1}\xi_{t-1-\Delta,MKT} + \\ & \gamma_{-1} VIX_{t-1-\Delta} + \delta_{-1} CV_{t-1-\Delta} + \zeta_{-1} VPIN_{t-1-\Delta} + \epsilon_t, \end{aligned}$$

where $\Delta = 0$ and $\Delta = 1$ correspond to the upper and lower panels, respectively. Tail estimates include a stock-quarter control for the probability of informed trading. Tail realizations are measured in counts of minutely returns exceeding basis point or spread thresholds. The count variable sums jumps exceeding 10 basis points or 5 half-spreads, and the sum variables are a weighted sum of jump sizes exceeding 5, 10, 25, and 100 basis points or 1, 5, 10, or 25 half-spreads. Continuous variation is estimated by hour with a 2.5 standard deviation threshold on minutely price movements. VPIN is constructed using bulk-volume classification with volume bins on front-month E-mini S&P 500 futures. Regressions consist of hourly observations for the 2005–2013 sample in the baseline specification and for 2005-2011 where VPIN is included. Standard errors are HAC with monthly (126 observation) bandwidth. $\xi_{t,MKT}$ and VPIN are normalized by their standard deviations in both panels. t-statistics are in parentheses.

SPY Basis-Point Jumps

	Jump Count			Jump Sum		
ξ_{MKT}	4.00***	3.61***	3.11***	64.07***	57.67***	44.98***
	(12.78)	(12.91)	(10.09)	(18.62)	(20.37)	(11.58)
VIX		0.95***	0.92***		15.96***	15.17***
		(7.05)	(6.09)		(12.49)	(9.23)
CV			1.48			41.72**
			(0.88)			(2.18)
$VPIN$			1.18***			28.76***
			(3.29)			(7.35)
Obs.	11280	11280	6398	11280	11280	6398
R^2	0.85	0.86	0.87	0.88	0.89	0.91

SPY Spread Jumps

	Jump Count			Jump Sum		
ξ_{MKT}	6.08***	5.20***	4.91***	54.84***	50.06***	42.18***
	(6.46)	(5.85)	(7.98)	(12.19)	(11.75)	(10.90)
VIX		1.78***	1.46***		13.34***	11.30***
		(8.17)	(5.85)		(10.39)	(7.43)
CV			-1.76			18.03
			(-1.46)			(1.19)
$VPIN$			4.73***			29.57***
			(6.87)			(7.19)
Obs.	11280	11280	6398	11280	11280	6398
R^2	0.77	0.78	0.82	0.83	0.84	0.87

*** $p < 0.01$, ** $p < 0.05$, $^{*}p < 0.1$

Table IV: **Contemporaneous and Predictive Regressions for Jump Realizations Controlling for PIN (Continued)**

	SPY Basis-Point Jumps					
	Jump Count			Jump Sum		
ξ_{MKT}	5.47***	4.26***	2.48***	86.26***	65.93***	37.23***
	(18.90)	(17.96)	(6.65)	(21.12)	(19.25)	(7.38)
VIX		1.19***	1.01***		20.28***	17.25***
		(7.81)	(5.38)		(9.61)	(6.15)
CV			9.22***			135.86***
			(2.69)			(3.91)
$VPIN$			0.73			24.46***
			(1.21)			(3.26)
Obs.	9016	9016	5114	9016	9016	5114
R^2	0.72	0.76	0.80	0.75	0.79	0.83

	SPY Spread Jumps					
	Jump Count			Jump Sum		
ξ_{MKT}	9.15***	6.65***	5.04***	76.37***	59.43***	39.77***
	(9.53)	(7.08)	(6.70)	(17.34)	(12.91)	(7.65)
VIX		1.97***	1.53***		17.62***	14.04***
		(7.98)	(5.48)		(11.77)	(7.04)
CV			2.38			76.11***
			(0.79)			(2.93)
$VPIN$			5.12***			27.06***
			(8.35)			(5.26)
Obs.	9016	9016	5114	9016	9016	5114
R^2	0.61	0.65	0.72	0.67	0.70	0.76

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

factor and a financial industry factor. Li, Todorov, Tauchen, and Lin (2017) find a one-factor structure around market jump events. Taken together, these papers suggest that other factors in the data are less important for explaining covariation in high-frequency returns, and by extension, my study of market and financial factors is unlikely to be threatened by omitted factors.

Second, several high-frequency signals used by practitioners do not necessarily have an aggregate impact that would generate persistent risk exposures. Chinco, Clark-Joseph, and Ye (2017) offer one such example. In their paper, the authors extract short-lived predictors using linear combinations of stock returns. The set of securities with predictable returns varies rapidly—on the order of minutes rather than months—as do the predicting stocks. They do not find stable factor relationships that would (1) deliver nonzero annual time-series betas and (2) contribute to omitted-variable bias in the cross-sectional slopes. Instead, such short-lived predictors would manifest as idiosyncratic picking-off risks to liquidity providers and show up in $\tilde{\xi}$.

Third, *all* of the largest increases in recovered extreme event risk for the market (financial sector) correspond with major market news (financial sector news), as described in Figure II (Figure VIII). At least for the most important market events, contamination by omitted factors seems unlikely—jumps in other factors with betas correlated with market or financial sector betas do not drive the innovations of greatest interest in this study.

VI. Applications

A. The 2010 Flash Crash

In a spectacular market episode, the May 6, 2010 Flash Crash saw equity indices decline by 5–6% and revert almost completely within a 30-minute period. Assessing welfare consequences associated with the 2010 Flash Crash has proved even more challenging than explaining the event's causes.[17] While Kirilenko, Kyle, Mehrdad, and Tuzun (2017) tabulate buyers and sellers in S&P 500 E-mini futures ("E-mini") during the Flash Crash, no corresponding data exists to evaluate the

[17] Explanations for the Flash Crash abound. Among these are that a single large trader's faulty algorithm caused a severe order flow imbalance (CFTC and SEC (2010)); extreme order flow toxicity drove away market makers and collapsed liquidity (Easley, López de Prado, and O'Hara (2012)); and a breakdown in cross-market arbitrage brought about an extreme price of immediacy (Menkveld and Yueshen (2017)).

redistributive consequences of the extreme turnover in equities and index products. Moreover, much popular discussion following the 2010 Flash Crash centers on distrust of the market mechanism and fears of future crashes, yet such concerns are inherently hard to quantify.

My measure of instantaneous jump risks is well-suited to evaluating the costs of rapid jump events. I require only that such events exceed the market makers' typical holding period and thus contribute to picking-off risk. Market makers fear picking off on both the initial price decline (or rise) and on the return because extended price disruptions of several minutes affect a security's "terminal value" with respect to the market maker's trading horizon.[18]

To demonstrate its utility in assessing the costs of flash crashes, I construct tail risk measures around and during May 6, 2010 using the one-factor market model. A one-factor market model is particularly apt in this instance because the 2010 Flash Crash originated in S&P 500 E-mini futures, a key price discovery market for the S&P 500. I estimate tail risks every 15 minutes to achieve high resolution on the crash interval (2:30–2:59pm) and surrounding trading hours.

Figure V plots market and idiosyncratic tail measures from 12:45pm on May 5, 2010 through 12:45pm on May 7, 2010 for each quarter hour from 9:45am to 3:45pm. To capture innovations and place risk changes in context of normal intraday and slow-moving macroeconomic variation, I difference the value during the same quarter hour on May 4, 2010, and divide by the standard deviation of differences for the same quarter hour over the preceding 63 trading days up to and including May 4, 2010 (a calendar quarter). Several new features are readily apparent. First, the Flash Crash itself is associated with extreme contemporaneous elevations of both the market (96 standard deviations) and idiosyncratic (73 standard deviations) tail risk measures. Second, jump risks remain elevated for the remainder of the trading day and throughout May 7, 2010, even after the initial shock subsides. Third, market tail risks increase a quarter hour *before* idiosyncratic tail risks, likely because the Flash Crash begins in the E-mini, a nearly ideal S&P 500 index proxy.[19] Volume-adjusted spreads widen and depth fall *proportionally to market betas* rather than in equal

[18] Kirilenko et al. (2017) find support for market makers not holding through "long" crashes. Rather than maintaining inventory during the 2010 Flash Crash, high-frequency market makers engaged in rapid turnover, or "hot potato" activity.

[19] The gap between the series is not driven by estimation error; the 95% confidence bands for the market and idiosyncratic tail risk series do not overlap over the 2:00–2:29pm window.

Figure V: **Standardized Deviations in Jump Expectations around the 2010 Flash Crash**
This figure plots standardized deviations in jump expectations around the May 6, 2010 Flash Crash. Tail risks are assessed with a market model with 15-minute increments. For each quarter hour, I normalize each value by subtracting the value during the same quarter hour on May 4, 2010, and dividing by the 15-minute specific standard deviation of this value across all dates in the 63 trading days up to and including May 4, 2010. The top figure plots the normalized value for the market factor before (green circles), during (red stars), and after (yellow diamonds) May 6, 2010. The dotted line plots the normalized 15-minute estimate for realized volatility. Black circles denote the 2:30–2:59pm interval during which the crash and reversion occur. The middle and bottom plots provide the corresponding information for the idiosyncratic jump factor and the VIX.

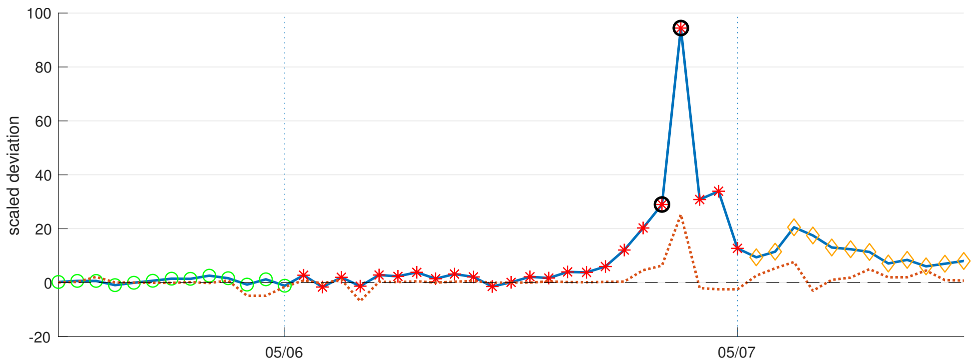

Anticipated Market Factor Jumps

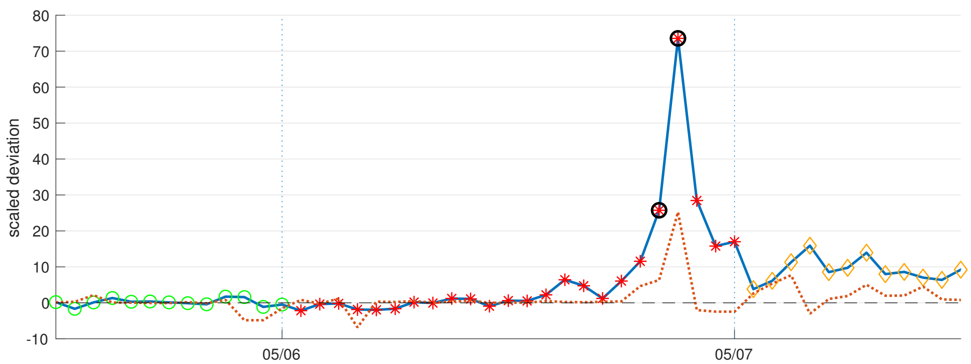

Anticipated Idiosyncratic Jumps

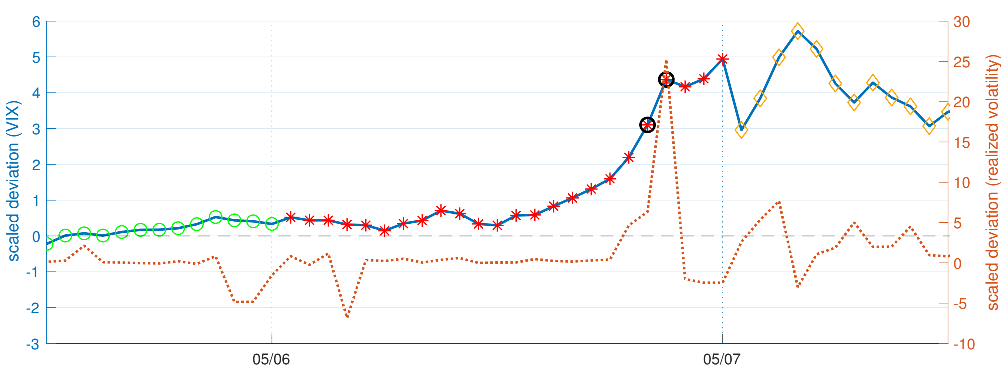

Intraday VIX

Figure VI: **Long-Term Effects of the 2010 Flash Crash on Implied Jump Risk**
This figure plots five-day backward-looking moving averages of the quarter-hour jump measure around the 2010 Flash Crash. Solid blue and dashed red lines correspond to implied market and idiosyncratic jump risks, with their associated scale on the left axis. The dotted yellow line is the VIX, and its scale is on the right axis. The vertical dashed line marks May 6, 2010, and the vertical dotted line marks two business days after the event, May 10, 2010.

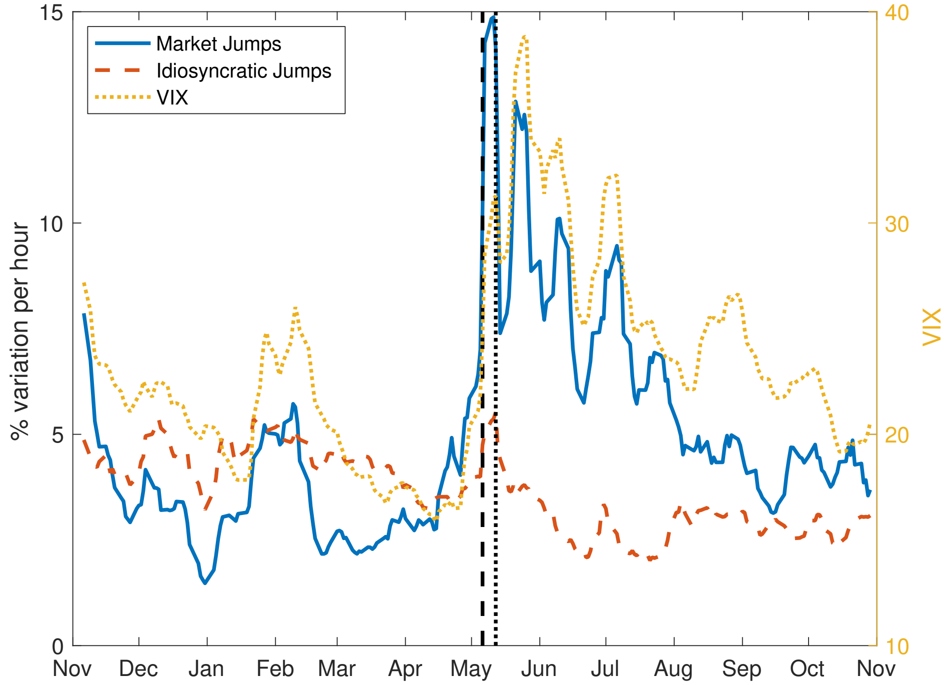

measure across stocks.

Intriguingly, the market tail risk measure increases by 20 standard deviations in the 2:15–2:30pm interval relative to its value on May 4, 2010. Market makers anticipate distress conditions even before Waddell & Reed initialized its trading algorithm at 2:32pm (Menkveld and Yueshen (2017)).[20] All told, the preceding relations align with several existing explanations of the Flash Crash and reassure that the proposed risk measure effectively anticipates near-term tail risks.

By contrast with my measure, options data used for constructing the VIX and other forward-looking risk measures incorporate volatility and jump information days or weeks beyond the duration of fleeting, mean-reverting flash crashes, and correspondingly are much less affected by such events. Although the VIX is somewhat elevated during the Flash Crash, the Flash Crash is not an extreme

[20] Moreover, as detailed in the Online Appendix, market buy depth systematically declines relative to market sell depth in line with market betas, suggesting that liquidity providers are especially concerned about the risk of market down jumps prior to the Flash Crash. The joint CFTC and SEC report on the 2010 Flash Crash documents a similar, lopsided deterioration of market depth for the E-Mini in the hour preceding the Flash Crash.

event for the VIX except in the rapidity of its increase intraday. An equally large and comparably sharp increase in implied volatility occurs in the same month: the normalized change-in-VIX measure achieves the same level on May 20, a day coinciding with a local maximum for the VIX (Figure VI). Likewise, my measure differs from VPIN in achieving an all-time high before the Flash Crash occurs and having no false positives for sufficiently high levels of implied market risk (Andersen and Bondarenko (2014a,b, 2015)).

The spread-implied measure also provides a longer-term view of changes in tail risk around the 2010 Flash Crash. The tail risk measure should remain elevated if the 2010 Flash Crash truly increases stability fears among market participants. Evidence for this effect is unambiguously negative. From Figure VI, we observe that both tail risk measures return to roughly their pre-Flash Crash levels only days later, and indeed tail risk in the week after the Flash Crash is statistically indistinguishable from tail risk in the week before it. Although longer-term average tail risks (and spreads) increase slightly in post-Crash weeks, these increases occur *after* May 10, 2010, several days after the crash. Subsequent tail risk elevations are inconsistent with a story of heightened perceived Flash Crash risk and likely arise from macroeconomic sources. In light of these results, it is difficult to argue that the 2010 Flash Crash had a persistent effect on market fears: high-frequency market makers should be among the most attuned to potential flash crash risk, yet their pricing of crash risks in spreads quickly reverts.

B. Federal Open Market Committee Announcements

The Federal Open Market Committee (FOMC) holds eight scheduled meetings per year to discuss salient economic and financial issues and policy responses. At the conclusion of each meeting, the FOMC releases a statement summarizing its views and actions. The release of these statements is among the most important scheduled macroeconomic news announcements. Several recent papers have documented empirical regularities associated with these announcements. Savor and Wilson (2013) and Lucca and Moench (2015) find that announcement-day average stock returns comprise a large fraction of the annual equity premium, and Savor and Wilson (2014) find that the CAPM works well for cross-sectional pricing during FOMC days.

Rational explanations for these phenomena require that risk or risk aversion be highly time-varying as measured in FOMC event time. Although realized market volatility is lower than average during the FOMC pre-announcement period, elevated and difficult-to-observe jump risk may offer a partial, rational explanation. High market jump risk requires a higher equity premium, and the increased importance of market jump risk can enforce the CAPM if the CAPM works for discontinuous returns (as suggested by Bollerslev, Li, and Todorov (2016)). Moreover, the sample of FOMC announcements may be too short for these jump risks to have been realized.

I apply my tail risk extraction methodology to analyze jump risks around FOMC announcements and find evidence against this hypothesis of elevated market tail risk during the high-return period. For each quarter-hour interval and calendar year, I compute the average FOMC announcement date tail risk, subtract the average non-FOMC announcement tail risk, and normalize by dividing by the respective standard deviation of tail risks across all days for each quarter hour and year. Importantly, the scheduling of announcements historically has not been precise enough to violate the assumption of Poisson arrivals of fundamental news: although market participants know the *planned* FOMC announcement time, there is significant uncertainty about precisely when the news comes out. This uncertainty in the past has been on the order of several minutes, as Figure VII illustrates, and translates in the context of the model into a sharply elevated jump risk for the interval containing the announcement.

Figure VII plots deviations in perceived tail risks around FOMC announcements. For every year in the sample, FOMC announcements indeed coincide with sharply elevated perceived tail risk relative to the non-FOMC dates in the same year.[21] Relative to the preceding quarter hour, most years also see a marked, anticipatory increase in implied tails in the quarter hour before the FOMC announcement (typically 2:00–2:14pm).[22] These anticipatory movements in tail risk can be explained by (1) uncertainty in the exact timing of the information release, as suggested by within-year dispersion of the announcement minute around the year's modal quarter hour, and

[21] Crisis years have a much larger unnormalized FOMC announcement effect, particularly in 2008. However, the large fluctuations in tail risk during 2008–2009 counterbalance the increased differences between FOMC and non-FOMC day means.

[22] These early-response results add color to Jiang, Lo, and Verdelhan (2011), who find increased spreads, decreased depth, and stagnant trading volume in the five minutes before major market news announcements in the U.S. Treasury bond market.

Figure VII: **Intraday Tail Risks around FOMC Announcements**

This figure plots intraday tail risks for the market factor for FOMC and non-FOMC announcement dates from 2005–2013. For each quarter-hour interval and calendar year, I compute the average FOMC announcement date tail risk and subtract the average non-FOMC announcement tail risk. I then normalize this quantity by the standard deviation of tail risks for all days in the same quarter hour and year. Stars indicate FOMC announcement times retrieved by minute from the first post-statement news article on Bloomberg or Dow Jones newswires following Fleming and Piazzesi (2005). The lower plot zooms in on the 1:45-2:45pm interval during which most announcements occur.

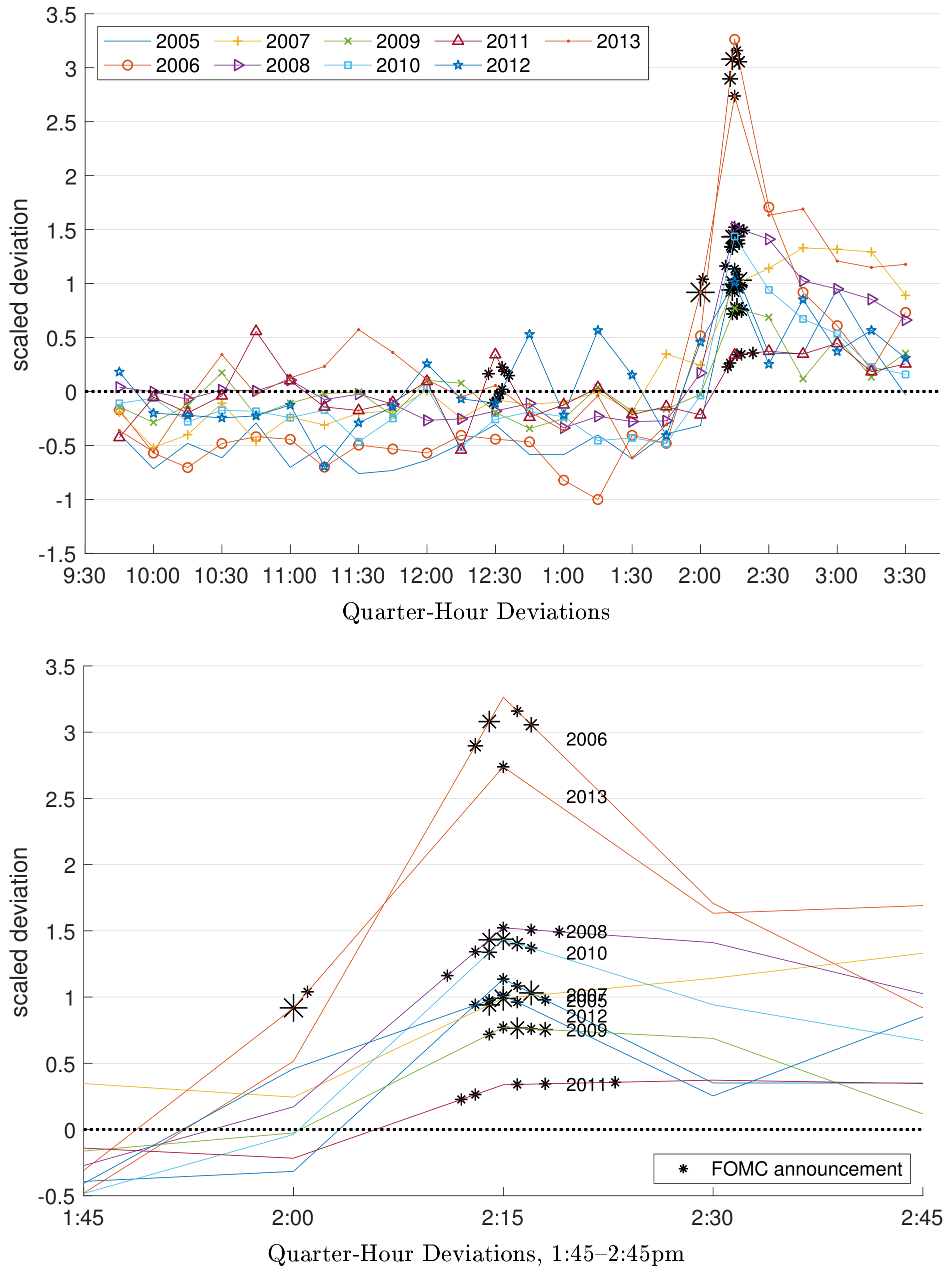

(2) fear of early information leakage and attendant price jumps, as suggested by the empirical investigations of Bernile, Hu, and Tang (2016) and Kurov, Sancetta, Strasser, and Wolfe (2017). Notably the measure does not simply reflect contemporaneous realized volatility around the FOMC announcement: Lucca and Moench (2015) instead find that volatility decreases monotonically in the hours prior to the FOMC announcement (Figure 3 of their work).

Although the tail risk measure registers increased risk in the quarter hours around FOMC news, implied tail risk is typically lower than average prior to the FOMC announcement, in parallel with the period's reduced volatility. There is little evidence that the high average returns the morning of FOMC announcements can be attributed to market jump fears. The pre-FOMC announcement drift and announcement-day success of the CAPM therefore cannot be attributed to an increase in the (physical) probability or magnitude of market jumps.

C. The 2007–2008 Financial Crisis

I study the 2007–2008 Financial Crisis to demonstrate the potential of my approach to identify factor- or sector-specific extreme-event risks. Specifically, I apply the jump extraction technique to estimate the magnitude of perceived jump risks to a "financials" factor independent from market risks. The choice of financial-sector risks is motivated by their economic importance during the 2000s as well as by the difficulty of disentangling financial sector risks from market risks using alternate methods; during this period, the rolling daily correlation of XLF, my financial-sector factor-mimicking portfolio, and the SPY often exceeds 90%. At the same time, Aït-Sahalia and Xiu (2018) demonstrate that the first two principal components of high-frequency returns correspond well with market and financial-sector innovations, respectively, suggesting that innovations in these factor risks should be detectable at high frequency.

The central regression in this analysis modifies Equation (6) to accommodate a financials factor:[23]

$$\left(\frac{Vh}{d}\right)_{it} = \tilde{\xi}_t + \xi_{t,MKT}\beta_{i,MKT} + \xi_{t,FIN}\left|\beta_{i,FIN}\right| + \epsilon_{it},\ \forall t. \tag{9}$$

[23] The Online Appendix introduces another bivariate risk model with market factor and value factor (HML) risks. Although HML explains less cross-sectional variation in returns at high frequencies, the tail risk methodology nonetheless recovers clearly interpretable HML jumps.

I exclude co-jump terms for the market and financial factors because (1) the joint risk of tail events in the market and financial factors is not of independent interest and (2) rank tests around market jumps find evidence against factor co-jumps with the aggregate market (Li, Todorov, Tauchen, and Lin (2017)). Appendix A discusses this point in additional detail.

shring Figure VIII plots the time series of implied financial sector tail risks. The recovered series of financial tail risks is visually similar to the one-factor market risks of Figure II, but it differs somewhat in the events corresponding with the largest changes in financial sector risks. Several events associated with large market risk increases in the one-factor model are in fact specific to the financial sector. Large-scale asset purchases, bank bailout legislation, and bank nationalization news feature prominently for financials, but not for the aggregate market in the two-factor model. Conversely, the FOMC interest rate target announcements of 2007–2008 and the S&P U.S. credit rating downgrade shock the aggregate market but not the financial sector separately.

I now test formally whether the recovered financials tail risks indeed correspond with jumps in the financials factor. By analogy with Table II, Table V compares medium-scale tail realizations in the XLF to ξ_{FIN} in the two-factor model of Equation (9). As before, I split specifications based on (1) the number of minutely differences of more than 10 basis points ("jump count") and the weighted sum of jumps of 5, 10, 25, and 100 basis points and (2) the number of minutely differences of more than 5 half-spreads ("jump count") and the weighted sum of jumps of 1, 5, 10, and 25 half-spreads. Rather than using the VIX and other controls, I include market jump tail risk and realizations as controls to quantify the degree to which jump types are successfully disentangled.

All specifications feature a strong relation between financial-sector tail risk and realized financial-sector jumps. A one standard deviation increase in ξ_{FIN} corresponds with 3.9 more basis-point jumps and 1.4 spread jumps on a baseline standard deviation of 10.7 and 2.5 XLF jumps per hour, respectively. Which tail risk measure is more strongly associated with financial sector jumps depends on the specification, but both ξ_{MKT} and ξ_{FIN} clearly possess explanatory power throughout. The strong relation between market tail risk and XLF jump realizations comes about because SPY and XLF often co-jump during the financial crisis, the primary source of variation in tail risk during the sample.

Figure VIII: **Largest Increases in Financial-Sector Tail Risk, 2004–2013**

This figure plots hourly estimated tail risks for financial risks (dark gray) and market risks (light gray) in a two-factor market and financials model. The fifteen largest increases in tail risks within a one-month window are overlaid with a black 'X.' Changes are measured as the tail risk at date t and hour h less the tail risk at date $t-1$ and hour h. The table below offers a brief description of coincident events on tail-risk news days. Standardized values divide by the full time-series standard deviation of changes. Bolded events correspond with extreme changes in both factors (using the two-factor model) within 24 hours. Implied tails on October 10, 2008 are truncated for visual clarity.

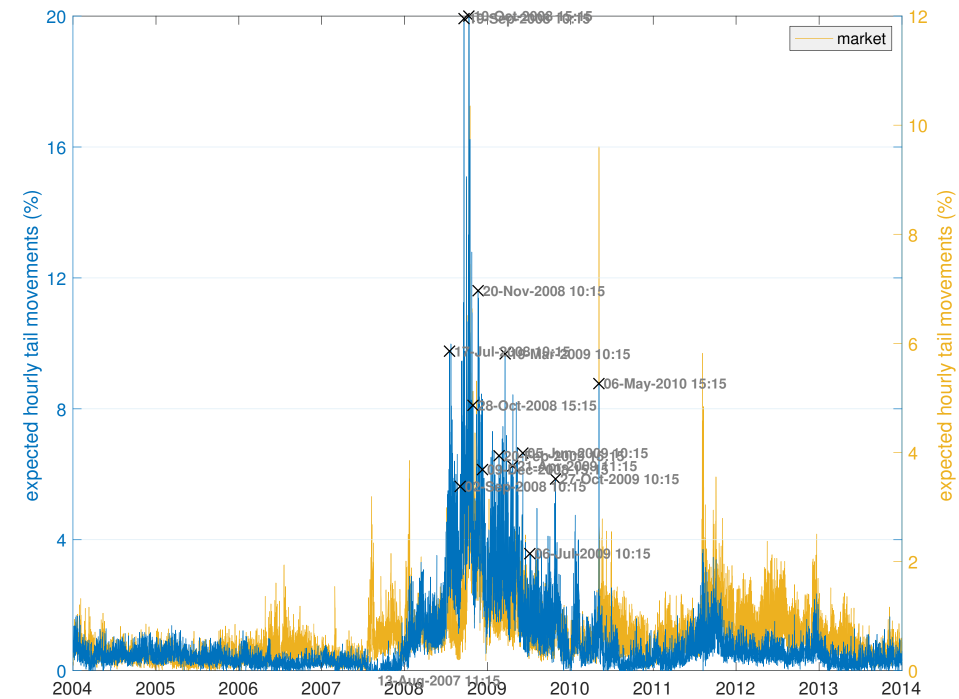

Date	Value	Event
13-Aug-07 11:15	4.13	Fed intervenes to stave off credit crisis, end "Quant Quake"
17-Jul-08 10:15	5.14	Large banks report collapsing profits and/or large losses
2-Sep-08 10:15	4.53	Korean Development Bank confirms interest in Lehman lifeline as potential partner banks express skepticism
19-Sep-08 10:15	15.73	TARP proposal becomes public; Treasury backs money-market funds
10-Oct-08 15:15	23.78	Stock market crashes in Asia, Europe, and the U.S.
28-Oct-08 15:15	6.61	First round of TARP bank bailouts ($115 billion)
20-Nov-08 10:15	9.87	Auto bailout rejected; sharp drop in financial stocks
9-Dec-08 15:15	4.35	WSJ reports first Congressional Oversight Panel review on bailouts harshly criticizes Treasury, TARP
20-Feb-09 15:15	6.38	Sen. Dodd suggests bank nationalization may be necessary
19-Mar-09 10:15	7.05	FOMC announces $1T in new bond and MBS purchases
21-Apr-09 11:15	4.85	BofA reports sharp rise in bad loans; financials down >10%
5-Jun-09 10:15	5.77	Rumors of FDIC push to increase control over Citigroup
6-Jul-09 10:15	4.17	FDIC proposes restrictions on takeovers of failed banks
27-Oct-09 10:15	4.25	House committee presents draft "Too Big to Fail" law
6-May-10 15:15	11.46	**2010 Flash Crash**

Table V: **Contemporaneous and Predictive Regressions for XLF Jump Realizations**

This table presents results from regressions of realized financial sector (XLF) jumps against estimated tail risks,

$$\begin{aligned} tail\_realization_t &= \alpha + \beta \xi_{t-\Delta,FIN} + \gamma \xi_{t-\Delta,MKT} + \delta SPY\_tail\_realization_t + \\ &\quad \alpha_{-1} tail\_realization_{t-1} + \beta_{-1} \xi_{t-1-\Delta,FIN} + \\ &\quad \gamma_{-1} \xi_{t-1-\Delta,MKT} + \delta_{-1} SPY\_tail\_realization_{t-1} + \epsilon_t, \end{aligned}$$

where $\Delta = 0$ and $\Delta = 1$ correspond to the upper and lower panels, respectively. Tail realizations are measured in counts of minutely returns exceeding basis point or spread thresholds. The count variable sums jumps exceeding 10 basis points or 5 half-spreads, and the sum variables are a weighted sum of jump sizes exceeding 5, 10, 25, and 100 basis points or 1, 5, 10, or 25 half-spreads. Regressions in the top panel consist of the 2005–2013 sample by trading hour, with one-month rolling HAC standard errors (126 observations). Regressions in the bottom panel average all variables within each year and hour of the trading day and use White standard errors. $\xi_{t,MKT}$ and $\xi_{t,FIN}$ are normalized by their standard deviations in both panels. t-statistics are in parentheses.

	XLF Basis-Point Jumps					
	Jump Count			Jump Sum		
ξ_{FIN}	3.91^{***}	1.39^{**}	0.69^{*}	71.71^{***}	28.42^{***}	17.78^{***}
	(5.43)	(2.19)	(1.88)	(8.96)	(3.32)	(4.04)
ξ_{MKT}		4.17^{***}	0.42		72.07^{***}	0.58
		(5.00)	(0.85)		(6.63)	(0.10)
SPY Jump			0.88^{***}			1.00^{***}
			(18.81)			(31.46)
Obs.	11280	11280	11280	11280	11280	11280
R^2	0.86	0.88	0.92	0.87	0.89	0.94

	XLF Spread Jumps					
	Jump Count			Jump Sum		
ξ_{FIN}	1.41^{***}	0.48	0.49	11.11^{***}	3.70^{**}	3.63^{**}
	(5.73)	(1.52)	(1.58)	(10.15)	(2.00)	(2.26)
ξ_{MKT}		1.56^{***}	1.42^{***}		12.34^{***}	3.45^{**}
		(4.67)	(4.60)		(6.04)	(2.28)
SPY Jump			0.02^{*}			0.13^{***}
			(1.68)			(10.80)
Obs.	11280	11280	11280	11280	11280	11280
R^2	0.60	0.65	0.65	0.76	0.80	0.86

*** $p < 0.01$, ** $p < 0.05$, $^{*}p < 0.1$

Table V: **Contemporaneous and Predictive Regressions for XLF Jump Realizations (Cont.)**

	XLF Basis-Point Jumps					
	Jump Count			Jump Sum		
ξ_{FIN}	5.98***	3.14***	2.17***	103.43***	53.46***	39.14***
	(8.16)	(5.07)	(4.09)	(10.90)	(6.14)	(6.18)
ξ_{MKT}		5.21***	1.60		92.19***	20.20
		(3.92)	(1.47)		(5.24)	(1.52)
SPY Jump			0.81***			0.96***
			(13.06)			(17.02)
Obs.	9016	9016	9016	9016	9016	9016
R^2	0.77	0.80	0.83	0.76	0.80	0.84

	XLF Spread Jumps					
	Jump Count			Jump Sum		
ξ_{FIN}	1.64***	0.79**	0.79**	13.83***	6.34***	5.93***
	(4.62)	(2.07)	(2.09)	(8.57)	(2.87)	(2.93)
ξ_{MKT}		1.65***	1.41***		14.21***	5.76**
		(4.25)	(3.95)		(5.05)	(2.51)
SPY Jump			0.02**			0.12***
			(2.14)			(8.43)
Obs.	9016	9016	9016	9016	9016	9016
R^2	0.41	0.46	0.46	0.60	0.65	0.68

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

To evaluate whether ξ_{FIN} picks up financial sector events that are not also aggregate market events, I add contemporaneous SPY jump counts and jump sums as an additional control. ξ_{FIN} typically remains significantly associated with XLF jumps net of market co-jumps, whereas ξ_{MKT} is often driven out. In the Online Appendix, I also extend the sample back to July 1999—the start date for the XLF plus six months to calculate betas—and I confirm that the ξ_{FIN} has significantly greater independent explanatory power in periods in which market and financial-sector dynamics are not so closely linked. These brief analyses suggest that the jump tail extraction technique separately identifies market and financial sector risks at high frequency.

Notably the Financial Crisis is characterized by large and asynchronous shocks to several economic sectors, including financials, energy, and real estate. Although a complete description of the extreme-event risks before and during the Financial Crisis is beyond the scope of this paper, one could in principle extend the factor model of Equation (9) to include sector-tracking securities for each industry, or even to decompose the market into sectors as in the Select Sector SPDR ETFs (akin to the analysis of sectoral co-jumps in Li, Todorov, and Tauchen (2017)). Importantly for this application, the omission of these other sectors from the estimation does not appear to contribute much to omitted variable bias in ξ_{FIN}: ξ_{FIN} forecasts XLF jumps well, and each of the 15 largest ξ_{FIN} changes corresponds with major financial-sector news.

VII. Conclusion

High-frequency market makers continually extract signals from order flow to optimize their provision of liquidity. Intermediaries must pay special attention to signals on potential discontinuous price movements, because such movements can generate losses from "picking off" by other fast traders. Securities with larger factor loadings are more exposed to discontinuous factor movements than are securities with smaller loadings. As a consequence, liquidity costs more for these securities, and the cross section of liquidity costs embeds significant information about near-term return factor risks. The key contribution of this paper is the development of a methodology for extracting some of this factor risk information in real time.

This methodology brings microstructure data to bear on measurement challenges in related fields. Bid-ask spreads are unique among existing data sources in their ability to reveal intraday changes in extreme-event risk for common factors in stock returns. By using these data, my approach differs from existing methods in its ability to obtain information about (1) a wide array of return factor risks (2) at an intraday frequency (3) for short look-ahead horizons.

This cross-sectional approach offers a valuable tool for researchers to evaluate extreme-event risks. High-frequency tail risk estimates provides a viable alternative to the VIX and other leading indicators of market turmoil. Regulators, too, might benefit: the dramatic rise in tail risk before the 2010 Flash Crash suggests that the measure may have predictive power for severe market disruptions.

By contrast with options-implied measures, my measure cannot anticipate tail realizations beyond a short forecasting horizon, nor does it provide direct information on the persistence or serial correlation of jump events. For these reasons, my measure is strictly speaking best interpreted as market expectations of tail realizations over short horizons. Notwithstanding market makers' limited planning horizon, however, empirical linkages between near-term factor risks and disasters of the Rietz (1988) and Barro (2006) variety are surprisingly strong. Seminal events of the Great Recession and 2007–2008 Financial Crisis manifest as large changes in implied tail risks in the immediate term, and recent work by Andersen, Fusari, and Todorov (2015) suggests that market crashes have their origins in sequences of "small" jumps explicitly captured by my methodology. Further study of these linkages is left for future investigation.

A. Derivations

Imposing the Factor Structure in Jump Returns

By assumptions 1–2, the jump intensity for stock i is given by

$$\lambda_{jump} = \sum_k \lambda_k + \tilde{\lambda}, \tag{10}$$

where λ_k is the jump intensity for factor k and $\tilde{\lambda}$ is the jump intensity for the stock's idiosyncratic component. This representation decomposes short-lived adverse-selection risks into factor- and idiosyncratic-news components.

Picking-off costs to the market maker integrate over the distribution of potential jumps larger than the half spread. By excluding co-jumps, when a factor k jumps, discontinuous returns have a simple form, $r_i^d = \beta_{ik} r_k^d$, as coincident jump returns from other sources are exactly zero. Consequently, we can sum over costs associated with each factor independently rather than integrating over a potentially complicated region associated with all potential combinations of jump returns exceeding h.

The market factor model is readily estimated in part because reliable negative betas are quite rare among common stocks. With this factor in mind, I assume for now that all betas are positive. Then substituting Equations (3) and (10) into Equation (2) delivers

$$h\lambda_{FT}\frac{q^*}{d} \quad = \quad \sum_k \lambda_k \int_{h/\beta_k}^{\infty} (\beta_k r_k - h)\, f(r_k)\, dr_k + \tilde{\lambda} \int_h^{\infty} (\tilde{r}_i - h)\, f(\tilde{r}_i)\, d\tilde{r}_i. \tag{11}$$

The salient region of the jump distribution for each asset-factor combination is determined by h/β_k. For each factor k, jump risks can be decomposed into two regions: jumps larger than the half-spread for all assets, i.e., $r_k \geq \bar{h}_k \equiv \max_i (h_i/\beta_{ik})$, and jumps larger than the half-spread for some assets but not for others:

$$\int_{h_i/\beta_k}^{\infty} (\beta_{ik} r_k - h_i)\, f(r_k)\, dr_k = \int_{h_i/\beta_{ik}}^{\bar{h}_k} (\beta_{ik} r_k - h_i)\, f(r_k)\, dr_k +$$

$$\int_{\bar{h}_k}^{\infty} (\beta_{ik} r_k - h_i) f(r_k)\, dr_k$$
$$= \beta_{ik} \underbrace{\int_{h_i/\beta_{ik}}^{\bar{h}_k} \bar{F}(r_k)\, dr_k - \bar{h}_k \bar{F}(\bar{h}_k)}_{\text{small jumps}} + \beta_{ik} \underbrace{\int_{\bar{h}_k}^{\infty} r_k f(r_k)\, dr_k}_{\text{large jumps}}, \tag{12}$$

where $\bar{F}$ denotes the counter-cumulative distribution function for r_k.

I assume that all jumps are large relative to half-spreads for my set of assets. In support of this assumption, Hendershott, Jones, and Menkveld (2011) show that by 2004, the first year in my sample, the effective half-spread for smaller-than-average (fourth size quintile) stocks is less than 6 basis points; using a β cutoff of 0.5 implies the smallest "large" jump can be less than a 12 basis point change in price, or equivalently, less than an 8–22 cent change in the price of the SPY market proxy during the time period considered. By contrast, the smallest (median) realized jumps detected in the SPY by the Lee and Mykland (2008) and Bollerslev, Todorov, and Li (2013) methodologies are 26.5 (46.6) basis points and 13.5 (32.7) basis points, respectively.[24] In empirical applications, I also impose loose restrictions on the set of assets considered to exclude stocks with extreme spreads and beta loadings to ensure that this assumption holds.[25]

Excluding "small" jumps delivers the following simplifying relation:

$$\lim_{\bar{h}_k \downarrow 0} \int_{h_i/\beta_k}^{\infty} (\beta_{ik} r_k - h_i) f(r_k)\, dr_k = -h_i + \beta_{ik} \int_0^{\infty} r_k f(r_k)\, dr_k. \tag{13}$$

Equation (11) then reduces to a linear relation between the liquidity consumer arrival rate for each asset and the distribution of jump risks for each factor multiplied by the asset's factor exposure:

$$h \lambda_{FT} \frac{q^*}{d} \overset{(3),(10)}{=} \sum_k \lambda_k \int_{h/\beta_k}^{\infty} (\beta_k r_k - h) f(r_k)\, dr_k + \tilde{\lambda} \int_h^{\infty} (\tilde{r} - h) f(\tilde{r})\, d\tilde{r} \tag{14}$$
$$\overset{(13)}{=} -\left(\tilde{\lambda} + \sum_k \lambda_k\right) h + \sum_k \beta_k \lambda_k \int_{\bar{h}_k}^{\infty} r_k f(r_k)\, dr_k + \tilde{\lambda} \int_h^{\infty} \tilde{r} f(\tilde{r})\, d\tilde{r}, \tag{15}$$

[24] Jump detection tests compare five-minute returns against estimates of local volatility. I use a 1% significance threshold in the Lee and Mykland (2008) methodology and $\tau = 4$ standard deviations in the Bollerslev, Todorov, and Li (2013) methodology. Corresponding values for the Financial Select Sector SPDR ETF (XLF) are 18.2 (60.7) and 12.5 (37.4) basis points.

[25] Empirically, I find that parameter estimates vary little with the choice of threshold $\bar{h}_k$ for market (SPY) and financial sector (XLF) test factors.

which, replacing integrals with tail expectations and moving the first term to the left-hand side, is Equation (4) up to the use of β_k rather than $|\beta_k|$,

$$(\lambda_{FT}q^* + \lambda_{jump}d)\,\frac{h}{d} = \sum_k \lambda_k E\left[r_k^d | r_k^d > \bar{h}_k\right]\beta_k + \tilde{\lambda}E\left[\tilde{r}^d | \tilde{r}^d > \bar{h}\right]. \tag{16}$$

I now revisit the possibility of negative factor betas. Doing so is important for non-market factors; for example, a jump in HML should drive "value" and "growth" stocks in opposite directions. Fortunately recovering expected jumps with non-positive betas entails minimal modification of our previous expressions. After estimating betas in the usual time-series regressions, simply define two separate "subfactors" for each factor k for which betas are sometimes negative, and replace $\beta_k r_k$ with two terms $\beta_k^+ r_k^+$ and $\beta_k^- r_k^-$, where $+$ and $-$ superscripts denote $x^+ \equiv x\mathbf{1}(x>0)$ and $x^- \equiv (-x)\,\mathbf{1}(x<0)$. Now that all subfactor betas are weakly positive we are back to the case with all $\beta_k > 0$, and the above derivation goes through unchanged to obtain

$$\begin{aligned}(\lambda_{FT}q^* + \lambda_{jump}d)\,\frac{h}{d} &= \sum_k \lambda_k^+ E\left[r_k^d | r_k^d > \bar{h}_k\right]\beta_k \mathbf{1}(\beta_k > 0) \\ &\quad + \sum_k \lambda_k^- E\left[r_k^d | r_k^d < -\bar{h}_k\right](-\beta_k)\,\mathbf{1}(\beta_k < 0) + \tilde{\lambda}E\left[\tilde{r}^d | \tilde{r}^d > \bar{h}\right].\end{aligned} \tag{17}$$

Note that by construction the subfactors for a given k cannot co-jump—we do not see a left- and right-tail realization simultaneously for a given factor.

Symmetry of left- and right-tail risks (assumption 3) implies that the expected costs $\lambda_k^+ E\left[r_k^d | r_k^d > \bar{h}_k\right]$ and $\lambda_k^- E\left[r_k^d | r_k^d < -\bar{h}_k\right]$ are the same. By this assumption Equation (17) becomes

$$(\lambda_{FT}q^* + \lambda_{jump}d)\,\frac{h}{d} \;=\; \sum_k \lambda_k E\left[r_k^d | r_k^d > \bar{h}_k\right]|\beta_k| + \tilde{\lambda}E\left[\tilde{r}^d | \tilde{r}^d > \bar{h}\right]. \tag{18}$$

Hence taking the absolute value of beta suffices to handle non-positive betas in the tail risk estimation equation (Equation (4)).

The price of sacrificing the symmetry assumption is that splitting jumps into up and down components must be matched by splits of the liquidity composite on the left-hand side of Equation (17). For example, higher depth at the best offer than at the best bid, all else equal, indicates low risk of extreme up jumps relative to extreme down jumps. Using the same liquidity composite for up and down jumps is then incorrect, and instead we would want to obtain signed versions of Vh/d for estimation. This is the approach I take in generalizing my methodology in the Online Appendix.

Allowing Co-Jumps

I now revisit assumption 2 in the context of Equations (11) and (4). Excluding co-jumps eliminates terms in Equation (11) associated with factors moving jointly. As an example, suppose that the econometrician considers only market and financial return factors and allows for co-jumps between them. Again denoting the financial return factor as FIN, the additional picking-off risk term associated with co-jumps is

$$+\lambda_{\{MKT,FIN\}} \int_R (\beta_{MKT} r_{MKT} + \beta_{FIN} r_{FIN} - h)\, d\,(r_{MKT}, r_{FIN}), \tag{19}$$

where R is the region described by $\beta_{MKT} r_{MKT} + \beta_{FIN} r_{FIN} - h \geq 0$. This additional term is readily converted into linear terms under the large-jumps assumption of the previous part if jumps are of the same sign. Under these conditions, the additional term in Equation (19) is decomposed as

$$\begin{aligned} -\lambda_{\{MKT,FIN\}} + \lambda_{\{MKT,FIN\}} E\left[r_{MKT} | r_{MKT}, r_{FIN} > 0\right] \beta_{MKT} \\ + \lambda_{\{MKT,FIN\}} E\left[r_{FIN} | r_{MKT}, r_{FIN} > 0\right] \beta_{FIN}. \end{aligned} \tag{20}$$

If jump signs differ, additional terms arise resulting from different combinations of the signs of the co-jump returns.

In my multivariate analysis of market and financial jumps, I omit these additional terms because (1) the correlation between SPY and XLF returns is positive and extremely strong, on the order of 89% in my sample, and (2) the absorption of the co-jump terms $\lambda_{\{MKT,FIN\}} \times E\left[r_{MKT} | r_{MKT}, r_{FIN} > 0\right]$ and $\lambda_{\{MKT,FIN\}} \times E\left[r_{FIN} | r_{MKT}, r_{FIN} > 0\right]$ has a clear associated economic intuition. The coefficient on β_{MKT} is the tail risk for the market factor with or without financial co-jumps, which arguably is of greater interest than either component of market factor tail risk independently.

B. Data Filters

The TAQ database occasionally records erroneous trade price and quantity information. I take three precautions to avoid contamination of the sample and mistaken detection of price jumps in Section V. First, I filter the trade price series following the trade-data methodology of Barndorff-Nielsen, Hansen, Lunde, and Shephard (2009). This methodology eliminates most obvious data errors. Second, I adapt the outlier-removal procedures of Brownlees and Gallo (2006) (similar to Rule Q4 of Barndorff-Nielsen et al. (2009)) to exclude price observations exceeding the centered median price (excluding the current observation) on $[t - 10m, t + 10m]$ by 2.5 mean absolute deviations plus a 15-basis point granularity parameter. This filter removes most data errors in the form of rapidly mean-reverting jumps in recorded prices. Finally, I use volume-weighted averages of prices within each minute as raw inputs rather than individual trades or quotes. This procedure smooths microstructure noise not of interest for my analysis.

The data is lightly filtered to exclude stocks with imprecisely estimated betas or extreme illiquidity. To be included in the sample, a stock must have:

1. Traded on at least half of the days in which the market has normal trading hours in the observation year;

2. Quoted spreads less than 5% of the price of the stock;

3. One-sided volume exceeding 200 shares in the trading interval, but less than the 95th percentile of one-sided volume; and

4. A market beta in $[0.1, 2.5]$.

Securities not satisfying all of these conditions are excluded from cross-sectional regressions.

The rationale for these filters is as follows. Filters 1 and 2 and the lower bound of filter 3 ensure that the stock is not too thinly traded to be reliable for risk estimation, either for computing betas in the time series or for estimating tail risks in the cross section. The upper bound on volume in filter 3 ensures that results are not driven by "influential" assets with extremely high volume (the distribution of volume is roughly lognormal). Filter 4 accounts for estimation error in the betas; especially large or small betas are likely to be the result of estimation error. In addition, β close to zero makes less tenable the assumption of all factor jumps being greater than h/β. Of these filters, the volume filter is the most stringent, and it eliminates 16.8% of security-hour observations; by contrast, the beta filter eliminates only 0.1% of observations. The average sample size after filtering is 2,820 distinct stocks for each trading hour.

References

Aït-Sahalia, Yacine and Dacheng Xiu. 2018. "Principal Component Analysis of High Frequency Data." *Journal of the American Statistical Association* :1–17.

Andersen, Torben G. and Tim Bollerslev. 1997. "Intraday Periodicity and Volatility Persistence in Financial Markets." *Journal of Empirical Finance* 4 (2–3):115–158.

Andersen, Torben G. and Oleg Bondarenko. 2014a. "Reflecting on the VPIN Dispute." *Journal of Financial Markets* 17:53 – 64.

———. 2014b. "VPIN and the Flash Crash." *Journal of Financial Markets* 17:1 – 46.

———. 2015. "Assessing Measures of Order Flow Toxicity and Early Warning Signals for Market Turbulence." *Review of Finance* 19 (1):1–54.

Andersen, Torben G., Nicola Fusari, and Viktor Todorov. 2015. "The Risk Premia Embedded in Index Options." *Journal of Financial Economics* 117 (3):558 – 584.

Backus, David, Mikhail Chernov, and Ian Martin. 2011. "Disasters Implied by Equity Index Options." *The Journal of Finance* 66 (6):1969–2012.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan. 2003. "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options." *Review of Financial Studies* 16 (1):101–143.

Barndorff-Nielsen, O. E., P. Reinhard Hansen, A. Lunde, and N. Shephard. 2009. "Realized Kernels in Practice: Trades and Quotes." *Econometrics Journal* 12 (3):C1–C32.

Barro, Robert J. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *The Quarterly Journal of Economics* 121 (3):823–866.

Bernile, Gennaro, Jianfeng Hu, and Yuehua Tang. 2016. "Can Information Be Locked-Up? Informed Trading Ahead of Macro-News Announcements." *Journal of Financial Economics* 121 (3):496 – 520.

Bollerslev, Tim, Sophia Zhengzi Li, and Viktor Todorov. 2016. "Roughing up Beta: Continuous vs. Discontinuous Betas, and the Cross-Section of Expected Stock Returns." *Journal of Financial Economics* 120 (3):464–490.

Bollerslev, Tim, George Tauchen, and Hao Zhou. 2009. "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies* 22 (11):4463–4492.

Bollerslev, Tim and Viktor Todorov. 2011a. "Estimation of Jump Tails." *Econometrica* 79 (6):1727–1783.

———. 2011b. "Tails, Fears, and Risk Premia." *The Journal of Finance* 66 (6):2165–2211.

———. 2014. "Time-Varying Jump Tails." *Journal of Econometrics* 183 (2):168–180.

Bollerslev, Tim, Viktor Todorov, and (Sophia) Zhengzi Li. 2013. "Jump Tails, Extreme Dependencies, and the Distribution of Stock Returns." *Journal of Econometrics* 172 (2):307–324.

Brownlees, C.T. and G.M. Gallo. 2006. "Financial Econometric Analysis at Ultra-High Frequency: Data Handling Concerns." *Computational Statistics & Data Analysis* 51 (4):2232–2245.

Brunnermeier, Markus K. and Lasse Heje Pedersen. 2009. "Market Liquidity and Funding Liquidity." *Review of Financial Studies* 22 (6):2201–2238.

Budish, Eric, Peter Cramton, and John Shim. 2015. "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response." *Quarterly Journal of Economics* 130 (4):1547–1621.

Cardella, Laura, Jia Hao, and Ivalina Kalcheva. 2017. "Liquidity-based Trading Fees and Exchange Volume." Working paper.

Carr, Peter and Liuren Wu. 2003. "What Type of Process Underlies Options? A Simple Robust Test." *The Journal of Finance* 58 (6):2581–2610.

CFTC and SEC. 2010. "Findings Regarding the Market Events of May 6, 2010." Report of the staffs of the cftc and sec to the joint advisory committee on emerging regulatory issues.

Chinco, Alexander, Adam D. Clark-Joseph, and Mao Ye. 2017. “Sparse Signals in the Cross-Section of Returns.” *Journal of Finance* Forthcoming.

Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam. 2000. “Commonality in Liquidity.” *Journal of Financial Economics* 56 (1):3–28.

Christie, William G. and Paul H. Schultz. 1994. “Why do NASDAQ Market Makers Avoid Odd-Eighth Quotes?” *The Journal of Finance* 49 (5):1813–1840.

Comerton-Forde, Carole, Terrence Hendershott, Charles M. Jones, Pamela C. Moulton, and Mark S. Seasholes. 2010. “Time Variation in Liquidity: The Role of Market-Maker Inventories and Revenues.” *The Journal of Finance* 65 (1):295–331.

Copeland, Thomas E. and Dan Galai. 1983. “Information Effects on the Bid-Ask Spread.” *Journal of Finance* 38 (5):1457–1469.

Cremers, K. J. Martijn and Jianping Mei. 2007. “Turning over Turnover.” *Review of Financial Studies* 20 (6):1749–1782.

Easley, David, Marcos Lopez de Prado, and Maureen O’Hara. 2016. “Discerning Information from Trade Data.” *Journal of Financial Economics* 120 (2):269 – 285.

Easley, David, Nicholas M. Kiefer, Maureen O’Hara, and Joseph B. Paperman. 1996. “Liquidity, Information, and Infrequently Traded Stocks.” *The Journal of Finance* 51 (4):1405–1436.

Easley, David, Marcos M. López de Prado, and Maureen O’Hara. 2012. “Flow Toxicity and Liquidity in a High-frequency World.” *Review of Financial Studies* 25 (5):1457–1493.

Easley, David and Maureen O’Hara. 1992. “Time and the Process of Security Price Adjustment.” *The Journal of Finance* 47 (2):577–605.

Fama, Eugene F. and James D. MacBeth. 1973. “Risk, Return, and Equilibrium: Empirical Tests.” *Journal of Political Economy* 81 (3):607–636.

Fleming, Michael J. and Monika Piazzesi. 2005. “Monetary Policy Tick-by-Tick.” Working paper.

Foucault, Thierry, Ohad Kadan, and Eugene Kandel. 2013. "Liquidity Cycles and Make/Take Fees in Electronic Markets." *Journal of Finance* 68 (1):299–341.

Foucault, Thierry, Sophie Moinas, and Erik Theissen. 2007. "Does Anonymity Matter in Electronic Limit Order Markets?" *Review of Financial Studies* 20 (5):1707–1747.

Foucault, Thierry, Ailsa Röell, and Patrik Sandås. 2003. "Market Making with Costly Monitoring: An Analysis of the SOES Controversy." *Review of Financial Studies* 16 (2):345–384.

Glosten, Lawrence R. 1987. "Components of the Bid-Ask Spread and the Statistical Properties of Transaction Prices." *The Journal of Finance* 42 (5):1293–1307.

Gorton, Gary B. and George G. Pennacchi. 1993. "Security Baskets and Index-Linked Securities." *The Journal of Business* 66 (1):1–27.

Han, Jungsuk, Mariana Khapko, and Albert S Kyle. 2014. "Liquidity with High-Frequency Market Making." *Swedish House of Finance Research Paper* (14-06).

Harris, Jeffrey H. and Paul H. Schultz. 1997. "The Importance of Firm Quotes and Rapid Executions: Evidence from the January 1994 SOES Rules Changes." *Journal of Financial Economics* 45 (1):135 – 166.

Hasbrouck, Joel and Duane J. Seppi. 2001. "Common Factors in Prices, Order Flows, and Liquidity." *Journal of Financial Economics* 59 (3):383 – 411.

Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld. 2011. "Does Algorithmic Trading Improve Liquidity?" *Journal of Finance* 66 (1):1–33.

Herskovic, Bernard, Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh. 2016. "The Common Factor in Idiosyncratic Volatility: Quantitative Asset Pricing Implications." *Journal of Financial Economics* 119 (2):249 – 283.

Holden, Craig W. and Stacey Jacobsen. 2014. "Liquidity Measurement Problems in Fast, Competitive Markets: Expensive and Cheap Solutions." *The Journal of Finance* 69 (4):1747–1785.

Jiang, George J., Ingrid Lo, and Adrien Verdelhan. 2011. "Information Shocks, Liquidity Shocks, Jumps, and Price Discovery: Evidence from the U.S. Treasury Market." *Journal of Financial and Quantitative Analysis* 46:527–551.

Kelly, Bryan and Hao Jiang. 2014. "Tail Risk and Asset Prices." *Review of Financial Studies* 27 (10):2841–2871.

Kirilenko, Andrei, Albert S. Kyle, Samadi Mehrdad, and Tugkan Tuzun. 2017. "The Flash Crash: High-Frequency Trading in an Electronic Market." *The Journal of Finance* 72 (3):967–998.

Koenker, Roger. 2005. *Quantile Regression.* Cambridge University Press.

Korajczyk, Robert A. and Ronnie Sadka. 2008. "Pricing the Commonality Across Alternative Measures of Liquidity." *Journal of Financial Economics* 87 (1):45 – 72.

Kurov, Alexander, Alessio Sancetta, Georg Strasser, and Marketa Halova Wolfe. 2017. "Price Drift Before U.S. Macroeconomic News." Working paper.

Lee, Charles M. C. and Mark J. Ready. 1991. "Inferring Trade Direction from Intraday Data." *The Journal of Finance* 46 (2):733–746.

Lee, Suzanne S. and Per A. Mykland. 2008. "Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics." *Review of Financial Studies* 21 (6):2535–2563.

Li, Jia, Viktor Todorov, and George Tauchen. 2017. "Jump Regressions." *Econometrica* 85 (1):173–195.

Li, Jia, Viktor Todorov, George Tauchen, and Huidi Lin. 2017. "Rank Tests at Jump Events." *Journal of Business & Economic Statistics* :1–10.

Lo, Andrew W. and Jiang Wang. 2000. "Trading Volume: Definitions, Data Analysis, and Implications of Portfolio Theory." *Review of Financial Studies* 13 (2):257–300.

Lucca, David O. and Emanuel Moench. 2015. "The Pre-FOMC Announcement Drift." *The Journal of Finance* 70 (1):329–371.

Mancini, Cecilia. 2009. "Non-Parametric Threshold Estimation for Models with Stochastic Diffusion Coefficient and Jumps." *Scandinavian Journal of Statistics* 36 (2):270–296.

Menkveld, Albert J. and Bart Zhou Yueshen. 2017. "The Flash Crash: A Cautionary Tale about Highly Fragmented Markets." *Management Science* Forthcoming.

Nagel, Stefan. 2012. "Evaporating Liquidity." *Review of Financial Studies* 25 (7):2005–2039.

Pelger, Markus. 2017. "Understanding Systematic Risk: A High-Frequency Approach." Working paper.

Rietz, Thomas A. 1988. "The Equity Risk Premium: A Solution." *Journal of Monetary Economics* 22 (1):117 – 131.

Savor, Pavel and Mungo Wilson. 2013. "How Much Do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements." *Journal of Financial and Quantitative Analysis* 48 (4):343–375.

———. 2014. "Asset Pricing: A Tale of Two Days." *Journal of Financial Economics* 113 (2):171 – 201.

Shanken, Jay. 1992. "On the Estimation of Beta-Pricing Models." *Review of Financial Studies* 5 (1):1–33.

Todorov, Viktor and Tim Bollerslev. 2010. "Jumps and Betas: A New Framework for Disentangling and Estimating Systematic Risks." *Journal of Econometrics* 157 (2):220 – 235.

Yan, Yuxing and Shaojun Zhang. 2014. "Quality of PIN Estimates and the PIN-Return Relationship." *Journal of Banking & Finance* 43 (0):137 – 149.

Yueshen, Bart Zhou. 2014. "Queuing Uncertainty in Limit Order Market." Working paper.