UNIL | Université de Lausanne
Département de finance

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

FINANCE RESEARCH SEMINAR SUPPORTED BY UNIGESTION

"Measuring Tail Risks in Real Time"

Prof. Brian WELLER
Northwestern University, Kellogg School of Management

Abstract

I develop a new methodology for measuring expected tail risks using the cross section of bidask spreads. Market makers embed tail risk information into spreads, because (1) middlemen lose to arbitrageurs when sharp price movements exceed the cost of liquidity and (2) price movements and potential costs are linear in factor loadings. Using this insight, simple crosssectional regressions relating spreads and trading volume to factor betas can recover factor tail risks in real time for priced or non-priced return factors. The recovered time series of implied market risks aligns closely with both realized market jumps and the VIX. In addition, the methodology quantifies a sharp, temporary increase in market tail risk before and throughout the 2010 Flash Crash; anticipates jump risks associated with Federal Open Market Committee announcements; and disentangles financial and aggregate market risks during the 2007–2008 Financial Crisis.

Friday, October 30, 2015, 10:30-12:00
Room 126, Extranef building at the University of Lausanne

Measuring Tail Risks in Real Time

Brian Weller*

Northwestern Kellogg

October 16, 2015

Abstract

I develop a new methodology for measuring expected tail risks using the cross section of bid-ask spreads. Market makers embed tail risk information into spreads, because (1) middlemen lose to arbitrageurs when sharp price movements exceed the cost of liquidity and (2) price movements and potential costs are linear in factor loadings. Using this insight, simple cross-sectional regressions relating spreads and trading volume to factor betas can recover factor tail risks in real time for priced or non-priced return factors. The recovered time series of implied market risks aligns closely with both realized market jumps and the VIX. In addition, the methodology quantifies a sharp, temporary increase in market tail risk before and throughout the 2010 Flash Crash; anticipates jump risks associated with Federal Open Market Committee announcements; and disentangles financial and aggregate market risks during the 2007–2008 Financial Crisis.

JEL: C58, G01, G12, G14, G17

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

*Kellogg School of Management, Northwestern University. b-weller@kellogg.northwestern.edu. Tel: +1 847 467 6176. Website: http://www.kellogg.northwestern.edu/Faculty/Directory/Weller_Brian.aspx. I thank Torben Andersen, Snehal Banerjee, Eric Budish, Anna Cieslak, Peter Cramton, Nicolas Crouzet, Ian Dew-Becker, Thierry Foucault, Arie Gozluklu (discussant), Ravi Jagannathan, Thomas Katzschner (discussant), Robert Korajczyk, Konstantin Milbradt, Marina Niessner, Dimitris Papanikolaou, John Shim, Viktor Todorov, and seminar participants at Northwestern Kellogg, HEC Paris, the 4th Industrial Organisation of Securities and Derivatives Markets Conference, and the 11th Annual Central Bank Conference on the Microstructure of Financial Markets for valuable comments and discussion.

I. Introduction

What does the behavior of market makers reveal about asset-pricing risks? In this paper, I develop a technique for recovering anticipated priced and non-priced risks from the cross section of bid-ask spreads. Rather than impeding the recovery of information about factor risks, spreads themselves embed a wealth of unexploited information about future price movements.

The market-making sector provides a natural setting for recovering conditional assessments of anticipated risks.[1] Unlike trades, which occur relatively infrequently, the best available quotes represent continuous, binding prices. High-frequency market makers continually update their orders so that the prevailing spread reflects available information on a nearly instantaneous basis for every exchange-traded security. Indeed, partly as a result of market-maker activity, the ratio of orders to trades (cancellations to trades) typically exceeds 30 (20) for U.S. stocks and 500 (90) for U.S. exchange-traded products.[2]

Liquidity suppliers at the best bid and offer collect and process price-relevant information quickly and effectively to avoid trading at unfavorable prices. Costly "picking off" occurs when the fundamental value of the asset jumps outside the bounds of the bid-ask spread, and the market maker fails to adjust her quote in response to the new information before another party trades against it. Overlooking tradable signals thus facilitates the picking off of stale quotes by better-informed traders at the market maker's expense. High-frequency liquidity providers therefore must be especially attuned to drivers of sharp price changes, even if such information is obtained only indirectly by divining from others' activities in the limit order book.

I exploit the frequent quote updates and forward-looking nature of modern market makers to extract real-time assessments of factor risks. The set of anticipated risks encompasses potential price movements of a few basis points, on the order of the median half-spread, to extreme price jumps, realized, for example, in the January 2015 scrap of the Swiss franc to Euro peg (+30% in CHF-EUR in minutes) and the December 2014 collapse of the Russian ruble (-13% in RUB-EUR within an hour).

Such tail realizations are particularly damaging to market makers because they are difficult to hedge and translate into losses with high probability, but they are also of broader concern to market participants and policymakers because they occur too quickly to be managed well. As a leading example in the United States, considered in Section VI, the 2010 Flash Crash sent shock waves from index futures to equities with potentially large distributional effects. 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.

[1] I use the term "market maker" to encompass all liquidity providers rather than designated market makers alone.

[2] Market Information Data Analytics System, http://www.sec.gov/marketstructure/datavis/ma_overview.html. Orders to trades are measured in share volume, and cancellations to trades are measured in counts of cancellations and trades. Exchange-traded products are defined as CRSP securities with share code 73 and primarily consist of exchange-traded funds.

My measure helps to fill this gap by providing intraday assessments of jump tail risks. More generally, this paper demonstrates that information recovered from market liquidity providers offers a rich new resource for understanding aggregate economic shocks and potential systemic threats.

Methodology Intuition

The intuition for my approach can be conveyed with a simple empirical system in the style of Fama and MacBeth (1973) regressions for estimating factor prices. The econometrician first recovers (rolling) factor exposures β_{ik} for each asset i and return factor k via time-series regressions:

$$r_{it} = \alpha_i + \sum_k \beta_{ik}^{(t)} f_{kt} + \epsilon_{it}, \ \forall i. \tag{1}$$

These estimated betas enter the second-stage cross-sectional regressions under the assumption that betas estimated from the longer time series apply for each subinterval over which cross-sectional slopes are estimated. For real-time assessment of risks, these intervals can be made quite small, on the order of hours or minutes. I modify the second-stage regression by replacing returns with the average effective half spread h_{it} (in percent per share) multiplied by volume (in shares), V_{it}:[3]

$$\overline{h_{it} V_{it}} = \gamma_t + \sum_k \xi_{kt} \beta_{ik} + \delta_{it}, \ \forall t. \tag{2}$$

Equation (2) relates average intermediation revenues to the asset's factor exposures β. ξ_t loses its interpretation as a price of factor risk in this context, but in its place, ξ_t gains the interpretation of anticipated factor tail risks. To see why, consider the optimization problem of competitive liquidity suppliers in the presence of picking-off risk.[4] In equilibrium, the half-spread is set such that expected gains per unit time from intermediating to liquidity consumers, $E\left[h_{it} V_{it}\right]$, exactly offset the expected cost per unit time from picking-off risks. Selecting a spread lower than this level does not recoup expected losses from picking off of quotes at stale prices, whereas selecting a spread higher than this level results in undercutting by other market makers.

If market makers are sufficiently fast, picking-off risk arises exclusively from factor or idiosyncratic jumps, and larger factor exposures β_{ik} generate greater picking-off risks for a given factor k. For example, a stock with a beta of zero with respect to the market is unaffected by market jumps, whereas a stock with a beta of two responds quite strongly to market movements. I show formally that the relation between the required compensation hV and betas is indeed linear in anticipated jump intensity and tail size for all but the smallest jumps. It follows that the cross-sectional es-

[3]Notably, the cross section of spreads alone is insufficient to extract factor risk information. The Internet Appendix considers spreads as the dependent variable and shows that recovered quantities are often negative, i.e., higher beta stocks have lower spreads.

[4]To facilitate exposition, I set aside issues of formal grounding for the main text and assume that picking-off risk is the sole source of the bid-ask spread.

timate for ξ_{kt} represents the tail expectation of the distribution of potential jumps for factor k at date t, or heuristically, the tail risk for each factor.

An important difference between Equation (2) and Fama-MacBeth cross-sectional regressions is that the researcher observes and estimates expectational variables for tail risks rather than stochastic tail realizations. Like options prices, which represent a risk-neutral expected payoff, spreads represent expectations of potential tail risks, and the cross section of spreads delivers an *expected* cost of jump realizations for each time interval rather than the realized cost. Consequently, coefficient estimates for anticipated tail risks ξ_{kt} are much more precise than analogous risk prices in a Fama-MacBeth framework, and cross-sectional slopes are of independent interest as conditional expectations for jump tails rather than inputs for a single average jump tail risk estimate.[5]

Applications

Figure I illustrates the output of the methodology under a market-factor model of returns. Figure I plots ξ_{mt} for each trading day in 2004–2013 with hourly estimates of the cross-sectional slopes. For each hour, I recover slopes from a filtered sample consisting of approximately 2,800 stocks to identify potential tail variation, resulting in narrow confidence intervals for anticipated tails (dashed blue). This tight estimation at high frequency distinguishes my approach from time-series methods that rely on high-frequency data series on the order of weeks or months.

The recovered 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%,[6] and a one standard deviation increase in the jump tail measure is associated with 5.47 more jumps per hour exceeding 10 basis points (t-statistic of 18.54). In addition, the extracted tail-risk measure correlates positively with the CBOE S&P 500 implied volatility index (VIX), a 30-day forward volatility measure, and realized return variation on the SPDR S&P 500 ETF (SPY), in part because both series feature significant comovement with realized jump intensities.

In addition to serving as a real-time barometer of market-factor risks, the measure performs well across diverse and challenging economic environments. I apply the methodology to the May 6, 2010 Flash Crash as a prototypical large and plausibly unexpected systematic jump. Existing tail estimation techniques do not have sufficient resolution to anticipate the Flash Crash or to reliably distinguish changes in tail risk ex post. My tail risk measure is a natural leading indicator for liquidity crashes because it draws directly from liquidity providers' revealed expectations for jump tail events. Market jump risk is elevated in the quarter-hour prior to the crash at 17 standard deviations above the previous day's tail risk, and it reaches 104 standard deviations above the

[5] This expectational interpretation of spreads comes at the cost of requiring a model of market-maker behavior. I exposit one motivating model in Section III, although this model is not unique in delivering the key equilibrium condition under which Equation (2) holds. I verify some of the underlying assumptions of this model in Section V.

[6] Although my approach and options-based techniques both measure jump tails, my measure operates in the physical measure rather than the risk-neutral measure. I compare theoretical properties of jump tail estimation methods in Section II and empirical properties in Sections V and VI.

previous day's tail risk at the height of the crash.[7] By contrast, implied idiosyncratic tail risk corresponding with level changes in spreads increases 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 SPY market index.

The Flash Crash also serves as an example for 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 after the ensuing weekend. This analysis serves as a first step toward an assessment of permanent impacts of liquidity-driven market meltdowns.

I next exploit the methodology's new intraday resolution on conditional tails 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.

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 evolution of aggregate market and financial risks during the 2007–2008 Financial Crisis using a two-factor model with market and financial sector risks. Despite the Financial Select Sector SPDR ETF (XLF) having an annual average daily correlation of 89% with the SPY over this period,[8] differences in factor loadings in the cross section nonetheless produce tight estimates for anticipated shocks specific to the financial sector. The most extreme changes in this series often differ from those of the market jump series and correspond with 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 sector risks.

[7]Incidentally, the only other occasions that register at least a 16 standard deviation increase in tail risk relative to the previous day correspond with: the largest stock market decline in four years and the rollout of the NYSE's Phase IV Hybrid Market (February 27, 2007); the market plunge at the height of the "Quant Quake" of August 2007 (August 9, 2007); the U.S. House's rejection of Paulson's financial stabilization plan (September 29, 2008); and the S&P downgrade of the U.S. federal government credit rating (August 5, 2011).

[8]The Financial Select Sector SPDR (XLF) is one of nine partitioning sector-specific ETFs associated with the S&P 500 index. Additional details on the XLF can be found at http://www.sectorspdr.com/sectorspdr/sector/xlf.

Outline

The paper proceeds as follows. Section II discusses related literature. Section III develops the equilibrium relation between bid-ask spreads and multifactor tail risks. Section IV describes data sources and empirical implementation. Section V obtains results from the market-factor model of jump risks and addresses potential confounds from other sources of the bid-ask spread. Section VI applies the model to the 2010 Flash Crash, Federal Open Market Committee announcements, and financial-sector risks during the 2007–2008 Financial Crisis. Section VII concludes.

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. 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), who use S&P 500 index options to recover *time-varying* jump tails. 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. Bollerslev and Todorov (2011a,b) instead exploit high-frequency data and extreme value in-fill arguments to estimate jump tails for the SPY market proxy.

Similarly to Kelly and Jiang (2014), this paper takes a cross-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 spreads rather than on tail return realizations. Because every bid-ask spread is informative at all times rather than only in the "rare event" states associated with return jumps, I significantly increase the conditioning frequency at which tails can be constructed. Secondly, my measure recovers tail expectations, which jointly summarizes 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.

I view my approach as complementary to 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,[9] 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 only recover 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. I elaborate on the relative applicability of these approaches when considering the 2007–2008 Financial Crisis, FOMC announcement, and 2010 Flash Crash applications of Section VI.

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. Budish et al. (2015) represent high-frequency market making as a story of two speed races: a race to be first in the order book to provide liquidity to uninformed traders, and a race to be first to modify orders in response to discontinuous changes or "jumps" in fundamental asset values. If the marginal liquidity provider(s) loses the second race, she offers intermediation services at stale prices and suffers losses when trades are executed at these old prices by other fast traders. These potential losses are "picking-off risk," an important source of the bid-ask spread. 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 et al. (2015) and this study can also be motivated using the quotes-as-options framework of Copeland and Galai (1983).

Budish et al. (2015) embed a compound jump process as the source of time variation in prices. 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. In so doing, I convert Budish et al. (2015)'s statement about spreads in continuous-time double auctions into a useful empirical relation between spreads and asset-pricing risks. Section III develops this relation in depth and consider several additional issues that arise when inverting the model to recover underlying jump risks.

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 the pre-HFT limit order book contains volatility information in addition to directional information, albeit at the individual asset level. Nagel (2012) shares the spirit of this paper in relating returns to intermediation to forward-looking *market* volatility. Specifically, Nagel (2012) shows that short-term reversal returns are very highly correlated with the VIX, and he interprets this relation as evidence that financially constrained intermediaries are less able to provide liquidity in times of financial uncertainty. I take

[9] Even Carr and Wu (2003) filters 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.

a different approach and instead derive a similar relationship as a robust consequence of picking-off risk associated with movements in the market factor. Jump risks in the market factor must be compensated in equilibrium by higher returns to market making for assets with greater market exposure. When anticipated jump risks—a component of the VIX—are greater, the level of equilibrium returns is higher and the slope with respect to factor loadings is steeper. I extend this intuition to a broad set of factor risks and show that bid-ask spreads embed much richer information about the underlying factor structure of realized asset returns. In this sense, my paper also relates to the broad literature on common factors in liquidity (e.g., Chordia, Roll and Subrahmanyam (2000), Hasbrouck and Seppi (2001), and Korajczyk and Sadka (2008)).

III. Spreads and Asset-Pricing Risks

In this section, I develop the relation between the cross section of bid-ask spreads and instantaneous factor tail risks. I first build intuition for the relationship between bid-ask spreads and jump factor exposure for individual assets. Competitive market makers set the spread to balance gains from liquidity supply against potential losses from picking off of stale quotes by other fast traders. This trade-off underlies the interpretation of recovered empirical quantities as jump tail risks. I then map theory to observable variables used in estimation. In particular, I address complicating issues relating to the arrival rate of liquidity consumers, non-unit liquidity demands, and discreteness in spreads. The third part of this section develops the two-stage cross-sectional approach for inverting the distribution of bid-ask spreads to recover factor information. In this part, I develop the empirical technique to recover information about asset prices from analysis of realized liquidity provision across securities.

A. Picking-Off Risks and Return Tails

The equilibrium bid-ask spread of a competitive market-making sector is set by balancing expected gains per unit time from liquidity provision against expected costs per unit time from displaying potentially stale quotes. I adapt the picking-off risk model of Budish et al. (2015) as formal motivation for this key condition. However, this condition also can be derived in alternative setups under different assumptions, e.g., those of Copeland and Galai (1983), so empirical results should not be viewed as requiring the Budish et al. (2015) framework in order to hold. In either case, I assume for now that picking-off risk is the sole source of cross-sectional variation in spreads.[10]

[10] I postpone discussion of this assumption until Section V, because assessing potential contamination by inventory and adverse selection risk is an empirical question that requires comparing estimates under different specifications. In fact, I require only that these sources of the spread not be strongly cross-sectionally correlated with factor exposures, but I make the stronger assumption for now to ease exposition of the model.

Jump Risks and Equilibrium Spreads

I now briefly review the basics of a lightly adapted Budish et al. (2015) model to provide a foundation for an enriched variant that can be taken to the data. The investible universe consists of a single security i. The fundamental value of asset i evolves as a compound Poisson process with arrival-rate parameter λ_{jump} and jump distribution $F(J)$. All market participants observe a public news source on the value of i. By contrast with continuous variation in prices, jumps represent large, discontinuous changes in asset value over an infinitesimal interval. For example, Federal Open Market Committee statement releases affect discount rates and are embargoed until a rigidly enforced dissemination time, and firm earnings are packaged and distributed to subscribers within milliseconds of release by Thomson Reuters and RavenPack.[11]

The economy consists of two types of traders: uninformed liquidity consumers with unit demands and a continuum of potential market makers, each capable of offering one unit of liquidity. Market makers "race" to enter the limit-order book to intermediate for the liquidity consumer at a price h, the half-spread. h arises from price competition among the competitive market makers. The order book respects price-time priority, although queuing order is random from the perspective of market makers. Potential market makers that lose the speed race to intermediate at h instead participate as stale-quote "snipers"—if the underlying value of the asset jumps outside of the NBBO, the fast non-market makers consume liquidity for a profit of $J - h$ at the offering market maker's expense.

The timeline of events is as follows. At time 0, the first "race" commences among potential market makers of equal speed. The winner is selected at random, e.g., by noise in the exchange clock. Once the book is filled, the economy advances until either of two events occurs to force an update to the limit order book. One possibility is that an information event for i creates a jump in its price. The fringe of stale quote snipers picks off the resting market maker with probability 1 if the jump exceeds the half-spread.[12] In either case, a subsequent order conditioning on the new information establishes a new bid and offer. Alternatively, a fundamental trader arrives and takes available liquidity from the order book. In this scenario the market maker earns the half-spread h, and the set of potential market makers again race to refill the order book centered at the same price as before. Because refreshing the book occurs in both event categories, the market making sector's optimal policy for liquidity provision entails providing exactly one unit of liquidity and replenishing if and only if liquidity is taken.

Figure II illustrates the sources of market-maker profits and losses per unit of offered liquidity under this model of picking-off risk. The rate of gains from intermediation equals the half-spread

[11] The RavenPack News Analytics service provides "analytics on more than 170,000 entities in over 100 countries and covers over 98% of the investable global market." As of this writing, associated documentation claims that the average overall latency from story publication to dissemination of machine-readable sentiment analysis is 250 milliseconds. By comparison, an average human blink takes 300–400 milliseconds.

[12] The assumption of picking off with certainty is a simplifying assumption only. The Budish et al. (2015) model obtains an identical equilibrium condition absent this assumption by making the participation of HFTs endogenous.

h multiplied by the arrival rate of liquidity consumers λ_{FT}. The probability of this event for a particular unit of liquidity is $\lambda_{FT}/(\lambda_{FT}+\lambda_{jump})$. The expected costs from being picked off equal the arrival rate of jumps in the fundamental value of the asset λ_{jump} multiplied by the tail expectation of jump sizes J above the spread h, as picking off is desirable only when the price movement exceeds the half-spread. Equating these quantities obtains the equilibrium condition of Budish et al. (2015):

$$\underbrace{\lambda_{FT} \times h}_{E\left[\text{benefit/time}\right]} \quad = \quad \underbrace{\lambda_{jump} \times \Pr\left(J > h\right) \times E\left[J - h | J > h\right]}_{E\left[\text{cost/time}\right]}. \tag{3}$$

Equation (3) relates anticipated price movements J to the half-spread h. Only price movements larger than h result in costs to the market maker in trading at unfavorable prices. Higher arrival rates of liquidity consumers drive h toward zero, whereas faster information arrivals or larger jumps conditional on information arrivals increase h. Although I motivate the empirical setup with the Budish et al. (2015) model, the trade-off of Equation (3) applies more generally to the class of models in which potentially stale quotes motivate the bid-ask spread.

Relating Picking-Off Risks to Return Tails

I overlay a factor structure on the simple jump process of Budish et al. (2015) (see, e.g., Todorov and Bollerslev (2010)). Jump or discontinuous returns are decomposed as

$$r_i^d = \sum_k \beta_{ik} r_k^d + \tilde{r}_i^d, \tag{4}$$

for a set of return factors k and idiosyncratic jump return $\tilde{r}^d$. Recent empirical work suggests that high-frequency traders quickly embed factor news into prices of individual stocks to enforce the factor structure of Equation (4). Brogaard et al. (2014) show that HFT order flow in individual stocks reflects market-factor information, and Ito and Yamada (2015) go further to claim that "HFTs are equipped to aggregate market-wide information and reflect [it] in prices [...] This makes their order flows informative for *market-wide* information" (my emphasis).

Because the factor structure relates returns in percentage terms rather than in dollar terms, I rewrite Equation (3) in return form as

$$\tilde{h}_i \lambda_{i,FT} = \lambda_{i,jump} \times \int_{\tilde{h}}^{\infty} \left(r_i^d - \tilde{h}_i\right) f\left(r_i^d\right) dr_i^d, \tag{5}$$

where Equation (5) converts all dollar quantities into returns by replacing dollar half-spreads h by percentage half-spreads $\tilde{h}$ and dollar jump sizes by discontinuous returns r^d. To economize on notation, I suppress asset subscripts, jump superscripts, and decorative marks from now on unless meanings are otherwise ambiguous.

At this point I make two simplifying assumptions 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.*

Both assumptions streamline the estimation procedure 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 the cross-sectional estimation, as described in Appendix C. Assumption 2 excludes heterogeneity in the rate of idiosyncratic information arrival among assets. This assumption, too, may be relaxed by controlling for realized idiosyncratic volatility or by including proxies for anticipated adverse selection at the stock level.

Assuming Poisson arrivals for jumps and assumptions 1–2, the jump intensity for stock i is given by

$$\lambda_{jump} = \sum_k \lambda_k + \tilde{\lambda}, \tag{6}$$

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 across the joint distribution of potential jumps larger than the half spread. By excluding co-jumps, discontinuous returns have a simple form, $r_i^d = \beta_{ik} r_k^d$, for the jumping factor k, as coincident jump returns in other factors are exactly zero. Consequently, we can sum over costs associated with each factor independently rather than integrating over the region associated with all potential combinations of jump returns. Substituting Equations (4) and (6) into Equation (5) delivers:

$$h\lambda_{FT} \;=\; \sum_k \lambda_k \int_{h/\beta_k}^{\infty} \left(\beta_k r_k - h\right) f\left(r_k\right) dr_k + \tilde{\lambda} \int_h^{\infty} \left(\tilde{r}_i - h\right) f\left(\tilde{r}_i\right) d\tilde{r}_i. \tag{7}$$

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 these two regions: jump sizes that exceed the spread for all assets, i.e., $r_k \geq \bar{h}_k \equiv \max_i \left(h_i/\beta_{ik}\right)$, and jump sizes that exceed the spread for some assets but not for others:

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

$$= \quad \underbrace{\beta_{ik}\int_{h_i/\beta_{ik}}^{\bar{h}_k} \bar{F}(r_k)\,dr_k - \bar{h}_k\bar{F}(\bar{h}_k)}_{\text{small jumps}} + \underbrace{\beta_{ik}\int_{\bar{h}_k}^{\infty} r_k f(r_k)\,dr_k}_{\text{large jumps}}, \tag{8}$$

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

I assume that all jumps are large relative to spreads for my set of assets. Hendershott et al. (2011) show that by 2005, the effective half-spread for smaller-than-average (fourth size quartile) stocks is less than 5 basis points; using a β cutoff of 0.5 implies the smallest "large" jump can be less than a tenth of a percent change in price, or equivalently, less than an 7–18 cent change in the price of the SPY market proxy during the time period considered. In estimation, I also impose loose restrictions on the set of assets considered to exclude stocks with extreme spreads and beta loadings.[13] Assigning zero density to small jumps delivers the following simplifying relation:

$$\lim_{\bar{h}_k\downarrow 0}\int_{h_i/\beta_k}^{\infty} (\beta_{ik}r_k - h_i)\, f(r_k)\,dr_k = -h_i + \beta_{ik}\int_0^{\infty} r_k f(r_k)\,dr_k. \tag{9}$$

Equation (7) then reduces to a linear relation between the liquidity consumer arrival rate and the distribution of jump risks for each factor.

$$\begin{aligned} h\lambda_{FT} &\overset{(4),(6)}{=} \sum_k \lambda_k \int_{h/\beta_k}^{\infty} (\beta_k r_k - h)\, f(r_k)\,dr_k + \tilde{\lambda}\int_h^{\infty} (\tilde{r} - h)\, f(\tilde{r})\,d\tilde{r} \\ &\overset{(9)}{=} -\left(\tilde{\lambda} + \sum_k \lambda_k\right) h + \sum_k \beta_k\lambda_k \int_{\bar{h}_k}^{\infty} r_k f(r_k)\,dr_k + \tilde{\lambda}\int_h^{\infty} \tilde{r} f(\tilde{r})\,d\tilde{r} \\ &= -h\lambda_{jump} + \sum_k \underbrace{\lambda_k E\left[r_k^d | r_k^d > \bar{h}_k\right]}_{\equiv \xi_k} \beta_k + \underbrace{\tilde{\lambda} E\left[\tilde{r}^d | \tilde{r}^d > \bar{h}\right]}_{\equiv \tilde{\xi}}. \end{aligned} \tag{10}$$

Each coefficient on β_k, $\lambda_k E\left[r_k | r_k > \bar{h}_k\right]$ or ξ_k, represents the upper tail risk for factor k. Equivalently, ξ_k summarizes the expected damage done to market makers by jumps in that factor. Larger risks or faster jump arrival rates must be compensated in equilibrium by a higher rate of fundamental trader liquidity consumption or by larger half-spreads h.

B. Empirical Implementation of the Model

Volume and Arrival Rates

The fundamental trader and total jump arrival rates are unobserved from the econometrician's perspective. Because active intermediaries should form unbiased expectations about future arrivals

[13] 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. The Internet Appendix describes an alternative approach to estimating likelihoods of factor moves without imposing the small jump assumption.

of fundamental traders, realized arrivals offer a noisy proxy for the anticipated arrival rate.[14] For this reason, I use realized volume (in units of 100-share round lots) as a proxy for expected arrival rates. The econometrician does not perceive expected volume $\bar{V}$ available to market makers, but under standard unbiasedness and independence assumptions on the error term, the heteroskedastic measurement error of $\bar{V}$ has no effect on recovered tail risk estimates asymptotically. Critically, substituting realized volume for expected volume does not violate the key expectational property of the dependent variable that facilities estimation of tail risk separately for each date. Realized volume adds noise to market maker expected volume, but the realized slope relating spreads, volumes, and betas across stocks converges to its true expected value ξ_t for large enough cross sections. By contrast, time variation in realized factor premia implies realized *return* slopes do not converge to their true expected value regardless of the size of the cross section.

I deflate realized volume to account for intermediation and reintermediation by market makers. Specifically, I deflate total volume by an estimated intermediation multiplier of $\kappa = 2.5$, as calibrated in a high-frequency setting for metals futures markets (Weller (2015)), to obtain an estimate for the correct scale of tail risk. However, because the intermediation multiplier may differ across asset classes, I refrain from interpreting the scale of implied risk and instead assume only that the intermediation multiplier is roughly constant across stocks within the 2005–2013 period. This assumption is tenable if increases in the share of "hot potato" trading brought about by HFT had already been achieved with their appreciable market share in most equity products by 2005. Under these conventions, the expected total arrival rate for traders of both types is given by

$$\widehat{\lambda_{FT} + \lambda_{jump}} = \frac{1}{2\kappa}\frac{V}{100} \tag{11}$$

for unsigned share volume V.

Non-Unit Depth and Trade Sizes and Discrete Spreads

The model supposes thus far that desired trading volume is fixed at one unit, or a round lot when applied to equities. However, even traders engaged in order splitting frequently trade larger quantities in a single order, and picking-off orders routinely execute against much of the displayed depth.

In the baseline model, the expected benefit of intermediating a (unit) order exactly offsets the expected cost of providing the option to trade to stale-quote snipers. If more depth is added, the potential cost to market makers of stale-quote sniping increases by a factor of d because snipers (in the aggregate) pick off all supplied liquidity at the stale bid or offer. Providing depth also increases the potential benefits of offering liquidity in that intended liquidity consumers can trade a larger number of shares before they themselves become liquidity suppliers.

[14] Appendix B discusses implicit assumptions on market makers' information sets.

Market makers' expected benefit per unit time of facilitating trade for a stochastic quantity q and an offered depth d becomes

$$\lambda_{FT} \times h_i \times (d + (E[q|q < d] - d) \Pr(q < d)). \tag{12}$$

Market makers only intermediate for desired trade sizes up to d—beyond that point, larger liquidity demands instead convert into resting limit orders or "walk the book" to consume liquidity at higher prices. By contrast, the cost of offering d shares increases linearly in d without bound, as stale-quote snipers pick off the entire offered depth across all venues when given the opportunity.[15] The net benefit per share of depth is weakly decreasing in d, because traders do not always consume up to $q = d$ units of liquidity.

If queue positions are known, market makers earn positive rents on inframarginal units, and market makers (or orders) toward the back of the book are closer to break-even on average. However, in high-frequency settings, continual churn in the limit order book makes position order uncertain (Yueshen (2014))—recall that the cancel-to-trade ratio exceeds 20 for stocks and 90 for exchange-traded products. In the random-sequencing limit, all units of liquidity offered satisfy a depth-adjusted equilibrium condition of Equation (3), because each unit of depth has the NBBO-average liquidity costs and benefits:

$$\lambda_{FT} \times h \times \underbrace{\{E[q|q < d] \times \Pr(q < d) + d \times \Pr(q \geq d)\}}_{\equiv q^*} = d \times \lambda_{jump} \times \Pr(J > h) \times E[J - h|J > h]. \tag{13}$$

Returning now to the empirical specification for the dependent variable, realized volume captures a quantity proportional to the volume-weighted arrival rate of both trader types, $\lambda_{FT} q^* + \lambda_{jump} d$. The equilibrium condition requires normalization by offered depth to relate liquidity variables on the left in terms of exogenous jump risk on the right. Combining Equations (10) and (13) and dividing by d (in round lots) obtains the adjusted empirical proxy for arrival rates:

$$\frac{1}{2\kappa} \frac{Vh}{d} \approx (\lambda_{FT} q^* + \lambda_{jump} d) \frac{1}{d} \times h = \sum_k \xi_k \beta_k + \tilde{\xi}. \tag{14}$$

Allowing for non-unit trading demands gives a second margin over which potential market makers can control their liquidity provision because optimal depth is not fixed at one unit. However, Equation (14) holds for all *marginal* market makers, i.e., those with the lowest spreads. Other market makers may play a strategy of offering more depth at higher costs of liquidity, but these traders satisfy a separate first-order condition.[16]

[15] Intermarket sweep orders or ISOs execute differently from Reg NMS protected orders and are quite popular among HFTs (e.g., HFT snipers). Sending ISOs simultaneously can clear all liquidity at the NBBO without incurring the delays associated with checking for potential price improvement at other market centers.

[16] These market makers balance expected intermediation gains (after depth at the NBBO is exhausted) against the

This simple adaptation also accounts for minimum tick sizes and discrete price increments. A half-spread exceeding the cost of providing a single unit of depth attracts additional market makers—the price of liquidity is artificially high—who in turn add depth to the book until the total depth at the NBBO equates expected intermediation gains and picking-off costs averaged across all orders as in Equation (13). For this reason I do not drop stocks typically at the minimum tick size, as depth offers sufficient adjustment in the context of the model. Nonetheless, dropping these stocks has minimal effect on the cross-sectional tail risk estimates.

C. Cross-Sectional Recovery of Tail Risks

Estimating tail risks first requires computing betas with respect to candidate realized return factors. I estimate backward-looking, rolling annual betas using daily returns r_{it} on candidate factor realizations f_{kt} for each stock in the filtered sample i:

$$r_{it} = \alpha_i + \sum_k \beta_{ik}^{(t)} f_{kt} + \epsilon_{it},\ \forall i. \tag{15}$$

Armed with these betas, I estimate (symmetric) tail risks cross-sectionally across stock-level observations via least absolute deviations regression:[17]

$$\left(\frac{Vh}{d}\right)_{it} = \tilde{\xi}_t + \xi_{t,MKT}\beta_{i,MKT} + \sum_{k \neq MKT} \xi_{t,k}\,|\beta_{ik}| + \delta_{it}. \tag{16}$$

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 100-share round lots, V is realized volume in 500-share units, h is the effective half-spread, and δ_{is} is a stock-specific error term for date t.[18] $\xi_{t,k}$ 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.

Equations (15) and (16) resemble Fama-MacBeth regressions for determining prices of factor risk. It differs in that all ξ_t estimates are of independent interest rather than only inputs into a single time-series average value. Unlike the realized factor premium at each date, the ex ante

expected jump cost conditional on the price jump exceeding their price of liquidity $h' > h$. In a richer model, these depth-providing market makers can facilitate intermediation chains (Weller (2015)).

[17] The market factor model is readily estimated in part because reliable negative betas are quite rare among common stocks. Consequently, excluding stocks with betas below a small, positive threshold excludes few stocks so little data is lost. However, imposing this non-negativity restriction with other candidate asset pricing factors is undesirable, as βs need not be centered around one, and they are frequently negative for several common factors, e.g., size (SMB) and value (HML). Appendix C shows that taking absolute values correctly accounts for negative betas with respect to non-market factors.

[18] The quoted spread sometimes overstates the true cost of consuming liquidity, as trades may execute inside the (displayed) spread. Replacing effective spreads with quoted spreads gives nearly-identical time series of implied risks, up to a scaling factor.

expected jumps are not (conditionally) stochastic, and their estimation relies on large N rather than on large T asymptotics. The key assumption needed for inference on each ξ_t term is that factor loadings β estimated over wider intervals are valid for each subinterval. The choice of annual betas trades off accuracy for precision in factor-loading estimates, but differences between betas for different horizons are likely to be small.[19]

For estimation, I use least absolute deviations (LAD, equivalently, median regression) rather than OLS in the second-stage regressions because the dependent variable is highly skewed, which results in too much weight being placed on fitting a small number of influential points. I compute standard errors via pairs bootstrap with 200 resamples to account for heteroskedasticity in the LAD errors (Koenker (2005)) and the slightly nonstandard two-stage methodology.[20]

IV. Data Description

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. The TAQ data aggregate orders from all Consolidated Tape Association exchanges and are timestamped to the second. I follow Holden and Jacobsen (2014) to recover cleaned effective spreads and market depths from the underlying TAQ data. Traded volume over each interval is directed observed. CRSP provides security attribute data (i.e., share codes), unique ticker-entity mappings, and daily shares outstanding for each security. In addition to CRSP and TAQ, I obtain intraday historical Chicago Board Options Exchange Volatility Index (VIX) data from Pi Trading.

The data sample consists of all common stocks (CRSP share code = 10 or 11) in the TAQ database from January 2004 to December 2013. Although TAQ starts in 1993, the spread may be too coarse prior to decimalization to provide a good guide for market-maker risks, and the large-jumps assumption cannot hold for bid-ask spreads on the order of several percent. In addition, the assumptions of continually updated spreads and minimal order processing costs are not plausible until algorithmic trading improvements in the mid-2000s (Foucault et al. (2003)). Slow market-maker responses to order flow through 2005 (Lyle et al. (2015)) suggests that even 2004 may be too early a start date; I focus on 2005–2013 for empirical tests for this reason and because of the model's reliance on HFT market making. Section V provides additional discussion of potential

[19] In principle, this estimation can be improved by estimating jump betas with respect to each set of factors (e.g., Todorov and Bollerslev (2010), Li, Todorov and Tauchen (2014)). I take this rather basic approach for three reasons. First, existing methods for computing jump betas are not well suited for multifactor models, especially in the presence of co-jumps. Second, high-frequency beta estimation is challenging across disparate liquidity environments, both over time and across stocks. Third, to a first approximation, jump betas are very similar to continuous betas at the stock level with only rare exceptions: for example, Todorov and Bollerslev (2010) find a cross-sectional correlation of 96.4% between jump and continuous betas for their sample of forty large stocks. Differences between betas associated with high- and low-frequency variation consequently should be a relatively minor source of measurement error.

[20] The two-stage procedure suffers from a standard generated regressors problem. The pairs bootstrap readily accounts for this issue by incorporating draws from Equations (15) and (16) simultaneously at a moderate computational cost.

time-variation in picking-off probabilities.

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 running 9:45–10:45am through 2:45–3:45pm. The filtered sample has 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 A.

V. Results

The empirical analysis in this study proceeds in two main steps. In the first step, I recover hourly tail risk estimates and compare these estimates with tail realizations and to alternative near-term forecasts such as the VIX. In the second step (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.

Figure III plots recovered market and idiosyncratic tail risks by hour over the 2004 to 2013 sample period. In estimating Equation (16), each hourly observation $\left(V\frac{h}{d}\right)_{it}$ is a volume-weighted sum across minutes of the within-minute average of effective spreads dividend by displayed depth. Total spread-implied market and idiosyncratic tail risks are clearly distinct.[21] Idiosyncratic jump risk, plotted in red, is dominated by market risk in several crisis periods. This result is reassuring as the period considered includes dominant market events such as the recent financial crisis and global recession.

Figure III illustrates the jump tail 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 in green. I compute changes by differencing jump tails for the same hour at date t and $t-1$ to account for intraday patterns and news that spans multiple trading hours, and I separate increases by a minimum distance of 10 trading days to isolate distinct events. The extracted set of events consists primarily of scheduled Fed-affiliated and macroeconomic announcements. 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.[22]

Figure IV 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)).

[21] This feature contrasts with the market and idiosyncratic tail shapes of Kelly and Jiang (2014), for example. The Kelly-Jiang measure exploits power law relations to show that market tail shapes inherit properties of asset-level tail shapes. However, average tail *locations* do not share this aggregation property.

[22] Several of these "peak news" days saw several large jumps realize in rapid succession. Notably, the implied tail risk measure cannot distinguish between market-maker expectations for one large jump or for a high arrival intensity of smaller jumps—the spread embeds minimal information about serial dependence of large price movements if all market-making risks are truly instantaneous. The success of the quarter-hour resolution applications in Section VI suggests that the tail measure may be of use in studying high-frequency tail dependence.

The trimmed sample eliminates the most pronounced jump patterns and the start and end of normal market hours, but the pronounced skewed U-shape pattern nonetheless obtains for each year in the sample.[23] Moreover, both plots rank 2007–2011 as among the most extreme years in the sample for expected and realized jumps and 2005–2006 and 2012–2013 as the least extreme years in the sample for these measures.

Table I quantifies the correlation relationships among spread-implied tail risks, volatility measures, and options-implied tail risks from Bollerslev and Todorov (2014). All measures are aggregated as weekly averages because weekly resolution is the highest frequency available under options-based alternatives. Tail measures are highly positively correlated, notwithstanding that risk-pricing information is embedded in the options-based approach but not in the spreads-based approach. Relative to the options-implied tails, high-frequency spread-implied tails are more similar to realized volatility and less similar to the VIX. This relationship is expected in that options-implied tails are "low frequency" (on the order of a week) and do not condition on intraweek information relevant to near-term volatility and jump tail risks. Likewise, the spread-based measure does not span the "long horizon" information embedded in options with more than a week to expiration that is also captured in the monthly VIX.

The market tail risk measure comoves especially strongly with realized market volatility at the weekly frequency. This correlation validates rather than indicts the tail risk measure. First, Bollerslev and Todorov (2011b) show that we cannot reject the hypothesis that realized continuous and jump variation have a one-factor structure at the weekly frequency. If the tail risk measure perfectly forecasts near-term tail realizations, we should expect this correlation to be indistinguishable from one at this level of granularity. Second, the tight link between anticipated jumps and anticipated volatility is paralleled by the relationship between options-implied tails and the VIX. Lower-frequency options-implied tails have an 88% correlation with its corresponding forward-looking volatility measure. Third, I show in the following section that the tail risk measure is associated with the component of market variation orthogonal to continuous variation, i.e., to the extent that the one-factor structure does not hold at higher frequencies, the tail risk measure explains residual jump variation in prices.

A. Empirical Tests

Regression Specifications

Taken together, these figures provide suggestive evidence that the spread-based tail risk measure captures high-frequency fluctuations in market risks. In this section, I formally test whether

[23] As suggested by Figures I and III, the tail risk measures achieve extreme values in 2008 for every hour of the trading day, even relative to other crisis years. I omit 2008 from the diagram to maintain resolution on the other years of the 2004–2013 sample.

implied tail risk coincide with jump realizations.[24] I also test whether the tail risk measure indeed captures jumps rather than volatility, because jump realizations and continuous variation are highly correlated.

I measure (jump) tail realizations both in units of spreads and basis points and in event counts and event sums (accounting for event size). The key verification regression takes the following form for the market factor (and is later repeated for a financials proxy in Section VI):

$$tail\_realization_t = \alpha + \beta \xi_{t-\Delta,MKT} + \gamma VIX_t + \delta CV_t + \epsilon_t. \tag{17}$$

Given the tail factor's close coevolution 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 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 (following the continuous and jump variation decomposition of Mancini (2009), among others). 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.

The tail realization measures used in the regression are as follows. Basis-point jumps count the number of events in which the minutely return exceeds 10 basis points, my implicit "large jump" threshold. The jump sum is a weighted average 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. The corresponding jump sum measure is a weighted average 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. Illiquidity-driven noise in jump estimates is not a concern for the SPY, because it trades at least once during every minute for every market hour of 2005–2013.

Table II presents summary statistics for tail realizations. 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 et al. (2013)), but they also include more frequent "medium-sized" jumps, as well: jump detection methodologies normalize by a measure of local volatility to distinguish clearly between continuous and discontinuous variation. By contrast, large price movements generate picking-off opportunities regardless of the underlying volatility environment so long as they move the "fundamental value" outside of the spread within a very short time period. For this reason, I employ relatively simple measures of realized tails, noting that the model does not distinguish between rare, truly discontinuous price movements and extremely rapid continuous ones

[24]Equation (3) operates under the physical measure if high-frequency market makers are risk neutral. Under this assumption comparing implied jump risks to jump realizations properly accounts for $\mathbb{P}$s and $\mathbb{Q}$s.

associated with high local volatility.

I repeat regression (17) for two choices of Δ. First, I run the contemporaneous regression using $\Delta = 0$. This regression should have high explanatory power if the model is true because market makers adjust their spreads nearly every instant rather than only once every hour. The recovered tail risk measure $\xi_{t,MKT}$ has the interpretation of the *within-hour average anticipated jump risk*.[25] I then assess within-period predictive ability with $\Delta \in (0, 1)$ periods ahead by using the average spread for the first minute of the hour for predicting the jump activity in the remainder of the period. This measure represents the *start-of-hour anticipated jump risk* under the assumption that anticipated volume is constant over the period from the market maker's perspective. This variant on the jump measure uses less forward-looking information (realized volume remains), but it also handicaps the jump measure because hourly average spreads capture within-hour variation in perceived jump intensities and magnitudes, whereas the first-minute spread does not.[26]

Recognizing that jump intensities and continuous variation are persistent, I repeat the beginning-of-period forecasts with lagged tail realizations and explanatory variables. This specification challenges the predictive ability of the cross section of spreads because the sizable persistence in volatility and jump risks is differenced out. In addition to running regressions for all trading hours from 2005 to 2013, I also average all variables by hour within each year to assess my measure's ability to pick up diurnal patterns (setting $\Delta = 0$).

Regression Results

Table III presents results from the baseline test of contemporaneous forecast jump tails on realized jumps. Standard errors are robust to heteroskedasticity (both panels) and serial correlation of up to 126 trading hours (top panel). I normalize $\xi_{t,MKT}$ by dividing by its standard deviation to facilitate interpretation of coefficients (for comparison, the standard deviation of the VIX in this period is 9.90). 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 and economically significant. For example, a one standard deviation increase in the market jump tail risk is associated with 5.47 additional realized basis-point jumps per trading hour and 69.7 additional weighted jumps (the measure predicts both intensity and size). The jump tail measure is only partly subsumed by the VIX, and they have roughly equal impact on the dependent variable for a one standard deviation change. The coefficient on continuous variation is inconsistent or driven

[25] Within-hour variation in tail risk does not introduce estimation bias for ξ because the product $V\frac{h}{d}$ varies linearly with ξ with a fixed exposure coefficient β.

[26] Although measured spreads and depth do not use forward-looking information, observed realized volume is a full-period measure. The Internet Appendix considers a third case of $\Delta = 1$ for which no forward-looking data is used. We should expect the measure to have weaker forecasting ability for an hour ahead because HFT market makers adapt their spreads and depth almost instantaneously to reflect near-term anticipated tail risks. Nonetheless, I find moderately-strong predictability of market jump realizations for the hour ahead. Section VI finds corroborating evidence for near-term predictive ability in the contexts of the 2010 Flash Crash and FOMC announcements.

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 jump tail risk rather than of contemporaneous or anticipated volatility.

Diurnal results are slightly smaller in coefficient magnitudes but are comparably significant statistically with and without the VIX control. Including continuous variation results in a serious multicollinearity problem because continuous and total variation are so strongly related with the coarsened hour-year averages. This low granularity destroys too much information to reliably distinguish between intraday patterns in jump intensities and continuous variation.

Table IV presents analogous results for beginning-of-period predictability (top panel) and lagged tail realizations and explanatory variables (bottom panel). Beginning-of-period predictability is nearly as strong as the full-period contemporaneous results of the preceding table because start-of-period spreads are very similar to full-period average spreads. These quantities differ substantially only if the perceived tail risk varies significantly throughout the hour. The bottom panel presents a more serious challenge to the forecasting power of the model. Controlling for lagged quantities removes the artificially high explanatory power of the forecasting variable that derives from persistent tail risks. As expected, coefficients decrease slightly with the inclusion of lagged tail risk and explanatory variables. However, the estimated predictive ability of the tail risk measure is nonetheless economically and statistically large. A one standard deviation increase in $\xi_{t-\Delta,MKT}$ is associated with an increase of 2.4 basis-point jumps per hour and 33.0 weighted basis-point jumps.[27] As before, these estimates survive the inclusion of measures of both future total volatility and contemporaneous continuous variation.

These tests confirm the suggestive evidence of Figures I, III, and IV. 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, it 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 jump tail risks. In the Internet Appendix, I provide additional tests comparing with a volatility-matched placebo factor to ensure that the estimation does not spuriously pick up variables unrelated to factor tails.

B. Other Sources of the Spread

Other sources of the spread constitute the primary empirical threat to interpretation of recovered coefficients as tail risks. An extensive market microstructure literature describes three primary sources of the bid-ask spread, of which picking-off risk is only a part. The bid-ask spread can be attributed to order processing, inventory, and adverse selection costs. In this section, I consider

[27] I confirm that results are not driven by extreme observations by recomputing both tables with log counts and log sums. All hourly coefficients remain statistically significant at the 99% level. Tables are available upon request.

inventory risks and adverse selection costs other than picking-off risks and their respective effects on quantities estimated in the model. In the algorithmic market-making era, order-processing costs are indistinguishable from zero and are ignored.[28]

Inventory Risk

Market makers must be compensated for exposure to price variation of assets in inventory. If market makers do not hedge inventory risks, e.g., using liquid factor-mimicking indexes, market-maker inventories may contaminate jump tail risks estimated from bid-ask spreads. To address potential contamination by inventory risk, I isolate the component of spreads contributed by adverse selection risks—including picking-off risk—rather than inventory risks. In place of effective spreads, I substitute the realized adverse selection component of effective spreads because it encompasses picking-off risk and plausibly does not relate to inventory risk associated with current or anticipated holdings.

A standard proxy for the adverse-selection or price-impact component of spreads is given by the difference between quote midpoints at time t and $t+5$ minutes (Glosten (1987)):

$$adv\_sel_{it} = q_{it}\left(m_{i,t+5m}/m_{it} - 1\right), \tag{18}$$

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. This quantity represents the "permanent" price impact of transactions as the difference between the 5-minute ahead price and the current price. I average this adverse selection value by stock and minute to obtain a continuous proxy for adverse selection costs. The updated specification of Equation (16) becomes

$$\frac{V_{it} \times adv\_sel_{it}}{d_{it}} = \tilde{\xi}_t + \xi_{t,MKT}\beta_{i,MKT} + \sum_{k \neq MKT} \xi_{t,k}\left|\beta_{ik}\right| + \delta_{it}. \tag{19}$$

Table V replicates Table III using the tail estimates from Equation (19). All coefficients retain comparable levels of economic and statistical significance. This table emphasizes that the tail risk measure is not contaminated by (omitted) inventory risk throughout the sample.

[28]Investigating these alternative sources of the spread takes on increased importance in light of the dramatic reductions in trading costs throughout the early 2000s. For example, Hendershott et al. (2011) reports that the average effective half-spread for stocks in the middle market cap quintile declines from roughly 14 basis points in 2001 to about 3 basis points in 2006. The reduction in effective spreads primarily derives from declines in adverse selection associated with better-controlled picking-off risk for *idiosyncratic* price movements (Lyle et al. (2015)), which calls into question the assumption of a constant picking-off probability given a jump. However, Lyle et al. (2015) also show that market makers closely monitor and anticipate *non-idiosyncratic* news events as early as 2002. For this reason, variation in picking-off probabilities for *factor jumps* is likely to be small and not a major driver of my time series results.

Adverse Selection

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

1. Intermediation against informed traders with long-lived information (slow);

2. Picking off by stale-quote snipers (fast).

This contrast has a theoretical basis in models of informed trading by insiders. In insider models such as Kyle (1985) or Back (1992), insider trades generate losses to the market maker at an $O(dt)$ rate. Foucault, Hombert and Rosu (2015) introduce a model of informed trading with $O(dz)$ innovations resulting from "news trading" on a flow of signals. However, as the time interval approaches zero, the likelihood of a price movement larger than a fixed $\epsilon > 0$ threshold nevertheless decreases toward zero. By contrast, picking-off costs as a response to jump realizations are $O(1)$ and do not scale even as the time interval becomes small. Hence long-lived information leads to slow erosion of market-maker profits, whereas price jumps contribute to a rapid deterioration of profits (conditional on an event).

The potential empirical confound posed by slow adverse selection is that cross-sectional variation in spreads reflects differences in exposures to long-lived informational risk in addition to picking-off risk. Economically, estimation bias arises when factor exposures align with the risk imposed by long-lived informed traders. For example, for market tail estimates to be biased, assets with high market betas must be more exposed to non-jump adverse selection risk. Formally, slow adverse selection biases tail risk estimates if and only if the volume- and depth-scaled component of half-spreads not associated with tail risk exposures h^- is cross-sectionally correlated with factor loadings, i.e., $cov\left(h^-\frac{V}{d}, \beta_k\right) \neq 0$.

Such alignment is tantamount to market participants having private information on the underlying factor. I follow Gorton and Pennacchi (1993) and others in suggesting that private informational advantages are unlikely for systematic factors. Gorton and Pennacchi (1993) argue that composite products, and by extension, the corresponding factor that is mimicked, are exposed to minimal risk of non-public adverse selection because insider information is typically known at the security level rather than at the aggregate level. By contrast with systematic factor risks, recovered idiosyncratic tail risk is likely to be biased upwards because non-jump adverse selection increases market-maker costs. However, over short time intervals, stock-level adverse selection should be roughly constant on average, and using differences to isolate local variation in common idiosyncratic jump risk obviates this concern.[29]

I support this assumption 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 et al.

[29] I take this approach in discussing the 2010 Flash Crash in Section VI.

(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 encapsulate 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).[30]

Table VI 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 control on the right-hand side of Equation (16). As for the jump tail measure net of potential inventory risk, the jump tail measure net of slow adverse selection risk performs very similarly to the baseline specification (Tables III) in matching time variation in market tail realizations; indeed, no tail risk coefficient in Table VI is statistically distinguishable at the 5% level from its counterpart in Table III.

Summary

The tail risk series closely resemble each other. Pairwise correlations with the baseline specification are 92.0% controlling for inventory risk and 98.7% controlling for slow adverse selection. Figure V depicts the recovered time series of market tail risks in the baseline, inventory-risk robust, and slow adverse-selection robust specifications. These strong correlations are borne out almost point by point in the normalized series. In sum, neither inventory risk nor slow adverse selection appear to meaningfully contaminate the measure of implied market tail risks.

VI. Applications

I supplement formal tests of Section V with comparisons to known jumps associated with macroeconomic events. These applications serve both as qualitative verifications that the measure picks up and anticipates tail realizations in a variety of challenging macroeconomic settings and as a demonstration of the broad range of settings to which this technique can be applied.

I first demonstrate that the market tail risk measure is elevated shortly before and during the hours of the 2010 Flash Crash and Federal Open Market Committee (FOMC) announcements. These settings offer evidence that the risk measure is anticipatory rather than a reflection of concurrent jumps and that it assesses tail risks rather than volatility. I then show that the tail risk extraction methodology disentangles aggregate market from financial risks. Using a simple two-factor return

[30] An active literature attends to issues encountered in estimating PIN in fast-moving equity markets. Boehmer, Grammig and Theissen (2007) discuss downward biases resulting from buy-sell misclassification errors. Lin and Ke (2011) describe a potential bias arising from maximum likelihood estimation with large numbers of trades serving as exponential terms. Yan and Zhang (2012) detail optimization failures on the boundary of the parameter space of the PIN model, which in turn translate into errors in the implied PIN measure. I find that results are robust across alternate estimation methodologies.

model, I find that extremes for the financial jump risk factor correspond with major banking events often distinct from those associated with broader stock market movements; importantly, changes in implied factor risks indeed correspond with factor-specific innovations.

These macroeconomic events pose forecasting challenges difficult to meet with other existing methodologies. Although my findings qualitatively match intuition, there are few existing benchmarks against which I can assess results quantitatively—no other methodology allows for measurement of conditional tail risks intraday. Indeed, a principal contribution of this paper is to develop a tool for quantifying short-lived jump risks for a variety of factors in real time for precisely these types of 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.[31] Kirilenko, Kyle, Mehrdad and Tuzun (2011) tabulate buyers and sellers in S&P 500 E-Mini futures ("E-Mini") during the Flash Crash, but no corresponding data exists to evaluate the extensive knock-on redistributive effects associated with 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,[32] yet fears of future crashes are inherently difficult 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 thereby bring about 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. Kirilenko et al. (2011) 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 extreme turnover, or "hot potato" activity.

To demonstrate its utility in assessing the costs of flash crashes, I construct contemporaneous tail risk measures around and during May 6, 2010 using the 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

[31] 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 (2015)).

[32] Such articles continue to be written several years after the crash, e.g., http://www.bloomberg.com/news/articles/2012-05-07/flash-crash-story-looks-more-like-a-fairy-tale and http://money.cnn.com/2014/05/06/investing/flash-crash-anniversary/.

to achieve high resolution on the crash interval (2:30–3:00pm) and surrounding trading hours.

Figure VI plots market and idiosyncratic tail measures from April 28, 2010 through May 14, 2010 for every quarter hour from 9:45am to 3:45pm. To capture innovations and place risk changes in context of normal diurnal and slow-moving macroeconomic variation, I difference the previous day's value at the same quarter hour and divide by the rolling standard deviation of differences for the same quarter hour over the preceding 63 trading days (a calendar quarter). Several features are readily apparent. First, the Flash Crash itself is associated with extreme contemporaneous elevations of both the market (104 standard deviations) and idiosyncratic (71 standard deviations) tail risk measures. Second, jump risks remain elevated for the remainder of the trading day and throughout May 7, 2010. Both risks return to near-normal levels after the May 8–9 weekend. Third, market tail risks increase a quarter hour *before* idiosyncratic tail risks, likely because the Flash Crash started in the E-Mini, a nearly-ideal S&P 500 index proxy. Importantly, market factor risks are identified using the cross section of spreads, depth, and volume rather than the characteristics of any particular security—market tail risk only reflects abnormal liquidity demands in the E-Mini to the extent that they spill over into all stocks in proportion to each security's market beta.

Intriguingly, the market tail risk measure achieves 17 standard deviations above its quarter-hour norm in the 2:15–2:30pm interval. Market makers appear to anticipate distress conditions even before Waddell & Reed initialized its trading algorithm at 2:32pm (Menkveld and Yueshen (2015)). 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.[33]

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 event for the VIX except in the rapidity of its increase intraday.[34] 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 VII).

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 VII, we observe that all three risk measures return to roughly their pre-Flash Crash levels only days later—these differences are neither statistically significant nor economically

[33] As noted in the introduction, the only other occasions during which tail risk increases by at least 16 standard deviations are the intraday crashes of February 27, 2007, August 9, 2007, September 29, 2008, and August 5, 2011. There are no false positives associated with such large changes in implied tail risks.

[34] Andersen, Bondarenko and Gonzalez-Perez (2015) find that the VIX breaks down as a reflection of market risk precisely during times of greatest market stress. Their alternative corridor implied volatility has virtually identical pre-Crash dynamics to the VIX, however.

large. 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 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—typically scheduled to within minutes—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 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 jump risk can make the CAPM appear to operate if the CAPM works for discontinuous returns (as suggested by Bollerslev, Li and Todorov (2014)). Moreover, the sample of FOMC announcements may be too short for these jump risks to have been realized.

Existing alternatives cannot provide resolution on such high-frequency tail variation. Cross-sectional methods such as Kelly and Jiang (2014)'s cannot identify daily or intraday variation in jump tails because they rely on realized jumps within the estimation window. Options-based approaches are also inapplicable for short-term jump tail analysis because they require a panel of securities that mature in trading hours surrounding FOMC announcements.

I apply my tail 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. I use full-period average spreads to measure jump risks because announcements are typically at the start of the quarter-hour, so the previous quarter hour's measure is a better assessment of anticipated risks.

Figure VIII 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.[35] 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).[36] 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 (2) fear of early information leakage and attendant price jumps, as suggested by the empirical investigations of Bernile, Hu and Tang (2015) and Kurov, Sancetta, Strasser and Wolfe (2015). Notably the measure does not simply reflect contemporaneous realized volatility around the FOMC announcement: Lucca and Moench (2015) instead find that volatility is monotone decreasing 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 conclude with a brief study of the 2007–2008 Financial Crisis. Specifically, I apply the jump extraction technique to discern the magnitude of perceived jump risks to a "financials" factor independent from market risks. The choice of the financial sector is driven by its economic importance during the 2000s as well as the difficulty of disentangling financial sector risks from market risks using alternate methods; during this period, the daily correlation of XLF, my financial sector factor-mimicking portfolio, and the SPY often exceeds 90%. At the same time, Aït-Sahalia and Xiu (2015) 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. In addition, because financials are potentially an unpriced factor, time variation in financial sector tail risks may not be detectable by standard cross-sectional approaches applied to average returns.

The central regression in this analysis is Equation (16), modified to accommodate a financials

[35] 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.

[36] These early-response results accord with the findings of Jiang, Lo and Verdelhan (2011) in the U.S. Treasury bond market. The authors find increased spreads, decreased depth, and stagnant trading volume in the five minutes before major market news announcements.

factor:

$$\left(\frac{Vh}{d}\right)_{it} = \tilde{\xi}_t + \xi_{t,MKT}\beta_{i,MKT} + \xi_{t,FIN}\left|\beta_{i,FIN}\right| + \epsilon_{it}, \ \forall t. \tag{20}$$

I exclude co-jump terms for the market and financial factors because the joint risk of tail events in the market and financial factors is not of independent interest.[37]

Figure IX 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 III, but it differs somewhat in the events corresponding with the largest changes in tail uncertainty. Several events associated with large market risk increases in the one-factor model (Figure III) 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 have pronounced effects for the aggregate market but not for the financial sector independently.

I now show quantitatively that the recovered financials tail risks indeed correspond with jumps in the financials factor. Table VII presents analogous results with the market-factor analysis in Section V. 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, I include market jump tails as a control to quantify the degree to which jump types are successfully disentangled.

Throughout, results are economically and statistically significant. A one standard deviation increase in the jump tail measure corresponds with approximately 8.2 more basis-point jumps and 1.6 spread jumps on a baseline standard deviation of 10.7 and 2.5 jumps per hour, respectively. Market tails have no incremental explanatory ability beyond that captured by $\xi_{t,FIN}$. The financial factor's explanatory ability is replicated in the averages by hour-year, suggesting that intraday patterns in tail risks are also captured by the jump tail measure. This brief analysis suggests that the jump tail extraction technique successfully and separately identifies market and financial risks and changes in risks at high frequency.

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 prices are higher for these

[37] Appendix C discusses this point in additional detail.

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 straightforward methodology for extracting some of this factor risk information in real time.

My methodology is distinguished from existing methods in its ability to obtain information about (1) priced and non-priced factor risks (2) at a sub-hour frequency (3) for short look-ahead horizons. This cross-sectional approach works both for evaluation of persistent tail risks and for assessment of their evolution intraday. Spreads are unique among existing data sources in their ability to reveal intraday changes in tail fears for return factors.

This methodology offers a valuable tool for researchers to evaluate tail risks in real time. High-frequency identification of tail risk changes in a wide range of return factors provides viable alternatives to the VIX and economic-uncertainty measures (e.g., the EPU index of Baker, Bloom and Davis (2013)). Regulators, too, might benefit: the extreme anticipatory rise in tail risk before the onset of the 2010 Flash Crash suggests that the measure may have predictive power for severe market disruptions. These properties are left for future investigation.

References

Aït-Sahalia, Yacine and Dacheng Xiu, "Principal Component Analysis of High Frequency Data," NBER Working Paper September 2015.

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

_ , Oleg Bondarenko, and Maria T. Gonzalez-Perez, "Exploring Return Dynamics via Corridor Implied Volatility," Working Paper 2015.

Back, Kerry, "Insider Trading in Continuous Time," *Review of Financial Studies*, 1992, *5* (3), 387–409.

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

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, "Measuring Economic Policy Uncertainty," Working Paper 2013.

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

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

Bernile, Gennaro, Jianfeng Hu, and Yuehua Tang, "Can Information Be Locked-Up? Informed Trading Ahead of Macro-News Announcements," Working Paper 2015.

Boehmer, Ekkehart, Joachim Grammig, and Erik Theissen, "Estimating the Probability of Informed Trading—Does Trade Misclassification Matter?," *Journal of Financial Markets*, 2007, *10* (1), 26 – 47.

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

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

_ and _ , "Time-Varying Jump Tails," *Journal of Econometrics*, 2014, *183* (2), 168–180. Analysis of Financial Data.

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

—, **Sophia Zhengzi Li, and Viktor Todorov**, "Roughing up Beta: Continuous vs. Discontinuous Betas, and the Cross-Section of Expected Stock Returns," Working Paper 2014.

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

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, "High-Frequency Trading and Price Discovery," *Review of Financial Studies*, 2014, *27* (8), 2267–2306.

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

Budish, Eric, Peter Cramton, and John Shim, "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," Working Paper February 2015.

Cardella, Laura, Jia Hao, and Ivalina Kalcheva, "Make and Take Fees in the U.S. Equity Market," Working Paper 2015.

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

CFTC and SEC, "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 2010.

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

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

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

—, **Marcos M. López de Prado, and Maureen O'Hara**, "Flow Toxicity and Liquidity in a High-frequency World," *Review of Financial Studies*, 2012, *25* (5), 1457–1493.

—, **Nicholas M. Kiefer, Maureen O'Hara, and Joseph B. Paperman**, "Liquidity, Information, and Infrequently Traded Stocks," *The Journal of Finance*, 1996, *51* (4), 1405–1436.

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

Fleming, Michael J. and Monika Piazzesi, "Monetary Policy Tick-by-Tick," Working Paper 2005.

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

_ **, Johan Hombert, and Ioanid Rosu**, "News Trading and Speed," Working Paper 2015.

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

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

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

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

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

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

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

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

Ito, Takatoshi and Masahiro Yamada, "High-frequency, Algorithmic Spillovers Between NASDAQ and Forex," Working Paper 21122 April 2015.

Jiang, George J., Ingrid Lo, and Adrien Verdelhan, "Information Shocks, Liquidity Shocks, Jumps, and Price Discovery: Evidence from the U.S. Treasury Market," *Journal of Financial and Quantitative Analysis*, 4 2011, *46*, 527–551.

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

Kirilenko, Andrei, Albert S. Kyle, Samadi Mehrdad, and Tugkan Tuzun, "The Flash Crash: The Impact of High Frequency Trading on an Electronic Market," Working Paper May 2011.

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

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

Kurov, Alexander, Alessio Sancetta, Georg Strasser, and Marketa Halova Wolfe, "Price Drift Before U.S. Macroeconomic News," Working Paper 2015.

Kyle, Albert S., "Continuous Auctions and Insider Trading," *Econometrica*, 1985, *53* (6), 1315–1335.

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

Li, Jia, Viktor Todorov, and George Tauchen, "Jump Regressions," Working Paper 2014.

Lin, Hsiou-Wei William and Wen-Chyan Ke, "A Computing Bias in Estimating the Probability of Informed Trading," *Journal of Financial Markets*, 2011, *14* (4), 625 – 640.

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

Lyle, Matthew, James Naughton, and Brian M. Weller, "How Does Algorithmic Trading Improve Liquidity?," Working Paper 2015.

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

Menkveld, Albert J. and Bart Zhou Yueshen, "The Flash Crash: A Cautionary Tale about Highly Fragmented Markets," Working Paper 2015.

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

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

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

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

Weller, Brian M., "Intermediation Chains," Working Paper 2015.

Yan, Yuxing and Shaojun Zhang, "An Improved Estimation Method and Empirical Properties of the Probability of Informed Trading," *Journal of Banking & Finance*, 2012, *36* (2), 454 – 467.

_ **and** _ , "Quality of {PIN} Estimates and the PIN-Return Relationship," *Journal of Banking & Finance*, 2014, *43* (0), 137 – 149.

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

Figure I: Hourly Market Tail Risks, 2004–2013

This figure plots rolling one-month means of hourly cross-sectional slope estimates ξ_{mt} (blue) for an order book depth-adjusted version of Equation (2) for each trading date in 2004–2013. Dashed blue bands depict corresponding 95% confidence intervals. Realized volatility (red) is estimated using minutely squared returns on the SPY and scaled to the hourly frequency. The VIX is plotted using the right axis (gold) for comparison.

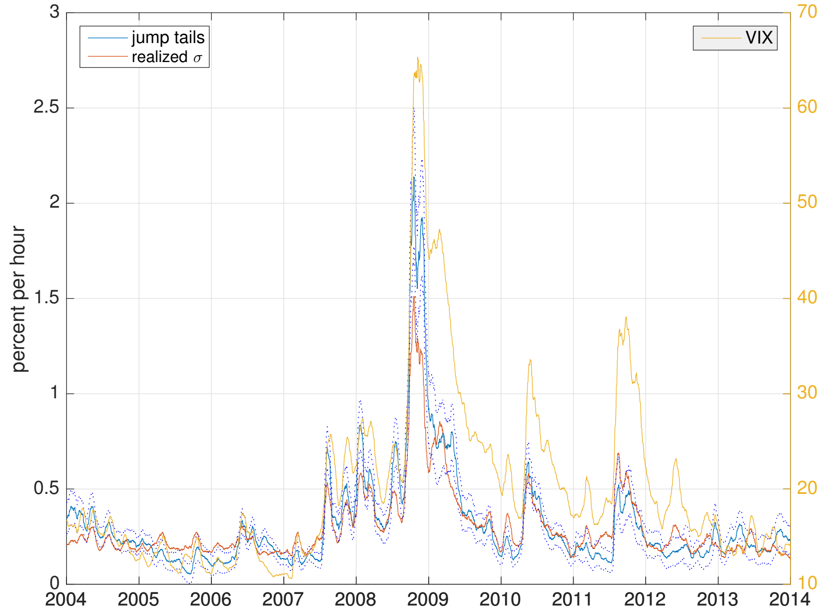

Figure II: Potential Intermediation Outcomes

This figure presents potential outcomes of offering liquidity with a half-spread h. Values at terminal nodes represent market-maker payoffs, and branch labels represent conditional probabilities. The green line indicates liquidity supply to a fundamental trader, whereas the red line indicates picking off by other market makers.

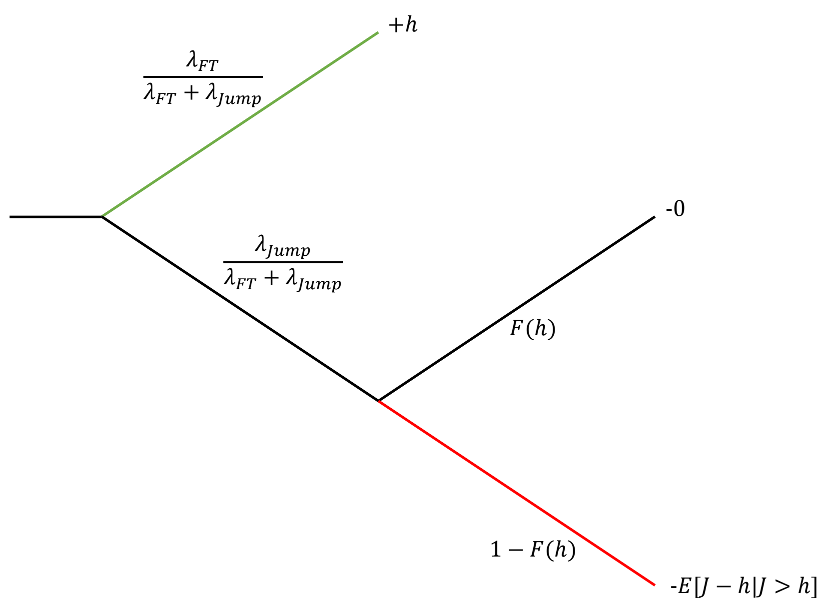

Figure III: Anticipated Jumps by Hour, 2004–2013

This figure plots hourly estimated tail risks for market and idiosyncratic risks in a one-factor market model. The fifteen largest increases in tail risks within a one-month window are overlaid with a green 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 of tail risk news days. Standardized values divide by the full time-series standard deviation. Blue 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.

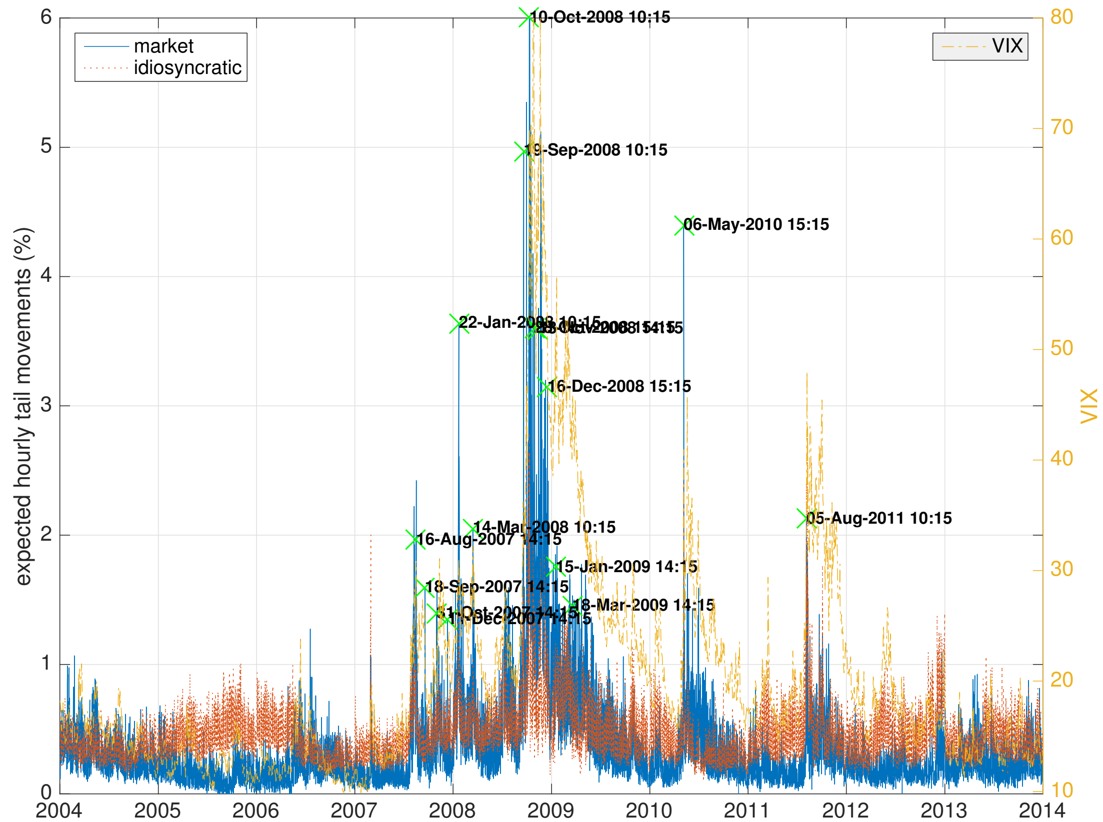

Hour	Std. Value	Event
16-Aug-07 14:15	6.54	Fed approves changes to its primary credit discount window facility
18-Sep-07 14:15	6.08	FOMC lowers Fed funds target 50 basis points to 4.75%
31-Oct-07 14:15	5.29	FOMC lowers Fed funds target 25 basis points to 4.5%
11-Dec-07 14:15	4.84	FOMC lowers Fed funds target 25 basis points to 4.25%
22-Jan-08 10:15	9.71	FOMC lowers Fed funds target 75 basis points to 3.5%
14-Mar-08 10:15	5.88	New York Fed drops bailout deal to save Bear Stearns
19-Sep-08 15:15	13.27	TARP proposed; short-selling ban; global campaign by central banks
10-Oct-08 15:15	23.71	Stock market crashes in Asia, Europe, and the United States
28-Oct-08 15:15	9.95	First round of TARP bank bailouts ($115 billion)
13-Nov-08 14:15	11.31	Large negative jobless claims surprise; most new claims since 9/11
16-Dec-08 15:15	7.65	FOMC lowers Fed funds target to 0-0.25%
15-Jan-09 14:15	4.93	Senate approves release of $350 billion of TARP funds
18-Mar-09 14:15	4.04	FOMC announces $1 trillion in new Treasury bond and MBS purchases
06-May-10 15:15	16.83	2010 Flash Crash
05-Aug-11 14:15	4.68	S&P downgrades U.S. government debt to AA+

Figure IV: Intraday Jumps by Hour, 2004–2013

This figure plots hourly means of market tail risks by hour and year (top) and of weighted realized spread jumps by hour and year (bottom). Realized spread jumps are a weighted average 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. The underlying economic model is a one-factor market model. 2008 is a positive outlier and is omitted from both diagrams to preserve resolution on the other years in the sample.

(a) Anticipated Jump Risks (ex 2008) (b) Weighted Realized Spread Jumps (ex 2008)

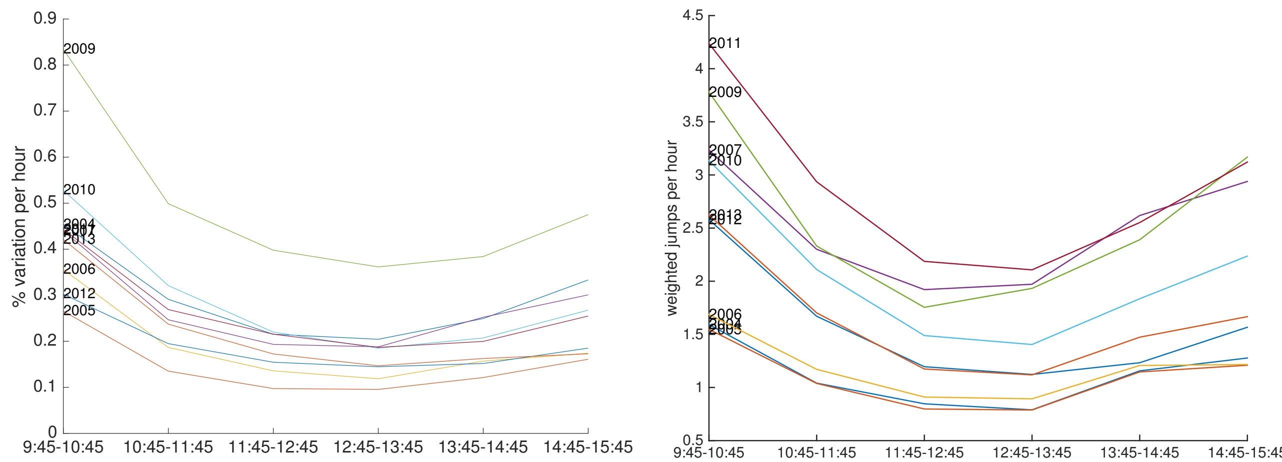

Figure V: Comparison of Implied Market Jump Risks Net of Alternative Sources of the Spread

(a) Forecast Market Tail Risks (b) Forecast Market Tail Risks (Normalized)

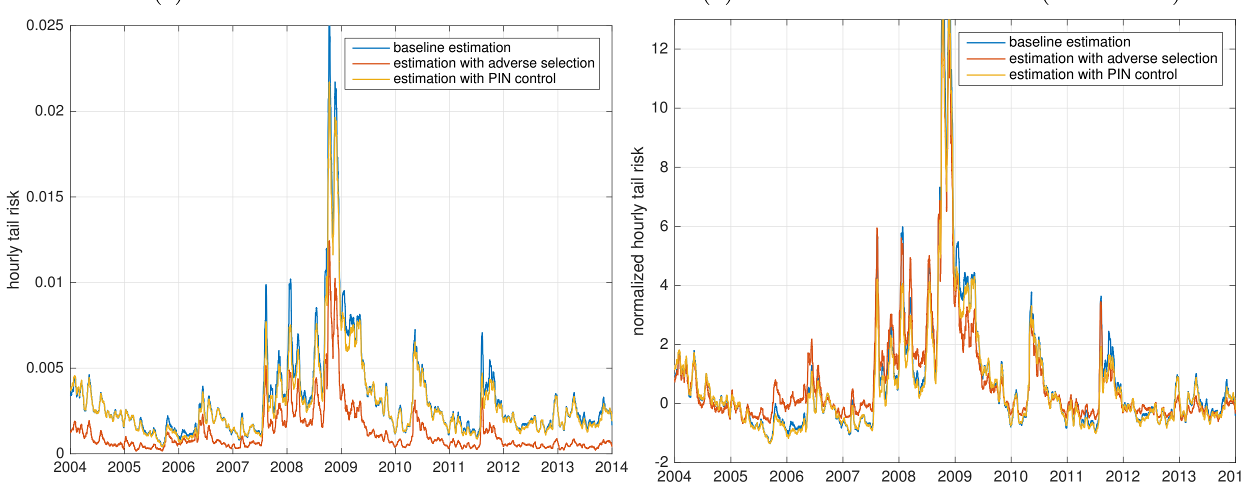

Figures plot rolling ten-day means of hourly estimated market tail risks for each trading date in 2004–2013. The blue line is the baseline estimation of Equation (16). The red line replaces the effective half-spread with realized adverse selection. The gold line adds a stock-quarter control for the probability of informed trading. The right plot aligns the series by subtracting series means and dividing by series standard deviations.

Figure VI: 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 previous day's value during the same quarter hour and dividing by the 15-minute specific standard deviation of this value across all dates in the preceding 63 trading days. The top figure plots the normalized value for the market factor before (green), during (red), and after (orange) May 6, 2010. The dotted purple line is the normalized 15-minute estimate for realized volatility. Black circles denote the 2:30–3:00pm 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.

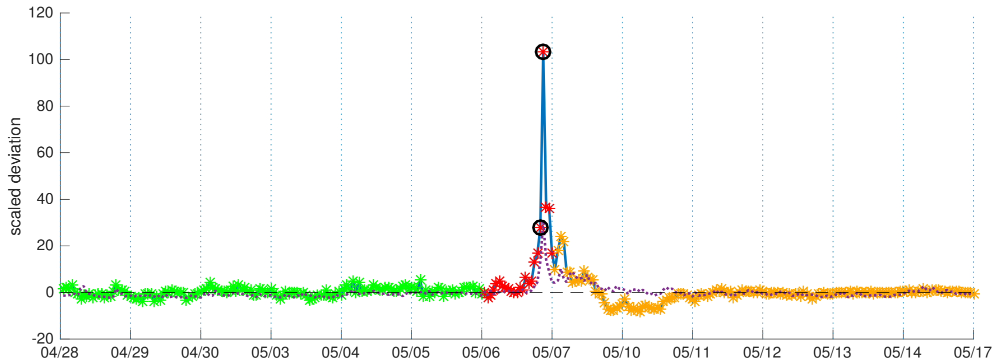

(a) Anticipated Market Factor Jumps

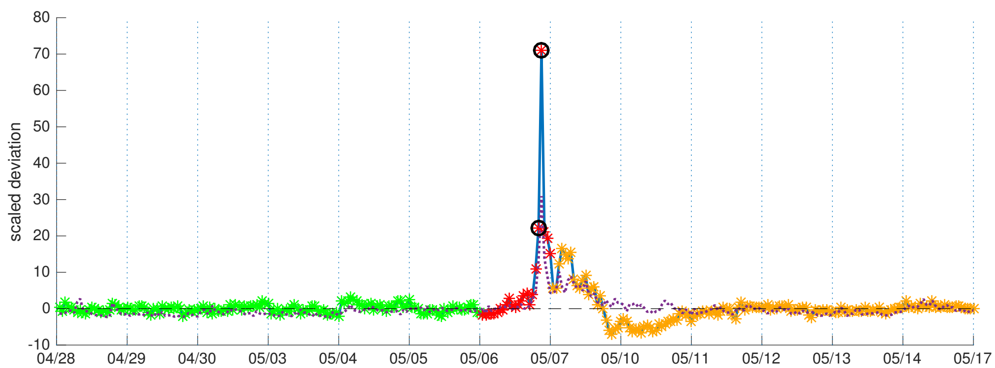

(b) Anticipated Idiosyncratic Jumps

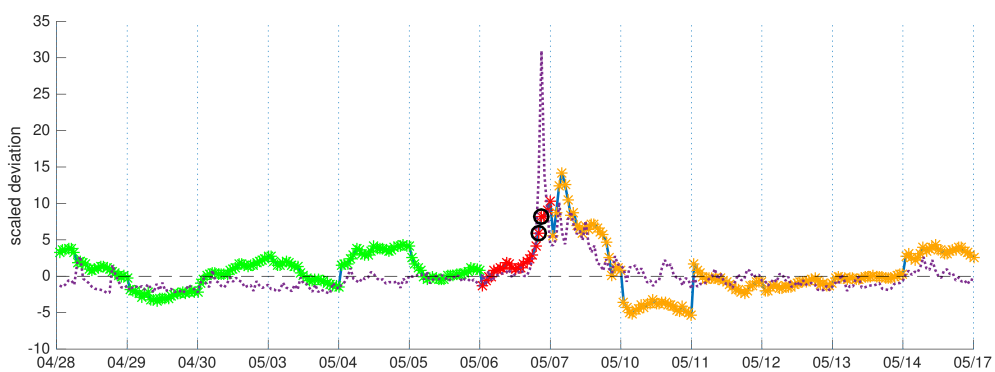

(c) Intraday VIX

Figure VII: 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. Blue and red lines correspond to implied market and idiosyncratic jumps, with their associated scale on the left axis. The orange line corresponds with the VIX, with its associated scale on the right axis. The dashed line marks May 6, 2010, and the dotted line marks two business days after the event, May 10, 2010. The associated table regresses the inferred tail risk measure against a constant and a post-Flash Crash indicator for quarter hours from April 28, 2010 through May 14, 2010. Both May 6, 2010 and May 7, 2010 are excluded. Standard errors are clustered by day.

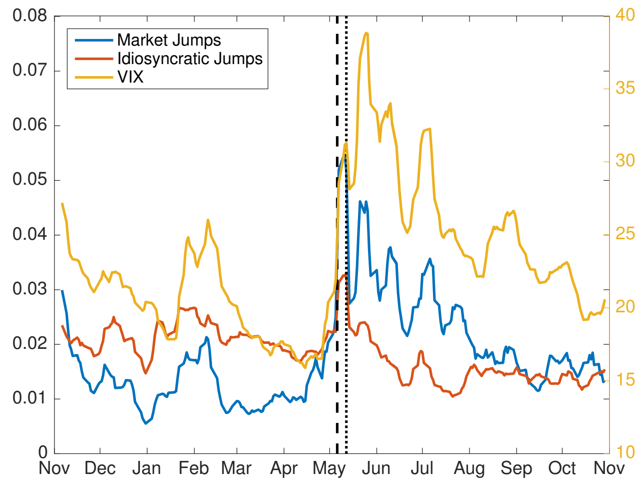

$$\xi_t = \alpha + \beta \mathbf{1}_{\text{After FC}} + \epsilon_t$$

	(1)	(2)
	Market Tail Risk	Idio. Tail Risk
Constant	0.025***	0.027***
	(10.122)	(25.218)
$\mathbf{1}$(after 5/6)	0.004	0.001
	(1.125)	(0.302)
Obs.	264	264
R^2	0.015	0.001

t-statistics are given in parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, $^{*}p < 0.1$.

Figure VIII: Intraday Tail Risks around FOMC Announcements

(a) Quarter-Hour Deviations

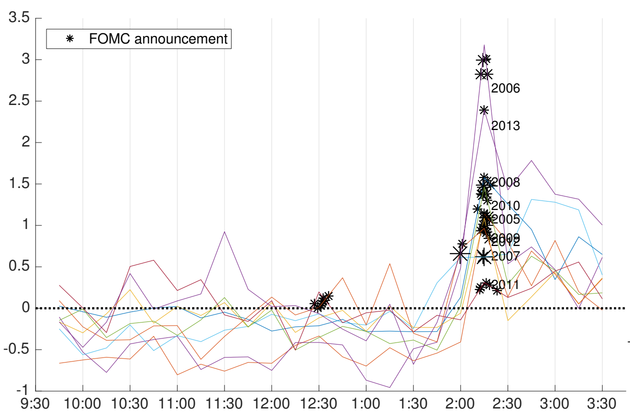

(b) Quarter-Hour Deviations, 1:00–3:00pm

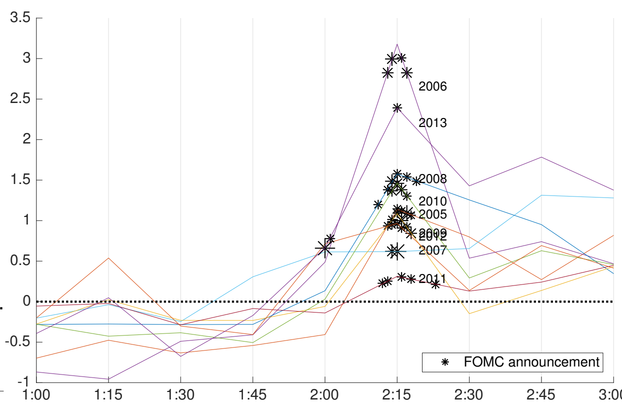

This figure plots intraday tail risks for the market factor for FOMC and non-FOMC announcement dates from 2005–2013. For each quarter-hour interval and calendar year, I compute the average FOMC announcement date tail risk and subtract the average non-FOMC announcement tail risk. I then normalize this 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 right plot zooms in on the 1:00-3:00pm interval.

Figure IX: Anticipated XLF Jumps by Hour, 2004–2013

This figure plots hourly estimated tail risks for financial and market risks in a two-factor market and financials model. The fifteen largest increases in tail risks within a one-month window are overlaid with a green 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 of tail risk news days. Standardized values divide by the full time-series standard deviation. Red events correspond with extreme changes in both factors (using the two-factor model) within 24 hours. Implied financial sector tails on October 10, 2008 are truncated for visual clarity.

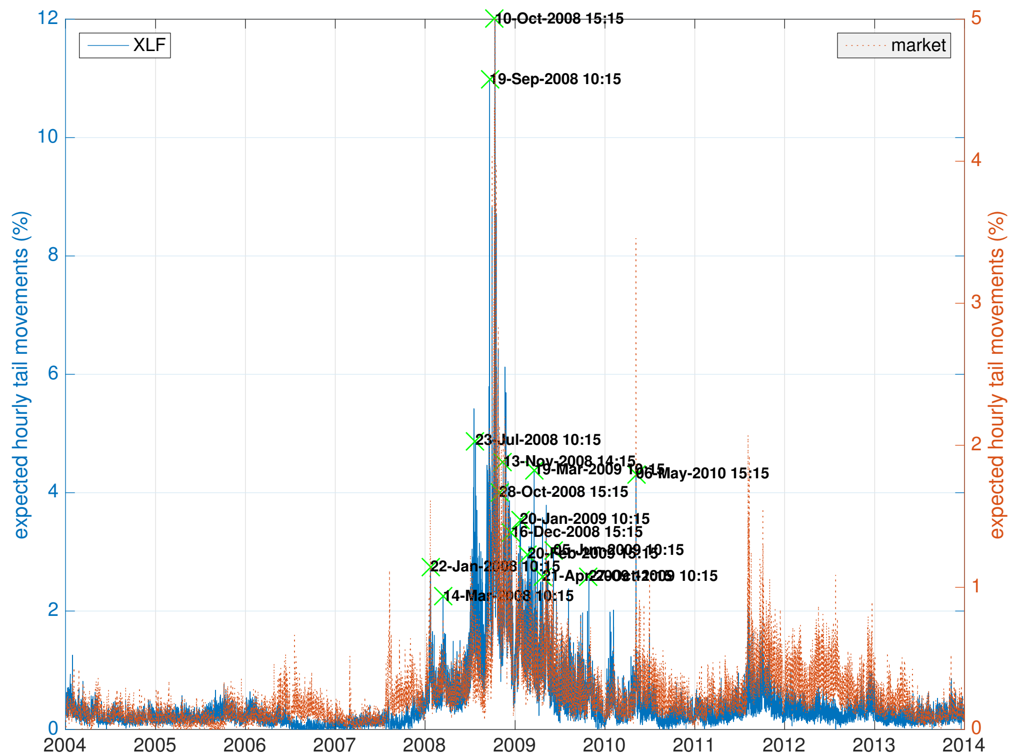

Hour	Std. Value	Event
22-Jan-08 10:15	5.88	Worldwide "Black Monday"
14-Mar-08 10:15	3.48	New York Fed drops bailout deal to save Bear Stearns
23-Jul-08 10:15	4.70	House passes Fannie and Freddie rescue bill after Bush drops opposition
19-Sep-08 10:15	19.89	TARP proposed; short-selling ban; global campaign by central banks
10-Oct-08 10:15	24.13	Stock market crashes in Asia, Europe, and the United States
28-Oct-08 15:15	6.53	First round of TARP bank bailouts ($115 billion)
13-Nov-08 14:15	9.18	Large negative jobless claims surprise; most new claims since 9/11
16-Dec-08 15:15	4.62	FOMC lowers Fed funds target to 0-0.25%
20-Jan-09 10:15	4.27	Discussion of bank recapitalization, bank stocks down 20–25%
20-Feb-09 15:15	5.98	Dodd suggests bank nationalization may be necessary
19-Mar-09 15:15	6.23	FOMC announces $1 trillion in new Treasury bond and MBS purchases
21-Apr-09 11:15	4.19	BofA reports sharp rise in bad loans, financial stocks drop more than 10%
05-Jun-09 10:15	4.84	Rumors circulate on FDIC push to gain greater control over Citigroup
27-Oct-09 10:15	3.92	House Financial Services Committee presents draft "Too Big to Fail" law
06-May-10 15:15	11.17	2010 Flash Crash

Table I: Weekly Correlation of Tail Measure with Other Volatility and Tail Measures

	Spread-Implied Tail	Options-Implied Tail	VIX	Realized Volatility
Spread-Implied Tail	–	0.75	0.85	0.94
Options-Implied Tail	0.75	–	0.88	0.80
VIX	0.85	0.88	–	0.92
Realized Volatility	0.94	0.80	0.92	–

This table reports weekly correlations of tail and volatility measures over the 2004–2011 sample period for which all measures are available. The spread-implied measure uses Equation (16) 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 are the weekly parametric left-tail risk estimates from Figure 7 of Bollerslev and Todorov (2014). All values are centered, equal-weighted averages by week to align with the weekly options-implied estimates.

Table II: Summary Statistics for Tail Realizations

This table presents summary statistics on the distribution of basis-point jumps, spread jumps, and volatility by hour and day for the SPY and XLF ETFs. The count variable sums jumps exceeding 10 basis points or 5 half-spreads, and the sum variables are a weighted average of jump sizes exceeding 5, 10, 25, and 100 basis points or 1, 5, 10, or 25 half-spreads. Realized volatility is the square root of the average squared one-minute SPY returns within each hour multiplied by $\sqrt{390}$ to obtain a daily measure (for market hours).

SPY Hourly	**Basis-Point Jumps**		**Spread Jumps**		**Volatility**
	Jump Count	Jump Sum	Jump Count	Jump Sum	
Mean	2.947	32.649	12.168	77.325	0.793
Std. Dev.	6.546	80.745	10.443	82.206	0.643
Median	0	0	9	49	0.608
Max	53	1310	55	754	11.535
SPY Daily					
Mean	17.627	195.314	72.793	462.596	0.791
Std. Dev.	36.111	444.137	55.160	440.557	0.587
Median	4	40	55	305.5	0.608
Max	273	4960	286	3416	7.636

XLF Hourly	**Basis-Point Jumps**		**Spread Jumps**		**Volatility**
	Jump Count	Jump Sum	Jump Count	Jump Sum	
Mean	6.394	76.610	0.690	3.686	1.181
Std. Dev.	10.735	145.440	2.472	13.972	1.102
Median	1	10	0	0	0.810
Max	59	1670	36	261	15.248
XLF Daily					
Mean	38.366	459.660	4.142	22.114	1.181
Std. Dev.	60.654	819.032	12.357	69.125	1.024
Median	12	120	0	0	0.812
Max	308	6610	144	905	10.293

Table III: Jump Tails and Contemporaneous Market Jumps

This table presents results from a regression of realized jumps against contemporaneous tail risks,

$$tail\_realization_s = \alpha + \beta\xi_{t,MKT} + \gamma VIX_t + \delta CV_t + \epsilon_t.$$

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 average 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. Regressions in the top panel consist of the 2005–2013 sample by trading hour, with one-month rolling HAC standard errors (126 observation bandwidth). Regressions in the bottom panel average all variables within each year and hour of the trading day and use White standard errors. $\xi_{t,MKT}$ is normalized by its standard deviation in both panels.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	SPY Basis-Point Jumps						**SPY Spread Jumps**					
Variable	Jump Count			Jump Sum			Jump Count			Jump Sum		
ξ_{MKT}	5.474***	3.487***	2.520***	69.766***	50.390***	27.123***	7.388***	4.641***	5.681***	65.184***	48.027***	43.946***
	(18.540)	(12.221)	(9.301)	(20.230)	(25.149)	(7.561)	(8.338)	(4.432)	(7.351)	(14.843)	(8.807)	(8.038)
VIX		0.269***	0.261***		2.620***	2.428***		0.371***	0.380***		2.320***	2.286***
		(7.665)	(5.839)		(7.771)	(5.017)		(6.178)	(6.607)		(5.249)	(4.640)
CV			4.144			99.655***			-4.455**			17.482
			(1.510)			(3.015)			(-1.977)			(0.710)
Obs.	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538
R^2	0.699	0.786	0.799	0.747	0.801	0.851	0.501	0.566	0.572	0.629	0.670	0.671

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	SPY Basis-Point Jumps						**SPY Spread Jumps**					
Variable	Jump Count			Jump Sum			Jump Count			Jump Sum		
ξ_{MKT}	3.404***	2.443***	1.005***	40.061***	30.297***	9.113***	5.903***	5.182***	4.318***	46.18***	41.612***	25.204***
	(10.274)	(6.201)	(8.362)	(9.418)	(5.571)	(6.896)	(11.608)	(10.022)	(10.366)	(11.357)	(8.283)	(9.961)
VIX		0.178***	0.130***		1.810**	1.110***		0.134***	0.105**		0.847**	0.305
		(3.217)	(3.727)		(2.632)	(3.261)		(3.652)	(2.651)		(2.131)	(1.193)
CV			14.811***			218.154***			8.890*			168.962***
			(6.701)			(9.814)			(1.684)			(5.005)
Obs.	54	54	54	54	54	54	54	54	54	54	54	54
R^2	0.884	0.941	0.975	0.887	0.930	0.984	0.830	0.840	0.844	0.851	0.858	0.882

t-statistics are given in parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table IV: Beginning-of-Period and Lagged Variable Market Jumps

This table presents results from regressions of realized jumps against beginning-of-period estimated tail risks with and without lagged copies of each variable,

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

corresponding 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 average 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. Regressions consist of the 2005–2013 sample by trading hour, with one-month rolling HAC standard errors (126 observation bandwidth). $\xi_{t,MKT}$ is normalized by its standard deviation in both panels.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	SPY Basis-Point Jumps						**SPY Spread Jumps**					
Variable	Jump Count			Jump Sum			Jump Count			Jump Sum		
ξ_{MKT}	5.250***	3.077***	2.042***	66.298***	43.336***	22.163***	7.174***	4.264***	4.400***	62.307***	42.029***	33.367***
	(22.040)	(17.308)	(11.044)	(14.826)	(14.560)	(8.001)	(10.817)	(5.408)	(6.805)	(19.367)	(10.564)	(7.810)
VIX		0.305***	0.272***		3.226***	2.544***		0.409***	0.413***		2.849***	2.570***
		(9.674)	(6.103)		(9.140)	(5.277)		(8.485)	(7.479)		(7.650)	(5.289)
CV			5.644**			115.504***			-0.746			47.252**
			(2.161)			(3.695)			(-0.349)			(2.010)
Obs.	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538
R^2	0.643	0.764	0.795	0.674	0.763	0.849	0.472	0.557	0.557	0.575	0.641	0.655

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	SPY Basis-Point Jumps						**SPY Spread Jumps**					
Variable	Jump Count			Jump Sum			Jump Count			Jump Sum		
ξ_{MKT}	2.392***	2.040***	1.446***	32.984***	28.300***	16.614***	4.077***	3.381***	3.367***	34.894***	29.483***	23.704***
	(11.213)	(11.722)	(10.625)	(11.050)	(10.418)	(9.180)	(6.347)	(5.493)	(6.328)	(10.426)	(9.250)	(7.081)
VIX		0.811***	0.672***		11.159***	7.536***		1.257***	1.203***		11.564***	9.550***
		(8.598)	(8.105)		(8.864)	(7.324)		(5.072)	(5.834)		(7.358)	(6.633)
CV			4.124*			93.973***			0.748			46.938**
			(1.962)			(3.571)			(0.388)			(2.155)
Obs.	13537	13537	13537	13537	13537	13537	13537	13537	13537	13537	13537	13537
R^2	0.808	0.834	0.849	0.812	0.837	0.882	0.722	0.739	0.740	0.767	0.787	0.798

t-statistics are given in parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Table V: Jump Tails and Contemporaneous Market Jumps with an Adverse Selection Proxy

This table presents results from a regression of realized jumps against contemporaneous tail risks,

$$tail\_realization_t = \alpha + \beta\xi_{t,MKT} + \gamma VIX_t + \delta CV_t + \epsilon_t.$$

Right-hand side tail estimates are constructed using realized adverse selection rather than half-spreads as in Equation (19). 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 average 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. Regressions in the top panel consist of the 2005–2013 sample by trading hour, with one-month rolling HAC standard errors (126 observation bandwidth). Regressions in the bottom panel average all variables within each year and hour of the trading day and use White standard errors. $\xi_{t,MKT}$ is normalized by its standard deviation in both panels.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	SPY Basis-Point Jumps						**SPY Spread Jumps**					
Variable	Jump Count			Jump Sum			Jump Count			Jump Sum		
ξ_{MKT}	4.907***	2.743***	1.723***	63.132***	39.556***	17.850***	6.669***	3.733***	4.162***	60.032***	39.162***	32.891***
	(20.533)	(9.942)	(6.693)	(20.396)	(17.298)	(6.169)	(8.196)	(4.152)	(5.658)	(14.618)	(7.426)	(6.170)
VIX		0.354***	0.315***		3.853***	3.023***		0.480***	0.496***		3.411***	3.171***
		(10.022)	(6.422)		(9.107)	(5.332)		(11.089)	(9.533)		(9.353)	(6.420)
CV			5.372**			114.304***			-2.259			33.022
			(2.009)			(3.650)			(-1.081)			(1.476)
Obs.	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538
R^2	0.562	0.762	0.786	0.612	0.768	0.840	0.408	0.553	0.554	0.533	0.651	0.657

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	SPY Basis-Point Jumps						**SPY Spread Jumps**					
Variable	Jump Count			Jump Sum			Jump Count			Jump Sum		
ξ_{MKT}	3.322***	2.250***	0.948***	39.491***	28.273***	8.489***	5.739***	4.536***	4.194***	45.720***	37.734***	25.590***
	(11.980)	(8.386)	(5.374)	(11.164)	(7.736)	(4.246)	(9.611)	(8.807)	(8.510)	(10.585)	(9.316)	(7.333)
VIX		0.222***	0.160***		2.325***	1.373***		0.249***	0.233***		1.655***	1.071***
		(5.287)	(4.029)		(4.659)	(3.507)		(7.639)	(4.654)		(6.080)	(3.175)
CV			13.824***			210.074***			3.632			128.958***
			(4.870)			(7.210)			(0.630)			(3.078)
Obs.	54	54	54	54	54	54	54	54	54	54	54	54
R^2	0.843	0.953	0.972	0.862	0.949	0.982	0.785	0.828	0.829	0.835	0.866	0.875

t-statistics are given in parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table VI: Jump Tails and Contemporaneous Market Jumps Controlling for PIN

This table presents results from a regression of realized jumps against contemporaneous tail risks,

$$tail\_realization_t = \alpha + \beta\xi_{t,MKT} + \gamma VIX_t + \delta CV_t + \epsilon_t.$$

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 average 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. Regressions in the top panel consist of the 2005–2013 sample by trading hour, with one-month rolling HAC standard errors (126 observation bandwidth). Regressions in the bottom panel average all variables within each year and hour of the trading day and use White standard errors. $\xi_{t,MKT}$ is normalized by its standard deviation.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	SPY Basis-Point Jumps						**SPY Spread Jumps**					
Variable	Jump Count			Jump Sum			Jump Count			Jump Sum		
ξ_{MKT}	5.365***	3.238***	2.064***	67.610***	46.109***	22.030***	7.201***	4.210***	4.301***	62.659***	42.993***	32.808***
	(24.561)	(15.709)	(6.303)	(15.864)	(11.444)	(5.520)	(10.423)	(5.112)	(5.690)	(17.062)	(9.469)	(5.559)
VIX		0.279***	0.260***		2.824***	2.432***		0.393***	0.394***		2.583***	2.417***
		(8.775)	(6.142)		(8.205)	(5.289)		(7.076)	(6.768)		(6.078)	(4.942)
CV			5.678**			116.460***			-0.441			49.257*
			(1.976)			(3.477)			(-0.167)			(1.769)
Obs.	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538	13538
R^2	0.672	0.759	0.789	0.701	0.760	0.844	0.476	0.544	0.544	0.581	0.629	0.643

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	SPY Basis-Point Jumps						**SPY Spread Jumps**					
Variable	Jump Count			Jump Sum			Jump Count			Jump Sum		
ξ_{MKT}	3.383***	2.413***	0.918***	39.702***	29.779***	8.357***	5.793***	5.019***	3.436***	45.151***	40.031***	19.392***
	(9.715)	(6.063)	(6.662)	(8.897)	(5.448)	(5.733)	(11.271)	(9.567)	(6.300)	(10.634)	(7.830)	(5.565)
VIX		0.176***	0.123***		1.806**	1.043***		0.141***	0.084*		0.932**	0.197
		(3.077)	(3.662)		(2.535)	(3.190)		(3.219)	(1.837)		(2.100)	(0.663)
CV			15.907***			227.877***			16.837***			219.559***
			(7.265)			(10.466)			(2.841)			(5.689)
Obs.	54	54	54	54	54	54	54	54	54	54	54	54
R^2	0.874	0.926	0.975	0.871	0.912	0.984	0.800	0.810	0.827	0.814	0.822	0.870

t-statistics are given in parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table VII: Jump Tails and Contemporaneous XLF Jumps

This table presents results from a regression of realized jumps against contemporaneous tail risks,

$$tail\_realization_t = \alpha + \beta \xi_{t,FIN} + \gamma \xi_{t,MKT} + \epsilon_s.$$

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 average 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.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	XLF Basis-Point Jumps				**XLF Spread Jumps**			
Variable	Jump Count		Jump Sum		Jump Count		Jump Sum	
ξ_{FIN}	8.190***	7.566***	119.631***	108.223***	1.558***	1.313***	8.819***	7.150***
	(6.661)	(6.047)	(8.658)	(6.240)	(8.188)	(3.367)	(7.308)	(3.214)
ξ_{MKT}		2.664		48.680		1.046		7.123
		(0.496)		(0.850)		(0.965)		(1.145)
Obs.	13541	13541	13541	13541	13541	13541	13541	13541
R^2	0.581	0.583	0.675	0.678	0.396	0.400	0.398	0.403

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	XLF Basis-Point Jumps				**XLF Spread Jumps**			
Variable	Jump Count		Jump Sum		Jump Count		Jump Sum	
ξ_{FIN}	6.983***	7.540***	91.871***	103.556***	1.150***	1.474***	6.226***	7.966***
	(13.153)	(9.157)	(16.617)	(12.807)	(19.561)	(7.288)	(17.934)	(6.966)
ξ_{MKT}		-3.650		-76.565***		-2.122		-11.402*
		(-1.649)		(-3.331)		(-1.725)		(-1.676)
Obs.	54	54	54	54	54	54	54	54
R^2	0.902	0.903	0.917	0.920	0.617	0.627	0.616	0.626

t-statistics are given in parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Table A.1: Data Filters

	Intraday (hour average)	**Daily**
Beta filter	96.0%	94.2%
Volume filter	83.0%	89.0%
Spread filter	99.9%	99.8%
All filters	80.8%	85.0%
Average sample size	**2776.8**	**3191.1**

A. Data Filters

The TAQ database occasionally records erroneous information about trade prices and quantities. 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 smoothes microstructure noise not of interest for my analysis.

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

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

2. A quoted spread 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. Market beta in $[0.1, 2.5]$.

Securities not satisfying all of these conditions are excluded from cross-sectional regressions. Table A.1 breaks down sample attrition by each filter. The average sample size is approximately $2,800$ distinct stocks for each trading hour and $3,200$ distinct stocks for each trading day.

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 jump tails in the cross section. For example, zero realized volumes in an interval clearly deviate from market makers' expectations. 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, but the underlying model cannot accommodate a log transform of the data to eliminate the potential right skew in the dependent variable. 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 factor jumps being greater than h/β, as required in the linear version of the model (Equation 10). Nevertheless, results are quantitatively almost identical when allowing for significantly narrower ranges of allowed βs.

B. Market-Maker Rationality and Informational Requirements

A. λ Known to Market Makers

Equation (10) leans on the rationality and attentiveness of the collective active market maker. Importantly, no single high-frequency market maker needs to enforce Equation (10). Individual market makers may perceive alternative statistical relationships among securities or no relationships at all. Provided that the product of expected factor tails and risk loadings is unbiased across market makers, the linearity of Equation (10) maintains the conjectured liquidity-provision–risk-pricing relationship by the law of large numbers. In short, for spreads to reflect asset-pricing risks, market makers must be fast and rational, but they need not be omniscient.

Although all market makers need not perceive the same risks, it is useful—but not essential—for market makers to have intimate knowledge of the arrival rate of traders and trading opportunities. Specifically, I assume the following:

Assumption 3. *Market makers know arrival rates λ_{FT} and λ_k for all factors k.*

Assumption 3 simplifies the equilibrium condition of Equation (10) by enabling the market maker to condition on expected volume in setting h. High-frequency market makers specialize in estimating near-term liquidity demands from the order book. By virtue of their fast connections and algorithmic processing, they have the opportunity to exploit a large set of potential signals in real time. The competitiveness of market making combined with the potential costs of misestimating demand contributes to detailed estimates of near-term demand outcomes.

Market Maker Information and Exogeneity

Equations (10) and (14) relate spreads and depth to expected jump arrivals. Under Assumption 3, the parameters λ_{FT} and λ_{jump} that govern expected volume and its composition are known to market makers. Because the conditional average volume $\bar{V}(\cdot; h, d)$ is in the market maker's information set, the competitive market-making sector sets spreads and depth such that jump risk costs exactly offset intermediation gains. From Equation (14), the market making sector therefore solves

$$\left(\frac{h}{d}\right)^* = \frac{\sum_k \xi_k \beta_k + \tilde{\xi}}{\lambda_{FT} q^* + \lambda_{jump} d} = \frac{\sum_k \xi_k \beta_k + \tilde{\xi}}{\bar{V}\left(\cdot; h, d\right)}. \tag{B.1}$$

Assumption 3 is useful because taking expectations over $1/\bar{V}\left(\cdot; h, d\right)$ otherwise introduces a concavity correction term for uncertainty in traders' arrival rate. The next section relaxes this assumption and discusses appropriate modifications to the estimation technique.

Market-maker liquidity prices h and quantities d are a function only of primitive economic risks and known expected volume. Multiplying both sides of the expression by $\bar{V}$ recovers expected intermediation revenues per unit of depth and maintains the left-hand side's dependence only on exogenous jump risks:

$$\left(\frac{h}{d}\right)^* \bar{V} = \sum_k \xi_k \beta_k + \tilde{\xi}. \tag{B.2}$$

B. λ Unknown to Market Makers

I now relax the assumption that market makers know near-term trader arrival rates with certainty (Assumption 3). Suppose instead that the market maker does not know trader arrival rates but observes with error the average potential realization of near-term volume. In particular, let the true average volume be conditionally lognormally distributed given a common signal s observed by market makers, where the disturbance term is orthogonal to perceived jump tail risks:

$$\log \bar{V} = s + \nu, \ \ \nu \sim N\left(0, \sigma_\nu^2\right). \tag{B.3}$$

This signal structure ensures a positive arrival rate and generates a empirically plausible conditionally lognormal distribution for volume.

The competitive market making sector sets expected intermediation profits equal to zero. Solving for the equilibrium spread obtains

$$\left(\frac{h}{d}\right)^* = \left(\sum_k \xi_k \beta_k + \tilde{\xi}\right) \frac{1}{E\left[\bar{V}\right]} = \left(\sum_k \xi_k \beta_k + \tilde{\xi}\right) \exp\left(-s - \frac{1}{2}\sigma_\nu^2\right), \tag{B.4}$$

where the second equality follows from the properties of the lognormal distribution. Multiplying both sides by the true average $\bar{V}$ obtains

$$\left(\frac{h}{d}\right)^* \bar{V} = \left(\sum_k \xi_k \beta_k + \tilde{\xi}\right) \frac{\bar{V}}{\exp\left(s + \frac{1}{2}\sigma_\nu^2\right)} = \left(\sum_k \xi_k \beta_k + \tilde{\xi}\right) \exp\left(\epsilon - \frac{1}{2}\sigma_\nu^2\right). \tag{B.5}$$

As in the known $\bar{V}$ case, using realized volume in place of expected volume contributes to measurement error on the left-hand side and does not affect ξ_k coefficient estimates asymptotically. However, the multiplicative scaling terms are potentially problematic for two reasons. First, uncertainty in average volume forecasts may vary over time. Implied tail risk coefficients are meaningful if σ_ν^2 is constant, but time variation in the precision of volume forecasting distorts estimates of relative levels of tail risk. Second, the precision of volume forecasts may differ across assets.

Perhaps surprisingly, obtaining correct tail risk coefficients up to a scaling term is still achievable for each cross-section under plausible independence assumptions. The image processing literature offers several solutions to the multiplicative noise or "speckle filtering" problem. For example, one such method, median filtering, is readily applied in the cross-sectional regression context under a suitable distance metric on asset betas. Left-hand side values not similar to those of "near" neighbors are excluded under this filter, as abnormally large local deviations are attributed to multiplicative noise.

With this filter in mind, Assumption 4 is a sufficient auxiliary condition for recovery of tail risks with market maker uncertainty about the true trader arrival rate:

Assumption 4. *The average uncertainty $\overline{\sigma_\nu^2}$ is constant over time, and σ_ν^2 is locally independent of factor loadings within each cross-section.*

The first component of Assumption 4 is tenable if HFT market makers' ability to predict future volume has been roughly constant over the sample period. The second component of the assumption holds if factor loadings do not have first-order effects on market makers' ability to forecast order flow.

C. Additional Empirical Considerations

Negative Betas and Return Symmetry

This methodology can accommodate negative asset betas. For illustration, suppose that the additional candidate factor is FIN or "financials," and that financial betas can be of either sign. If jump tails are symmetric, estimating expected jumps entails minimal modification of our previous expressions,

$$\begin{aligned}\left(\frac{Vh}{d}\right)_{it} &= \tilde{\xi} + \xi_{MKT}\beta_{im} + \xi_{FIN}\beta_{i,FIN}\mathbf{1}\left(\beta_{i,FIN} > 0\right) - \xi_{FIN}\beta_{i,FIN}\mathbf{1}\left(\beta_{i,FIN} < 0\right) + \epsilon_i \\ &= \tilde{\xi} + \xi_{MKT}\beta_{im} + \xi_{FIN}\left|\beta_{i,FIN}\right| + \epsilon_i. \qquad \text{(C.1)}\end{aligned}$$

The first line acknowledges that negative betas only increase half-spreads for buying the asset when the factor's jump return is negative. Collapsing the expression in the first line suggests that taking the absolute value of beta suffices for estimation.

I implicitly assume symmetry of jumps under $\mathbb{P}$ because the distribution of realized jumps for individual stocks and the SPY and XLF ETFs is very close to symmetric.[38] This condition can be weakened further: the density need not be symmetric around zero for every jump size. It is sufficient that the tail expectation of jump sizes above $\bar{h}_k$ is equal for positive and negative factor moves. Alternatively, Equation (C.1) can be generalized to allow for asymmetric tail expectations:

$$\left(\frac{Vh}{d}\right)^R = \tilde{\xi}^R + \xi^R_{MKT}\beta_{i,MKT} + \xi^R_{FIN}\beta_{i,FIN}\mathbf{1}_{\beta_{i,FIN}>0} - \xi^L_{FIN}\beta_{i,FIN}\mathbf{1}_{\beta_{i,FIN}<0} + \epsilon_{i1}, \quad \text{(C.2)}$$

$$\left(\frac{Vh}{d}\right)^L = \tilde{\xi}^L + \xi^L_{MKT}\beta_{i,MKT} - \xi^L_{FIN}\beta_{i,FIN}\mathbf{1}_{\beta_{i,FIN}<0} + \xi^R_{FIN}\beta_{i,FIN}\mathbf{1}_{\beta_{i,FIN}>0} + \epsilon_{i2}, \quad \text{(C.3)}$$

where R and L denote right- and left-jump tails and the corresponding volume and spread proxies. These expressions are readily stacked to obtain a single regression equation in the unknown left and right tail coefficients. I estimate left- and right-jump tails separately in the Internet Appendix and confirm that jump tails are typically (nearly) symmetric, but symmetry breaks down around extreme events such as the 2010 Flash Crash.

Co-jumps

I now revisit assumption 2 in the context of Equations (7) and (10). Excluding co-jumps eliminates terms in Equation (7) 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

$$+\lambda_{\{MKT,FIN\}}\int_{\beta_{MKT}r_{MKT}+\beta_{FIN}r_{FIN}\geq h}^{\infty}\left(\beta_{MKT}r_{MKT}+\beta_{FIN}r_{FIN}-h\right)d\left(r_{MKT},r_{FIN}\right). \quad \text{(C.4)}$$

This additional term is readily converted into linear terms under the large jumps assumption of Equation (10) if jumps are of the same sign. Under these conditions, the additional term in Equation (10) is decomposed as

$$\begin{aligned}&-\lambda_{\{MKT,FIN\}} + \lambda_{\{MKT,FIN\}}E\left[r_{MKT}|r_{MKT},r_{FIN}>0\right]\beta_{MKT}\\&\quad+\lambda_{\{MKT,FIN\}}E\left[r_{FIN}|r_{MKT},r_{FIN}>0\right]\beta_{FIN}.\end{aligned} \quad \text{(C.5)}$$

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

[38] Jumps detected using Lee and Mykland (2008) and Bollerslev et al. (2013) methodologies share this symmetry property. Table III of Bollerslev and Todorov (2011b) also confirms a single stochastic factor for left- and right-jump tails under $\mathbb{P}$.

$\lambda_{\{MKT,FIN\}} \times E\left[r_{MKT}|r_{MKT}, r_{FIN} > 0\right]$ and $\lambda_{\{MKT,FIN\}} \times E\left[r_{FIN}|r_{MKT}, r_{FIN} > 0\right]$ is desirable. The total coefficient on β_{MKT} is the tail risk of the market with or without financial co-jumps, which arguably is of greater interest than either component of market tail risk independently.

Liquidity Maker and Taker Fees

I add 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 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 (2015)). I omit consideration of taker fees. Liquidity taker fees drop out in the linear model of Equation (10) because they are subsumed into the lower bound of integration for each factor.

Because rebates are typically small relative to the spread, reasonable alternatives for the level of the rebate have minimal effect on results. For example, reducing the rebate from \$0.0022 per share to zero maintains the same shape in the time series of jumps while shifting the recovered tail risk down imperceptibly. In addition, rebate measurement errors contaminate the dependent variable and are likely to be uncorrelated with asset betas, so rebate mismeasurement should have minimal effect on factor tail risk point estimates.