

Option-Implied Information and Predictability of Extreme Returns*

Grigory Vilkov[‡]

Yan Xiao[§]

This version: September 24, 2012

We study whether option-implied conditional expectation of market loss due to tail events, or *tail loss measure*, contains information about future returns, especially the negative ones. Our tail loss measure predicts future market returns, magnitude, and probability of the market crashes, beyond and above other option-implied variables. Stock-specific tail loss measure predicts individual expected returns and magnitude of realized stock-specific crashes in the cross-section of stocks. An investor, especially the one who cares about the left tail of her wealth distribution (e.g., disappointment-averse), benefits from using the tail loss measure as an information variable to construct managed portfolios of a risk-free asset and market index. The tail loss measure is motivated by the results of the extreme value theory, and it is computed from observed prices of out-of-the-money put as the risk-neutral expected value of a loss beyond a given relative threshold.

Keywords: extreme value theory, tail measure, implied correlation, variance risk premium, option-implied distribution, predictability, portfolio optimization

JEL: G11, G12, G13, G17

*We are grateful for a number of critical comments and stimulating suggestions from Raman Uppal.

[‡]Goethe University Frankfurt, Finance Department, Grüneburgplatz 1, D-60323 Frankfurt am Main, Germany; email: vilkov@vilkov.net.

[§]Graduate Program “Finance and Monetary Economics,” Goethe University Frankfurt.

1 Introduction

The question is not “if the market breaks”; the question is “when the market breaks”. We all know the market is going to correct some day. On the other hand “top picking” is a treacherous pursuit. I have been amazed at the way the stock market has withstood an onslaught of negativity - all on the assumption that the Federal Reserve is guaranteeing a bull market, PERIOD... Don’t believe it. This market is going to correct. The question is “when” and “by how much”.

www.seekingalpha.com, September 20, 2012

The level of the market index in the future is random, and its distribution can be viewed by each investor differently. Market prices of options reflect common assessment of the probability distribution of the underlying asset on the expiration day, adjusted for the investors’ tolerance for bearing risk. Researchers have long recognized the possibility of inferring such forward-looking probability assessments from option prices to improve the prediction of future probability distribution; however, so far we have not been very successful in predicting market and stock crashes, in terms of both their magnitude and timing. Several approaches have been proposed for forecasting low-probability tail events based on the option-implied moments of returns and historical variables.

In this paper, instead of using moments of returns to predict a crash magnitude, we suggest a forward-looking variable that directly quantifies the expected market or stock crash size, conditional on its realization. With the help of this variable we make an attempt to predict the magnitude and probability of a crash in the market index and in the cross-section of stocks. The main contribution of our work is estimating a market-wide and stock-specific tail loss measure (TLM) from the observed prices of traded out-of-the money put options and showing that its *ex ante* magnitude is positively related to the magnitude of the future realized crashes, and that, when no crash occurs, the expectation of a higher loss due to the sudden decline in asset value is associated with a higher premium to be earned for taking the risk.

Tail loss measure contains an incremental information about the future stock returns beyond variance risk premium (VRP) and implied correlation (IC), which have been

shown to predict future market returns.¹ Even though the tail loss measure is significantly related to the probability of a market crash, and though it also explains the magnitude of the realized market losses, in the out-of-sample portfolio optimization exercise, we mostly benefit from its ability to predict risk premium, and not the crashes. It is worth mentioning that it is extremely hard to predict the timing and magnitude of a low-probability market crash, partly because any crash is by definition rare, and even having a high probability of happening it may not come to a realization soon enough.

We proceed in several steps. First, using the results of the Extreme Value Theory (EVT), we define and show how to estimate the tail loss measure from observed out-of-the-money option prices. The theoretical advantage of the tail loss measure compared to many other option-implied quantities is that it is unique for a given minimal size of a crash. Second, we perform the time-series tests, using the robust regression approach and “crash” dummies to show that one standard deviation increase in the tail loss measure causes a marginal positive change in the weekly market return of 0.29% (or 15% p.a.), and that if the crash occurs, then one standard deviation difference in tail loss measure amounts to -1.35% per week for a one-sigma event, and -1.84% for a two-sigma event. A similar direction of the results holds for monthly and quarterly horizons. The binary probit model estimation indicates that the weekly change in tail loss measure is positively related to the probability of a crash (with a significant relation for 1-week and 1-month horizons), though the effect is not economically important. Third, we carry out an out-of-sample portfolio optimization exercise, using option-implied variables and their combinations to construct “*managed portfolios*, in which one invests more or less in an asset according to some signal,” as suggested in Cochrane (2005, page 20).² We compare the performance of the option-implied variables to that of a non-informative portfolio, and the only variable that significantly adds useful information for optimization (above a non-informative case) is the implied correlation. The tail loss measure beats the non-informative portfolio in economic terms for most cases, but the difference is not statistically significant. Finally, we study the cross-section of stock returns and stock-specific tail loss measures, and we come to similar conclusions as in the time-series

¹See Bollerslev, Tauchen, and Zhou (2009), and Driessen, Maenhout, and Vilkov (2012) for details.

²The approach is also similar to the parametric portfolio policy of Brandt, Santa-Clara, and Valkanov (2009), but applied in the time-series settings.

analysis for the market. Tail loss measure is positively related to the expected return and to the magnitude of a crash in the case of its realization.

We conclude this introduction by discussing the relation of our work to the existing literature. The idea that option prices contain information about future distribution of asset returns has been understood ever since the work of Black and Scholes (1973) and Merton (1973); Christoffersen, Jacobs, and Chang (2011) provide a splendid and comprehensive survey of this literature. Using implied moments is helpful in predicting future volatility, correlations, expected returns of the market, as well as in the cross-section of stocks, and other attributes of return distribution, however, only a few studies are somewhat successful in predicting the probability and magnitude of future crashes. Most authors use a proxy of implied skewness as a major instrument for analysis. For example, Carson, Doran, and Peterson (2006) and Doran, Peterson, and Tarrant (2007) find that option-implied skew has significant forecast power for assessing the degree of market crash risk, though the effect is not economically relevant; Bates (1991) finds no strong crash fears as expressed by the skewness of the risk-neutral distribution during two months immediately preceding the crash. We include model-free implied skewness as the explanatory variable in all types of tests and do not find any overwhelming ability of the currently observed implied skewness to predict future crashes.

We rely in our derivations on the results of the Extreme Value Theory, and especially on the suggestion of Hamidieh (2011) to use out-of-the money options to infer the parameters of the excess loss distribution. Hamidieh uses only one parameter in his analysis, namely, the shape of the tail distribution, and documents a thickening of the left tail of the S&P 500 index just prior to the market turmoil of September 2008. Instead, we are using fully parametrized excess loss distribution to compute the expected value of the asset value loss beyond a certain threshold and relate it to the magnitude and probability of market-wide and stock-specific crashes. Another study that relies on the Extreme Value Theory to derive a related tail measure is Kelly (2011). He suggests a market-wide measure of tail risk that can be estimated from the cross-section of stock returns, exploiting stock-specific crashes every month to identify common fluctuations in tail risk across stocks. The results of his and our studies are consistent in direction and in

magnitude, however, our tail loss measure is forward-looking and can be estimated in real time by simply processing current option prices; hence, we can react to new information more rapidly.

A number of papers look at the relation among variance risk premium, fear of jumps, and future returns. Bollerslev, Tauchen, and Zhou (2009) show that the variance risk premium has a strong predictive power for future market returns, and from Bollerslev and Todorov (2011) we know that “the compensation for rare events accounts for a large fraction of the average equity and variance risk premia.” Implied correlation is also related to the risk of negative jump through the notion of a loss of diversification, and it also has been shown to both predict future returns and bear a systematic negative risk premium (Driessen, Maenhout, and Vilkov (2009, 2012)). Our results indicate that neither variance risk premium nor the implied correlation contain all the “fearful” information about the future crashes, and tail loss measure proves to be useful in identifying crash expectations and associated risk premium in asset returns.

Some studies derive option-based systematic factors related to the risk of market crashes and then empirically verify the sign of the associated risk premiums. For example, Chang, Christoffersen, and Jacobs (2012) propose a market skewness factor and show that it commands a negative risk premium. Cremers, Halling, and Weinbaum (2012) derive an option-based factor, which is sensitive to the jump (crash) risk only, and show that it also bears very significant negative risk premiums. The significant negative risk premium for both market skewness and jump factors is consistent with the positive relation between the tail loss measure and future returns in the case that no crash occurs, however, we remain agnostic about the systematic nature of the jump risk and its pricing, and rather we suggest a variable, which is *ex ante* related to crash (negative jump) size and its probability.

The last, but not least, strand of literature deals with the theoretical foundations for risk premiums related to fears of crash. Bates (2008) derives an equilibrium with options markets with heterogenous crash-averse investors and finds that “specification of crash aversion is compatible with the static option pricing puzzles, while heterogeneity partially explains the dynamic puzzles.” Gabaix (2012) develops a rare-disasters model

and successfully resolves several option-related puzzles on “high price of deep out-of-the-money puts; and ... high put prices being followed by high stock returns.” The latter result is directly related to our findings; moreover, we provide additional empirical support for it, using not just prices of puts but also the quantification of the expected market losses conditional on a crash.

The rest of the paper is organized as follows. In Section 2, we discuss the construction of the left tail loss measure and its economic interpretation, and we also specify the predictions to be tested later on. In Section 3, we describe the data on stocks and options that we use. In Section 4, we proceed with an empirical analysis, first studying the time-series predictability of market returns, then carrying out an out-of-sample portfolio optimization exercise and a cross-sectional analysis of individual stocks. Our conclusion is found in Section 5, and we collect the proofs and the details about the estimation of the tail loss measure in Appendix A.

2 Tail Loss Measure

In this section, we first discuss the theoretical foundation for constructing the tail loss measure; then we provide its economic interpretation and formulate three testable hypotheses for analysis.

2.1 Constructing the Tail Loss Measure

Instead of looking at the full distribution of asset price changes, we concentrate on its left tail only, in the form of the forward-looking expectation of the drop in asset price below a certain threshold conditional on a crash occurring within a short period of time. Such a tail loss measure (TLM) characterizes not the probability, but the severity of a crash in the case of its realization. To construct the tail loss measure we employ the results of the Extreme Value Theory (EVT), which mainly deals with the distributional properties of the extreme, or low probability, events. Without knowing the true underlying distribution of the asset value, we can utilize (under some technical conditions) the *second theorem* of EVT, also known as the Pickands-Balkema-de Haan theorem, to describe the distribution

of the asset value x below an extreme (low) enough threshold h by a generalized Pareto distribution:

$$G_{\beta,\xi}(h-x) = \begin{cases} 1 - \left(1 + \xi \frac{h-x}{\beta}\right)^{-1/\xi} & \text{if } \xi \neq 0 \\ 1 - \exp\left(-\frac{h-x}{\beta}\right) & \text{if } \xi = 0, \end{cases}$$

where β is the scale, and ξ is the tail shape parameter.³ The mean excess tail value of the asset, provided $\xi < 1$, can be computed as

$$E(h-x|h > x) = \frac{\beta}{1-\xi},$$

and we define the tail loss measure $\Omega_{h,t}$ at time t conditional on the threshold h as the expected excess tail value of an asset relative to the current value of the asset:

$$\Omega_h \equiv \frac{E_t(h-x|h > x)}{x_t}.$$

Following Hamidieh (2011), we infer the parameters β and ξ of the tail value distribution of an asset from the observed prices of options on this asset, and hence all the expectations, as well as the tail loss measure itself, are defined under the risk-neutral probability measure. Refer to Appendix A for derivations and description of estimation procedures.

Economically, our tail loss measure is equivalent to the Conditional Value-at-Risk, computed under the risk-neutral probability measure, and in terms of information content, our tail loss measure is somewhat comparable to risk-neutral skewness or other option-implied tail measures quantifying the fatness of the left tail of the asset value distribution. The advantage of TLM over value-at-risk is that it is forward-looking and inferred from the current prices of traded options without restrictive assumptions about the distribution of an asset. Compared to the other option-implied measures, TLM also has a significant advantage: it is unique, and even though we have sparse strikes of options traded at any given time, we are able to identify the tail loss measure uniquely. In other words, the tail loss measure resolves two problems common to the most option-implied measures: a mathematical one that the inverse problem of estimating the risk-neutral

³Please see the proofs related to the EVT results and generalized Pareto distribution properties in the original papers of Balkema and Haan (1974), and Pickand (1975), or the review material in Section 7 of McNeil, Frey, and Embrechts (2005).

moments is ill-posed, and a economic one that, while in incomplete markets there exist multiple risk-neutral asset distributions and hence theoretically there may be multiple valid estimations of non-traded implied moments, there is only one TLM for a given threshold.

Theorem 1 *Assuming that the risk-neutral distribution of asset value x belongs to the maximum domain of attraction of an extreme value distribution, the tail loss measure $\Omega_h \equiv \frac{E(h-x|h>x)}{x}$, inferred from the prices of traded options on asset x , with threshold h near the endpoint of x , is unique.*

2.2 Economic Interpretation and Testable Hypotheses

Because economic agents dislike losses and negative skewness, they require additional risk premium for holding unfavorable assets, e.g., the market index in unfavorable times of high probability of a market decline (as follows from Barberis and Huang (2008), Brunnermeier and Parker (2005), and Brunnermeier, Gollier, and Parker (2007), among others). Hence most of the time we expect a high tail loss measure to correspond to the high equity risk premium, or high market return. However, because TLM also reflects the expectations of the investors about the magnitude of future crashes, when these expectations come true, we expect to document a positive relation between the TLM and the magnitude of (now negative) asset returns. Thus, when no crash occurs, the relation between TLM and the market return resembles the dependence between the level of risk and the insurance premium that we want to receive to hold the risk. When the crash occurs, i.e., the risk is being realized, we are expected to pay out to reimburse the insured side of the transaction, and the payment is typically higher, the higher the premium was received. This observation leads to our first testable set of predictions.

Hypothesis 1 *When no market crash occurs, the market tail loss measure $\Omega_{h,t}^M$ is positively related to the future market return $r_{t,t+\Delta t}^M$, i.e., it is associated with a high expected market risk premium. In the case that the low-probability crash is realized, the tail loss measure is positively related to the magnitude of the negative realized market return.*

Because the tail loss measure quantifies the severity of a crash as seen by investors, we also expect the TLM to be helpful in identifying the probability of a crash, even though an expected crash with a high probability may still remain unrealized. This particular feature of rare events complicates the estimation of any related expected measure empirically, but based on theoretical grounds we can expect the following:

Hypothesis 2 *The dynamics of the market tail loss measure are positively related to the probability of a low-probability market crash in the future.*

Being aware of a steadily growing number of option-based measures quantifying the forward-looking expectations of the market distribution,⁴ we should be careful not to produce another redundant measure that is especially sensitive to extreme negative returns, i.e., crashes. We will therefore include in all tests the previously derived robust predictors of market returns, such as implied correlation, market variance risk premium, and implied skewness measures. Before we continue, we should fill in a gap in the discussion, namely, the proper definition of a low-probability market crash. In the empirical section we define a low-probability market crash as a negative return larger in magnitude than a predetermined multiple of the expected market volatility. To conform with reality we use the common “fear index” VIX, smoothed over a short historical period as the expected market volatility.⁵

If we fail to reject the predictions regarding the relation of the tail loss measure and the market return dynamics, we can expect to use the return- and crash-predicting abilities of the TLM in portfolio optimization. We will study these questions in empirical sections, though we do not stipulate any testable hypotheses with respect to this as of now.

In addition to studying the market tail loss measure, we also compute the same quantity for the individual stocks and test how the individual TLMs are related to the cross-section of expected stock returns. We study what is the direction and magnitude of the cross-sectional relation between stock returns and TLMs. Motivated by the finding in Kelly (2011), who derived the backward-looking market tail measure under the true probability and found that stocks with higher tail exposure earn higher returns, we also expect our option-implied TLM to be positively related to stock returns. However, as in the case with the market TLM analysis above, we want to differentiate a situation when there is a high probability of a negative return, and when this negative return is actually realized. Thus, our third set of testable predictions is as follows:

⁴To name just a few, Driessen, Maenhout, and Vilkov (2012) study implied correlation; Bollerslev, Tauchen, and Zhou (2009) look at the variance risk premium, and Han (2008) links implied skewness and investor sentiment.

⁵It has been shown in several studies that the implied volatility can serve as a good predictor of the future volatility. See Poon and Granger (2003) for a literature review.

Hypothesis 3 *Individual stock tail loss measure is positively related to the cross-section of expected stock returns, when there is no stock-specific low-probability crash realized. In case the stock-specific low-probability crash is realized, the stock tail loss measure is positively related to the magnitude of the negative realized stock return.*

We devote Section 4.1 to testing the first two hypotheses, and we proceed in Section 4.2 with a portfolio optimization exercise; then we study the cross-sectional predictions in Section 4.3.

3 Data Description

In this section, we describe the data on stocks and stock options that we use in our study. Our sample period is from January 1996 to October 2010. Our data on stocks are from the Center for Research in Security Prices (CRSP) and NYSE's Trades-and-Quotes (TAQ) database. To construct some stock-specific characteristics we also use data from Compustat; the data for the index and individual stock options are from OptionMetrics.

3.1 Stock and Market Returns and Characteristics

We study stocks that are in the S&P 500 index at any time during our sample period. The *monthly* stock returns of the S&P 500 constituents and the index itself are from the monthly CRSP Stock Prices file. We use the S&P 500 prices to construct the overlapping returns for a number of periods (one week; one, two, and three months). We also use high-frequency *intraday* transaction price data for an exchange traded index proxy (SPY) to compute the market variance; these data are from the NYSE's Trades-and-Quotes database. We apply a number of commonly used filters to the high-frequency data and use the second best variance estimator with the resampling and averaging of Zhang, Mykland, and Ait-Sahalia (2005), with the resampling frequency of 15 minutes to compute on each day t the realized variance (RV_t^2) over the last 21 trading days.⁶

We also compute a number of stock-related characteristics. The market value of equity (SIZE) is defined by the natural logarithm of stock price per share multiplied

⁶See, for example, Barndorff-Nielsen, Hansen, Lunde, and Shephard (2009) for an overview of the high-frequency filters.

by shares outstanding; both variables are obtained from the CRSP database. Book-to-market ratio (BTM) is computed as the book value divided by the market capitalization. At the beginning of July of each year, we compute a book-to-market ratio using the market value of equity at the end of December of the preceding calendar year and the book value for the firm’s latest fiscal year ending in the previous calendar year. Book value is calculated as the sum of total common/ordinary equity and the balance-sheet deferred taxes. The twelve-month momentum (MOM) is measured for the end of each month t using monthly return data from CRSP as the cumulative return from month $t - 12$ to month $t - 1$, and one-month reversal (REV) is also computed for the end of each month as the past month return.

3.2 Option Returns and Option-Implied Characteristics

For stock and index options we use IvyDB (OptionMetrics), which contains data on all U.S.-listed index and equity options.

To compute the option-implied measures we use the standardized volatilities for the maturity of 30 days from the Volatility Surface File, which contains a smoothed implied-volatility surface for a range of standard maturities and a set of option delta points. We select out-of-the-money implied volatilities for calls and puts (we take calls with deltas smaller or equal to 0.5, and puts with deltas bigger than -0.5) to estimate on each day the model-free implied variance (MFIV), skewness (MFIS), and kurtosis (MFIK), using the method of Bakshi, Kapadia, and Madan (2003).⁷ We use the MFIV for index and individual options, as well as index weights, to compute for each day the forward-looking implied correlation (IC), applying the method of Driessen, Maenhout, and Vilkov (2012). We also use the MFIV as a proxy for VIX² throughout the paper. Following Bollerslev, Tauchen, and Zhou (2009), we compute the value of the variance risk premium (VRP) of

⁷To compute the integrals that give the required values of the variance, cubic, and quartic contracts precisely, we need a continuum of option prices. We discretize the respective integrals and approximate them using the available options. We normally have 13 out-of-the-money call and put implied volatilities from the surface data for each maturity. Using cubic splines, we interpolate them inside the available moneyness range, and extrapolate using the last known (boundary for each side) value to fill in a total of 1001 grid points in the moneyness range from 1/3 to 3. Then we calculate the option prices from the interpolated volatilities using the known interest rate for a given maturity and use these prices to compute the model-free implied volatility, skewness, and kurtosis.

the index proxy, defined as the difference between the VIX² and realized index variance RV², computed over the last 21 days from the high-frequency transactions data.

For the tail loss measure we first estimate on each day t the tail shape ξ_t and scaling β_t parameters of the generalized Pareto distribution under the risk-neutral probability measure by solving the optimization problem for all out-of-the-money puts with strikes weakly smaller than $K_{0,t}$:

$$\{\xi_t, \beta_t\} = \arg \min_{\xi, \beta \in (\mathbb{R} \times \mathbb{R}^+)} \sum_{i=0}^{n-1} \left| \frac{P_{i,t} - P_{i,t}^*}{P_{i,t}^*} \right|,$$

where $P_{i,t}$ is the observed price of a put option with strike $K_{i,t}$ on day t , and $P_{i,t}^*$ is the “theoretical” price of a put option with the corresponding strike and maturity, calculated by the option pricing formula based on the Extreme Value Theory:

$$P_{i,t}^* = P_{0,t}^* \left[\frac{\xi_t}{\beta_t} (K_{0,t} - K_{i,t}) + 1 \right]^{1-1/\xi_t}. \quad (2)$$

The strike $K_{0,t}$ is the threshold strike, and we assume that for all put options below this strike we can apply equation (2) to compute the put option prices, from the observed option price P_0^* with the threshold strike. We select the threshold strike as current stock price S_t adjusted downwards by the two monthly volatilities, where volatility is approximated as the average value of implied volatility (VIX) over the last quarter, i.e., $K_{0,t} = S_t \times (1 - 2 \times \frac{\widehat{VIX}_t/100}{\sqrt{12}})$ with $\widehat{VIX}_t = \frac{1}{63} \sum_{i=0}^{62} VIX_{t-i}$. For derivations of the TLM and its estimation routine, refer to Appendix A.

4 Empirical Investigation

In this section, we provide the empirical tests of the hypotheses stipulated earlier. We start, in Section 4.1, by studying the predictability of time-series market returns and magnitude of crashes in the case of their realization, using as predictors the TLM and a set of other option-implied variables; we also analyze the predictability of the crash probability. In Section 4.2, we perform an out-of-sample portfolio optimization routine, then provide the details of the cross-sectional analysis in Section 4.3.

4.1 Market TLM and Time-Series Return Predictability

4.1.1 Time-Series Explanatory Regressions

To test Hypothesis 1 we simply run a time-series explanatory regression of the market return on a number of variables. However, taking into account the nature of the study, namely our assumption about the existence of low-probability events with extremely negative returns without the exactly offsetting events with extremely positive returns, we should *a priori* reject a linear model, because simple OLS estimators can be heavily influenced by a small fraction of anomalous observations in the data (see Huber (2005) for a review). To be able to squeeze the data into the desired linear model design, we need somehow to identify the outliers to treat them and “normal” data separately. To stay scientific and avoid arbitrary outlier removal or winsorization procedures, we use robust M-estimators, where, in effect, the data are down-weighted depending on their distance from the median absolute value (the estimator with high resistance to outliers), and extreme outliers are given zero weight. We apply the biweight (also called bisquare) data transformation, which combines the properties of resistance with relatively high efficiency in many applications (see Mosteller and Tukey (1977) for a general discussion and details).⁸

Thus, we consider the linear explanatory model

$$y_{t+\Delta t} = \beta_0 + \beta_1 x_{t1} + \dots + \beta_k x_{tk} + \varepsilon_t = \mathbf{x}'_t \beta + \varepsilon_t, \quad (3)$$

where $y_{t+\Delta}$ is the market holding period return from day t to $t + \Delta t$, \mathbf{x}_t is the set of regressors observed at time t , and β are the respective estimated coefficients. As predictors we use all option-based variables that have been shown elsewhere to predict market returns.⁹

⁸Bisquare weighting function is one of the standard choices for robust regression preprogrammed in the most popular statistical packages. We use the default settings for computing the bisquare weights with the tuning constant value of 4.685, which produce 95% efficiency when the errors are normal, and still offer protection against outliers.

⁹We also included in the set of predictors the market fundamentals—earnings growth rate and price-dividend ratio, following the sum-of-the-parts (SOP) method of Ferreira and Santa-Clara (2011), but the improvement over the option-based predictors for short-term returns is marginal. These results are available upon request.

We first study how the TLM is related to the future market returns in the states, in which no “crash” occurs, and, following our discussion above, we expect to receive a premium for holding the market risk. The results of the estimation are provided in Table 1, and the sign of the TLM coefficient is positive across all models, as expected, though it is significant only for a very short horizon return of one week. Thus, the higher the market losses anticipated by investors are (and adjusted for the willingness to bear the risks), the higher the market return is in the next period, but only for the periods without crash realizations. We also compute in Table 2 the marginal effect of the one-volatility change in each variable, keeping all other variables constant, on the respective expected return. The volatility of the market TLM in our sample equals to 1.88%, and hence, a marginal change in expected short-term market return from the change in TLM by one volatility value will be 0.29% per week, which is 15% per year. Using TLM as the only predictor gives us the \bar{R} of 1.34% for weekly and 3.01% for quarterly returns, while including the other option-based variables boosts it to about 2.5%, and to 17%, respectively. It is important, however, that the TLM does not add any explanatory power to the model for the long-term returns, and adding the TLM to the other option-based variables even decreases the \bar{R} slightly.

The other option-based variables perform very well in explaining the market returns in a no-crash state. The most robust measure seems to be the implied correlation (IC), which demonstrates a positive and significant coefficient for all three return horizons; moreover, the coefficient scales almost linearly with the holding period. The one-volatility change in IC adds 0.27% per week to the expected short-term market return, and it also adds 3.09% per quarter to the expected quarterly return. Variance risk premium (VRP) does not work well for a short period of one week, though it also demonstrates a positive and significant sign for longer holding periods of one and three months. The significance of VRP increases with the time period, and it is consistent with the observation of Bollerslev, Tauchen, and Zhou (2009, page 1) that “the magnitude of the [VRP] predictability is particularly strong at the intermediate quarterly return horizon.” The marginal effects of the VRP are two to three times weaker than those demonstrated by the IC. The higher (third and fourth) moments of risk-neutral distribution, namely MFIS and MFIK, are not significantly related to future market returns. Implied Volatility Index (VIX), however,

typically named a “fear index” of the market, is negatively related to expected returns in the future. While it is consistent with the leverage effect (e.g., Christie (1982), among others), it might seem to be contradicting the economic theory. Fear should be associated with demanding a higher risk premium for holding the market, and not vice versa.

To study what happens in case of a crash realization, we include in the model the interaction “crash” dummies for the TLM and the low-probability negative market return. The dummy is equal to 1 when the market has a decline in the following holding period larger than one ($D1_{const}$ for constant and $D1_{slope}$ for the slope) or two ($D2_{const}$ and $D2_{slope}$, respectively) sigmas, where sigma is defined as the average value of VIX over the last quarter rescaled to match the duration of the holding period (one week, one month, or three months). When we include either set of “crash” dummies into the regression (3), we get a significantly negative adjustment to the constant, indicating that some of the realized negative returns are left unexplained by our current right-hand-side variables. More importantly, the TLM slope gets adjusted with a highly significant negative coefficient for the regressions with 1-week and 1-month returns (and for $D1_{slope}$ for 3-month returns). The adjustment is several times higher than the original TLM coefficient, and with one-volatility difference in the TLM, keeping all other variables constant, it amounts to -1.35% per week for one-sigma decline, and -1.84% per week for two-sigma decline. For a monthly return regressions the one- and two-sigma slope dummies are equivalent to the decline in future market return by 2.40% and 3.70% per month, respectively, for a one-volatility increase in TLM.

Thus, we fail to reject the composite Hypothesis 1, and indeed a higher TLM indicates a higher required return on the market (especially so for short-term holding periods), when no low-probability crash occurs; however, when the crash occurs, a higher TLM is related to larger losses for the market index. The implied correlation (for all selected holding periods) and variance risk premium (for longer horizons) prove to be very important and significant, both economically and statistically, in explaining future market returns. Still, a part of the extreme negative returns remains unexplained by the selected variables.

4.1.2 Probit Regressions

Our main objective here is to understand, whether it is possible to link the probability of a market crash under the objective (or true) probability measure and the observed tail loss measure, which is computed under the risk-neutral probability, and whether the relationship is positive. To study Hypothesis 2, we use the binary probit model estimated over the whole sample period.¹⁰

The categorical variable y_t on the left-hand side of our model is defined in terms of future market movements beyond a certain negative threshold. We consider a case with two possible values of y_t depending on the expected market volatility $\widehat{VIX}_t = \frac{1}{63} \sum_{i=0}^{62} VIX_{t-i}$ and a realized market return $r_{t,t+\Delta t}$ in the next holding period window. More precisely, we define the categorical variable at time t as follows:

$$y_t = \begin{cases} 0 & \text{if } r_{t,t+\Delta t} \geq -2\sigma_t, \\ 1 & \text{if otherwise,} \end{cases}$$

where $\sigma_t = \widehat{VIX}_t/100 \times \sqrt{\frac{\Delta t}{251}}$ is the expected market volatility rescaled to the holding period.

As explanatory variables we use the same option-based variables as for the time-series regressions in the previous section, with the only exception being the crash dummies. For the TLM we separately use the level and the weekly changes combined with the other variables in levels. In Table 3 Panel A you will find the coefficient estimates, and in Panel B—the marginal effect estimates expressed as a change in the probability of a 2-sigma market crash within the next holding period arising from increasing a given variable by one unconditional volatility and keeping all other variables at their unconditional mean values.

Analyzing the sign and statistical significance of variables in explaining the probability of a market crash, we should confess that it is indeed not easy to predict sudden market declines. The tail loss measure has an expected positive sign, i.e., increasing TLM

¹⁰As in Glick, Moreno, and Spiegel (2011), when studying the link between the probability of a low-probability event and a set of variables, we use the overlapping windows of returns. The rationale for using the overlapping windows is that crashes are rare, and if they do not last long enough, we may even oversee some of them by sampling infrequently.

indicates a higher probability of crash, but it is not quite significant. The weekly change TLM also demonstrates a positive sign, which is significant at the 5% level for 1-week holding period, and at 10% for the 1-month holding period. Being statistically significant, the Δ TLM is not really significant economically: its one-volatility increase boosts the probability of a 2-sigma crash only by less than 0.5% across all holding periods. The other variables are also not extremely helpful in predicting the low-probability crashes. VIX has a positive coefficient for 1-week horizon, and it turns negative and insignificant for longer horizons. Both MFIS and MFIK have negative coefficients, which become significant for longer holding periods. The negative sign is well understood for skewness, because less negative skewness is associated with lower probability of a crash, but that a higher kurtosis is associated with a lower probability of a crash comes as a surprise. The same negative (and surprising) signs have IC and VRP, though for both the effects are not significant. An explanation may be that MFIK, IC, and VRP increase only *after the crash* has occurred as an overreaction of investors.¹¹ By any rate, significant or not, the marginal influence of each variable on the probability of a crash is tiny, and, while we fail to reject Hypothesis 2, we doubt that option-based variables will turn to be very useful in pre-seeing the crashes in portfolio optimization and taking the respective preemptive actions.

4.2 Using TLM in Portfolio Optimization

We should understand that the return (and probability of crash) predictability exercises that we performed in the previous section 4.1 are in-sample ones, and hence the results of the estimation cannot be used by an investor in portfolio optimization or applications alike. To see how much information useful for predicting future market performance is contained in various option-implied variables *ex ante*, we carry out a portfolio optimization exercise.

We start by using two assets in a portfolio: a risk-free one, and the market. We specify the portfolio weight (investment into the market index) as a function of selected variables, which can be interpreted as instruments used to form managed portfolios described in

¹¹This explanation goes along the lines of Bates (1991), who observed that after the 1987 crash the various S&P 500 index options developed their smirk, while it was not pronounced before.

Cochrane (2005). In particular, we use weight w_t of the following form:

$$w_t = \sum_{j=1}^J \theta_{j,t} \times x_{j,t},$$

where $x_{j,t}$ is the value of the instrument j , and $\theta_{j,t}$ is the loading on the respective managed portfolio at time t .¹² The net return of the portfolio in the period from t to $t + \Delta t$ is equal to $r_{t,t+\Delta t}^p = r_{t,t+\Delta t}^f + w_t \times r_{t,t+\Delta t}^{em}$, where $r_{t,t+\Delta t}^f$ is the net riskfree rate, and $r_{t,t+\Delta t}^{em}$ is the market return in excess of the riskfree rate. At each point in time t we optimally select the vector of loadings θ_t on the managed portfolios by maximizing average utility of final wealth from investing a unit of consumption on each day during the estimation window of length T :

$$\theta_t = \arg \max_{\{\theta_t\}} \frac{1}{T - \Delta t} \sum_{\tau=t-T}^{t-\Delta t} u(1 + r_{\tau,\tau+\Delta t}^f + w_\tau \times r_{\tau,\tau+\Delta t}^{em}). \quad (6)$$

For the empirical implementation we use the constant relative risk aversion (CRRA) utility function with the relative risk aversion $\gamma = 5$, and the estimation window $T = 251$ trading days, or one calendar year. Moreover, we restrict the maker weights to be between -0.5 and 1.5 , thus limiting borrowing and shortsales at reasonable levels. After we find the optimal portfolio weight for each day, we calculate the out-of-sample portfolio return and a number of performance metrics, in particular, we compute mean return, volatility, certainty equivalent for a respective utility used for optimization problem (6), and portfolio turnover.

Empirically, we consider several combinations of option-based variables; however, instead of creating a large number of the managed portfolios, we always create only one, either from a single instrument (=single variable) or from the synthetic instrument, which we extract from several variables as their first principal component for each estimation window. The reason for such a procedure is twofold: (i) the first principal component is very informative, and it typically explains more than 85% (and often more) of the total variance for all the combinations of the variables we used, and (ii) having more than one

¹²Another interpretation of our approach is the parametric portfolio policies by Brandt, Santa-Clara, and Valkanov (2009), though instead of reducing the dimensionality of the portfolio optimization with many risky assets, we potentially increase the dimensionality of the problem with one risky asset only.

managed portfolio constructed from our variables considerably decreases the stability of the optimal weights and deteriorates the out-of-sample portfolio performance.

To test whether tail loss measure and other variables have incremental information about the future market distribution beyond the history of the market returns, we specify a “non-informative” instrument (NI) equal to 1 at any point in time, and we compare the out-of-sample performance of each portfolio to that of the NI portfolio. Because we evaluate the portfolio performance for overlapping holding periods, to calculate the p -values we use the bootstrapping methodology described in Efron and Tibshirani (1993), with additional adjustment, as in Politis and Romano (1994), to account for the autocorrelation in overlapping returns. The results of the estimation are provided in Table 4. In Panel A we show the return statistics,—mean, standard deviation (std), and the certainty equivalent return (ce) for the CRRA investor with $\gamma = 5$, and in Panel B we provide the statistics for the portfolio market weights (mean and standard deviation), and loadings on the managed portfolios (mean).

The out-of-sample performance estimates indicate that a non-informative managed portfolio performs relatively well, bringing 9.38% and 9.98% annualized return for monthly and quarterly rebalancing, respectively. Note that the average monthly risk-free rate and market return for the sample period are 3.67% and 5.70% p.a. Moreover, the certainty equivalent return of the market investment with monthly rebalancing turns out negative -2.51% p.a. Three of the five managed portfolios outperform the non-informative optimal portfolio: TLM, IC, and composite one (TLM, VIX, IC, VRP); however, for only two of them the outperformance is statistically significant (IC for quarterly rebalancing, and borderline significant for monthly rebalancing, and composite for quarterly rebalancing). By point estimates, the TLM earns 60 basis points per annum more than the NI, though after taking into account the risks by comparing the certainty equivalent, the difference decreases to 28 b.p. The IC-based managed portfolio earns almost 1% p.a. more than the NI portfolio, with the risk-adjusted return (ce) difference of 79 b.p., for quarterly rebalancing. The composite portfolio performs slightly worse than the IC-based one.

Looking at the characteristics of market weights in Panel B, we observe that on average investors are long the market, with the weights from minimum 0.48 to maximum

0.75 across all portfolios and holding periods. The NI portfolio has the lowest market weight, and VIX portfolio has the highest one. The weights are very volatile during the sample period, and their behavior can be well seen in Figure 1, where we depict the time series of weights for two portfolios (NI and composite ones) with the monthly holding period, and the level of S&P 500 index for the same time interval. The weights for both portfolios mostly stay in the boundary states, and hence neither strategy requires a lot of day-to-day rebalancing. Visually, the NI portfolio weights are slightly more stable, and they diminish faster and by more in the case of a deteriorating market index performance; NI weights also have a tendency to stay low slightly longer than those of the portfolio based on option-implied variables.

By analyzing the sign and significance of the loading on the managed portfolios, we can also draw inferences about the average relation between the option-implied instrumental variables and the return characteristics beneficial to the optimizing investor.¹³ For all of our portfolios the loadings are positive and highly significant, which indicates that investors are on average long the managed portfolios (note that for the non-informative portfolio this average loading is equivalent to the mean market weight), and hence high values of all characteristics are associated with the expected conditions beneficials to the investor. It confirms the positive sign of TLM, IC, and VRP in the market return explanatory regressions from section 4.1; however, the sign for the VIX is the opposite of the one in the regressions. This can happen due to two reasons: (i) the conditioning of information is important in a particular application, and the unconditional explanatory regression (3) is somewhat misspecified; and (ii) the VIX is important for predicting not only the first moment of returns but also other return distribution characteristics, and there is a conditional positive relation between VIX and return characteristics that the investor benefits from (e.g., skewness). The positive sign of the managed portfolio loadings indicates that high values of the implied variables are associated with the high conditional values of the equity premium; in other words, using these variables, we can reap higher premiums by going long the market risk, but we still have no overwhelming

¹³Our loading on the managed portfolio is similar to the loading on the cross-sectional characteristic in the parametric portfolio policies of Brandt, Santa-Clara, and Valkanov (2009) and can be interpreted in a similar fashion.

indication of future crash that would allow us to reduce the exposure *before* the crash occurs.

Thus, the results of the out-of-sample portfolio analysis suggest that by using information option-implied characteristics, we can improve the portfolio performance, and the improvement is significant in economic terms (up to 1% p.a. of extra return). However, it is still not clear how much the implied measures can be useful in identifying the future crashes with high enough precision to be taken into account for investment decisions.

4.3 The Individual TLM and Cross-Section of Returns

To see how the stock-specific tail loss measure is related to the cross-section of stock returns, we perform two standard exercises. First, we compute the performance of the monthly rebalanced long-short portfolios, and, second, we estimate the two-stage Fama and MacBeth (1973) regressions.

At the end of each month, we sort the stocks into terciles, quintiles, or deciles based on individual stock TLMs, and create a long-short portfolio by taking a long (short) position in stocks within the top (bottom) quantile. We hold this long-short portfolio until the end of the subsequent month, and then proceed with next rebalancing. We apply three weighting rules for the long and short portfolios, namely, value-weighting (VW), equal-weighting (EW), and prior gross-return weighting (RW).¹⁴ The return analysis of the long-short portfolio is provided in Table 5. We observe that there is a clear positive relation between the TLM and future stock returns, which is also statistically significant for EW and RW portfolios, and which is stronger for portfolios with more extreme values of TLM.

The VW portfolios demonstrate positive, but insignificant, returns and alphas, which may be due to the fact that some other size-related characteristics affect the returns. To see how the TLM characteristic interacts with the other option-implied and traditional

¹⁴The RW is intended to reduce the possible bias arising due to the errors in recorded returns (e.g., due to bid-ask bounce), and it is suggested by Asparouhova, Bessembinder, and Kalcheva (2010, 2012) as the first best-method for bias correction.

stock characteristics, we refer to the Fama-MacBeth regression results in Table 6.¹⁵ Following our intuition from the time-series regressions, we expect a positive relation between the TLM and expected stock returns, when there is no realization of the stock-specific low-probability crash. When the crash occurs, however, we expect the magnitude of the crash to be positively related to the TLM. To test these relations, we include two sets of “crash” dummies in the first-stage cross-sectional regressions. The dummy is equal to 1 for a given stock when this stock return has a decline in the subsequent month larger than one ($D1_{const}$ for constant and $D1_{slope}$ for the slope) or two ($D2_{const}$ and $D2_{slope}$, respectively) sigmas, where sigma is the expected volatility given by the average value of stock-specific model-free implied volatility over the last quarter. If no stock experiences a crash in the following month, we set the dummy values to missing. We proceed with the second stage as usual, by computing average values for each coefficient and testing their significance.

The results confirm that without additional explanatory variables the TLM is on average positively related to expected returns (with a p-value of 0.06). For the 1-sigma crash, a higher stock-specific TLM is associated with a larger magnitude of a crash, and otherwise the TLM is positively related to future returns (both relations are very significant). As we have seen in the time-series tests, a large part of the negative “crash” return is still unexplained by the selected variables. When we separately treat only the 2-sigma events, we still get a significant positive relation between the TLM and the magnitude of the crash, but the TLM itself is no more (significantly) related to the expected returns. We explain this by observing that average individual volatility is about two times higher than the index volatility, and hence a 1-sigma event for individual stocks is equivalent to quite a large negative return. Not treating this kind of event separately distorts the estimation of the average relation, especially taking into account that negative idiosyncratic jumps in individual stocks’ sample happen more often than in one time-series of the market index.

Thus, we fail to reject composite Hypothesis 3, and the individual stock tail loss measure is positively related to the cross-section of expected stock returns, when there

¹⁵Note that results for EW and RW (hence, bias-corrected) sorted portfolios are almost identical, and hence we do not expect any bias in the Fama-MacBeth procedure, which is implicitly an equal-weighted test, and hence would be biased in the presence of errors in the stock returns.

is no stock-specific low-probability crash realized. In the case that the stock-specific low-probability crash is realized, the stock tail loss measure is positively related to the magnitude of the negative realized stock return.

5 Conclusion

Using results of the Extreme Value Theory we estimate from the observed option prices a forward-looking tail loss measure (TLM). It quantifies the expectations of the investors (adjusted for the investors' tolerance for bearing risk) of the magnitude of low-probability crashes in the near future. In addition to the TLM, we consider the most popular option-implied predictors of the future moments of asset returns—model-free implied variance, skewness, and kurtosis (MFIV, MFIS, and MFIK), implied correlation (IC), and variance risk premium (VRP).

We test whether the TLM and the other option-implied, and hence forward-looking, measures, are able to explain the magnitude and probability of future market and stock-specific crashes, as well as the future asset returns, when no crash occurs. We find that TLM is on average positively related to the future market returns, especially in the short run, moreover, when a crash occurs, the TLM is positively related to the magnitude of the realized negative return. An increasing TLM is associated with increasing probability of a crash in the near future. In effect, high expectations of the tail loss make investors demand a higher premium for taking market risk. The other variables that unambiguously and significantly explain future market returns are the VRP (positive and significant for longer periods), and the IC (positive and significant for all periods).

We also perform an out-of-sample portfolio optimization exercise. We show that using the IC as an instrumental variable to construct a managed portfolio leads to significantly better results than not using any information beyond the past market performance. TLM also provides an investor with incremental useful information, though the performance is not significantly better than in the non-informative optimization. Using all implied variables together works well, though all variables can barely catch up with the IC-based portfolio.

Cross-sectional sorting and Fama-MacBeth procedures show that stock-specific TLM is positively related to the stock returns in the cross-section, and, when a stock-specific crash occurs, the magnitude of the negative stock return is greater, the higher the stock-specific TLM is.

Appendices

A Estimation of the Generalized Pareto Distribution Parameters from Option Prices

Theorem 1 Assuming that the risk-neutral distribution of asset value x belongs to the maximum domain of attraction of an extreme value distribution H_ξ , the tail loss measure $\Omega_h \equiv \frac{E(h-x|h>x)}{x}$, inferred from the prices of traded options on asset x , with h near the (left) endpoint of x , is unique.

Proof. If the market is complete, then there exists only one risk-neutral measure (up to the difference in beliefs, if any). If the market is not complete, there exist multiple risk-neutral measures. If all the existing risk-neutral distributions belong to the MDA of the extreme value distribution H_ξ , then by the Pickands Balkema De Haan theorem, we can find for each distinct risk-neutral distribution a generalized Pareto distribution (GPD) $G_{\xi,\beta}$ describing the excess value distribution, with the same tail shape parameter ξ as of the extreme value distribution H_ξ . The scaling parameter β depends on the threshold only. Hence, for the threshold h , extreme enough to qualify as threshold for all risk-neutral distributions, and, given the same tail shape parameter ξ , there will be a unique GPD for the excess value under all existing risk-neutral distributions. Thus, even though the option prices are discrete, we only need two put option prices for a given maturity, with strikes below the threshold to identify the parameters of the respective unique GPD and then compute the tail loss measure, unique for a given threshold. ■

We estimate the parameters of the generalized Pareto distribution from option prices using the procedure similar to the one described in Hamidieh (2011). We use the following two facts from the Extreme Value Theory (EVT) for the derivations:

1. For most continuous distributions of a random variable x , we can describe the excess value over (or under) a high (low) threshold u conditional on being above (below) the threshold by a generalized Pareto distribution $G_{\beta,\xi}$, and hence the conditional expectation of the excess value can be estimated as $E[x - u|x > u] = \beta/(1 - \xi)$ (or

$E[u - x|x < u] = \beta/(1 - \xi)$, if we are talking about the left tail), where β and ξ are the scale and tail shape parameters of the distribution.¹⁶

2. Assuming that we can describe the conditional excess value of a random variable x over an extreme threshold u by the GPD with parameters β and ξ , one can show that the conditional excess value of x over any more extreme threshold v can also be described by the GPD with the respective scale $\beta + \xi \times |v - u|$ and tail shape ξ parameters.¹⁷

Observe that the price $P_t(K)$ of an out-of-the-money put option with strike K can be computed as the discounted risk-neutral expectation of the future payoff, i.e., $P_t = e^{-r(T-t)} E_t^Q[(K - S_T) \times 1_{K > S_T}]$. Using the rules of conditional probabilities, we can rewrite

$$P_t(K) = e^{-r(T-t)} E_t^Q[(K - S_T)|K > S_T] \times E^Q[1_{K > S_T}]. \quad (\text{A1})$$

Assuming that the risk-neutral distribution Q belongs to the MDA of the GEV, and that K is far enough to qualify as a “extreme enough threshold,” we use fact #1 from above and rewrite equation (A1) in terms of the parameters of the GPD under the probability measure Q :

$$P_t(K) = e^{-r(T-t)} \frac{\beta(K)}{1 - \xi} \times E^Q[1_{K > S_T}]. \quad (\text{A2})$$

Now consider the put option with a strike $K_1 \leq K$, and use fact #2 to infer the distribution and the expectation of the payoff of the option on maturity, given that it expires in-the-money. The distribution of the excess conditional payoff of the option will be described by the GPD with parameters $\beta + \xi \times (K - K_1)$ and ξ , and hence

$$P_t(K_1) = e^{-r(T-t)} \frac{\beta(K) + \xi \times (K - K_1)}{1 - \xi} \times E^Q[1_{K_1 > S_T}]. \quad (\text{A3})$$

¹⁶It follows from the Pickands Balkema De Haan theorem (a.k.a. the second theorem of EVT), which states that the distributions for which normalized maxima converge to a Generalized Extreme Value (GEV) distribution, and which belong to the Maximum Domain of Attraction (MDA) of GEV distribution, constitute a set of distributions for which the excess distribution converges to the GPD as the threshold is raised. It is known (i.e., Resnick (1987)) that a vast majority of known continuous distribution commonly used in the statistics and finance belongs to the Maximum Domain of Attraction of GEV distribution, i.e., the theorem “proves to be a widely applicable result that essentially says that the GPD is *the canonical distribution* for modeling excess losses over high thresholds.” (McNeil, Frey, and Embrechts (2005, p. 278)).

¹⁷The second fact is proven as a Lemma 7.22 on page 279 in McNeil, Frey, and Embrechts (2005).

Note that the expectation of the indicator function $E^Q[1_{K>S_T}]$ in the equations above is just the risk-neutral probability of the event $K > S_T$, i.e., $\mathbf{P}^Q(K > S_T)$. One can show (e.g., McNeil, Frey, and Embrechts (2005, p. 283)) that by making the above assumption about the Q -distribution being in the MDA of the GEV, we can express the probability of a more extreme event $K_1 > S_T$ as a probability of a less extreme event $K > S_T$ and the parameters of the GPD as follows:

$$\mathbf{P}^Q(K_1 > S_T) = \mathbf{P}^Q(K > S_T) \times \left(1 + \xi \times \frac{K - K_1}{\beta(K)}\right)^{-1/\xi}. \quad (\text{A4})$$

Putting equations (A2), (A3), and (A4) together, we can rewrite the price of a farther out-of-the money (OTM) put as a function of a less OTM one:

$$P_t(K_1) = P_t(K) \left(1 + \xi \times \frac{K - K_1}{\beta(K)}\right)^{1-1/\xi}. \quad (\text{A5})$$

We can use equation (A5) to derive the theoretical prices of far out-of-the-money put options using the price of a put option with a “boundary” strike price K , which is low enough to qualify as suitable threshold.¹⁸

The procedure for estimating the tail loss measure has three steps. First, we use the pricing equation (A5) to price far out-of-the-money puts relative to a “boundary” put with the strike K assumed to be a suitable threshold for EVT to work. Second, we compare the theoretical prices to the observed market prices of the respective put options, and infer the parameters $\beta(K)$ and ξ of the generalized Pareto distribution by minimizing the pricing errors. Third, using the estimated parameters of GPD, we compute the tail loss measure (TLM) conditional on a given threshold. It is given by a risk-neutral expectation of the excess value of the random variable, which is the market loss beyond the given threshold K conditional on market falling lower than level K on the maturity date of the options:

$$\text{TLM} = \frac{\beta(K)}{1 - \xi}.$$

¹⁸Refer to Hamidieh (2011) for the proofs that the pricing formula in (A5) satisfies required bounds and possesses necessary properties.

Table 1: Market Return Predictability

In this table, we report the results of time-series explanatory regression of the market return on a number of option-based predictors: tail-loss measure (TLM), traded Implied Volatility Index (VIX), model-free implied skewness (MFIS), model-free implied kurtosis (MFIK), implied correlation (IC), and variance risk premium (VRP). We also include the interaction dummies for the TLM and the low-probability crash. The dummy is equal to 1 when the market has a decline in the following holding period larger than one ($D1_{const}$ for constant and $D1_{slope}$ for the slope) or two ($D2_{const}$ and $D2_{slope}$, respectively) sigmas, where sigma is defined as the average value of VIX over the last quarter rescaled to match the duration of the holding period (one week, one month, or three months). We estimate the model with overlapping observations using robust bisquare data transformation with the default tuning constant of 4.685, which makes our estimates robust to outliers (see Mosteller and Tukey (1977) for details) and provides 95% efficiency. The two-sided p-values in the parentheses are based on the Newey and West (1987) autocorrelation-adjusted standard errors with the lag equal to the number of overlapping periods.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
	1-week return			1-month return			3-month return								
Const	-0.0037 (0.01)	-0.0027 (0.84)	0.0012 (0.93)	0.0043 (0.68)	0.0023 (0.87)	-0.0067 (0.22)	0.0449 (0.45)	0.0565 (0.36)	0.0610 (0.14)	0.0443 (0.44)	-0.0044 (0.79)	0.0754 (0.63)	0.0961 (0.55)	0.0680 (0.58)	0.0969 (0.54)
$D1_{const}$				-0.0073 (0.01)					-0.0280 (0.03)					-0.0839 (0.00)	
$D2_{const}$					-0.0182 (0.04)					-0.0502 (0.01)					-0.3798 (0.00)
$TLM \times D1_{slope}$				-0.7174 (0.00)					-1.2792 (0.00)					-1.8455 (0.00)	
$TLM \times D2_{slope}$					-0.9768 (0.00)					-1.9669 (0.00)					1.5405 (0.09)
TLM	0.1124 (0.00)		0.1537 (0.04)	0.2039 (0.00)	0.1648 (0.02)	0.2947 (0.01)		0.2623 (0.28)	0.1708 (0.42)	0.2210 (0.35)	0.4826 (0.17)		0.4740 (0.32)	0.5296 (0.22)	0.4931 (0.29)
VIX		-0.0113 (0.25)	-0.0414 (0.03)	-0.0435 (0.01)	-0.0436 (0.02)		-0.0670 (0.08)	-0.1196 (0.08)	-0.0725 (0.18)	-0.1072 (0.10)		-0.2617 (0.01)	-0.3471 (0.01)	-0.3040 (0.01)	-0.3465 (0.01)
MFIS		0.0010 (0.81)	0.0022 (0.60)	-0.0017 (0.60)	0.0015 (0.72)		-0.0037 (0.82)	-0.0035 (0.82)	-0.0144 (0.20)	-0.0025 (0.87)		-0.0237 (0.60)	-0.0245 (0.58)	-0.0091 (0.79)	-0.0252 (0.56)
MFIK		-0.0004 (0.92)	-0.0015 (0.74)	-0.0028 (0.41)	-0.0019 (0.67)		-0.0178 (0.35)	-0.0214 (0.28)	-0.0225 (0.09)	-0.0175 (0.34)		-0.0331 (0.51)	-0.0399 (0.44)	-0.0259 (0.51)	-0.0401 (0.43)
IC		0.0220 (0.00)	0.0205 (0.00)	0.0197 (0.00)	0.0212 (0.00)		0.0877 (0.00)	0.0848 (0.00)	0.0716 (0.00)	0.0825 (0.00)		0.2409 (0.00)	0.2349 (0.00)	0.1928 (0.00)	0.2314 (0.00)
VRP		0.0489 (0.09)	0.0466 (0.12)	-0.0019 (0.92)	0.0285 (0.31)		0.2091 (0.04)	0.2200 (0.03)	0.0465 (0.48)	0.2240 (0.02)		0.4966 (0.00)	0.4719 (0.00)	0.4927 (0.00)	0.4711 (0.00)
\bar{R}_t , %	1.34	2.28	2.49	40.78	18.59	2.52	7.67	7.84	46.51	26.36	3.01	17.02	16.92	48.76	37.87

Table 2: Market Return Predictability: Marginal Effects from One-Standard Deviation Change in Variable

In this table, we report the marginal effects of a one-standard deviation increase in an explanatory variable (keeping all other variables constant) on the expected market return, for three holding periods (a week, a month, or a quarter). The results are based on the unconditional volatility of the explanatory variables and the coefficients from a fitted time-series explanatory regression of the market return on a number of option-based predictors: tail-loss measure (TLM), traded Implied Volatility Index (VIX), model-free implied skewness (MFIS), model-free implied kurtosis (MFIK), implied correlation (IC), and variance risk premium (VRP). We also include the interaction dummies for the TLM and the low-probability crash. The dummy is equal to 1 when the market has a decline in the following holding period larger than one ($D1_{const}$ for constant and $D1_{slope}$ for the slope) or two ($D2_{const}$ and $D2_{slope}$, respectively) sigmas, where sigma is defined as the average value of VIX over the last quarter rescaled to match the duration of the holding period (one week, one month, or three months). We estimate the model with overlapping observations using robust bisquare data transformation with the default tuning constant of 4.685, which makes our estimates robust to outliers (see Mosteller and Tukey (1977) for details) and provides 95% efficiency.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
	1-week return			1-month return			3-month return								
TLM \times $D1_{slope}$				-0.0135					-0.0240						-0.0347
TLM \times $D2_{slope}$					-0.0184					-0.0370					0.0289
TLM	0.0021		0.0029	0.0038	0.0031	0.0055	0.0049	0.0032	0.0042	0.0042	0.0091	0.0089	0.0100	0.0093	0.0093
VIX		-0.0010	-0.0036	-0.0038	-0.0038		-0.0058	-0.0104	-0.0063	-0.0094		-0.0228	-0.0303	-0.0265	-0.0302
MFIS		0.0002	0.0004	-0.0003	0.0003		-0.0007	-0.0006	-0.0026	-0.0004		-0.0043	-0.0044	-0.0016	-0.0045
MFIK		-0.0001	-0.0003	-0.0006	-0.0004		-0.0035	-0.0042	-0.0045	-0.0035		-0.0066	-0.0079	-0.0051	-0.0080
IC		0.0029	0.0027	0.0026	0.0028		0.0116	0.0112	0.0094	0.0109		0.0317	0.0309	0.0254	0.0305
VRP		0.0012	0.0012	-0.0000	0.0007		0.0052	0.0055	0.0012	0.0056		0.0123	0.0117	0.0122	0.0117

Table 3: Explaining the Probability of a Market Crash: Probit Regression

In this table, we report the estimation of the binary probit regression, where the left-hand-side variable takes the value of 1 in the case of a low-probability market crash during the holding period (which is a week, or a month, or a quarter). Low-probability market crash is defined as a index decline in the following holding period larger than two sigmas, where sigma is defined as the average value of VIX over the last quarter rescaled to match the duration of the holding period. The right-hand side is represented by a number of option-based predictors: tail-loss measure (TLM), weekly change in the tail-loss measure (Δ TLM), traded Implied Volatility Index (VIX), model-free implied skewness (MFIS), model-free implied kurtosis (MFIK), implied correlation (IC), and variance risk premium (VRP). Panel A provides the estimates of the coefficients, with t-statistics presented in the parenthesis, and McFadden R^2 in the last row. Panel B presents the marginal changes in probability (in %) of the 2-sigma market crash for a one-volatility increase in a given variable, keeping all other variables constant at their unconditional mean values.

	1-week		1-month		3-month	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Probit estimation</i>						
Const	-1.0841 (0.61)	-0.6086 (0.78)	9.3863 (0.09)	11.3340 (0.06)	16.3973 (0.00)	15.3110 (0.00)
TLM	7.9447 (0.29)		17.0653 (0.46)		48.1962 (0.10)	
Δ TLM		17.7884 (0.01)		35.0813 (0.07)		16.2824 (0.37)
VIX	1.9709 (0.22)	3.0162 (0.00)	-2.6910 (0.57)	-0.8325 (0.73)	-11.3750 (0.11)	-1.8251 (0.44)
MFIS	-0.4926 (0.52)	-0.5806 (0.45)	-1.7272 (0.32)	-1.9119 (0.27)	-2.2067 (0.14)	-2.6096 (0.07)
MFIK	-0.6088 (0.40)	-0.7249 (0.32)	-3.8735 (0.04)	-4.4006 (0.03)	-5.8693 (0.00)	-5.5500 (0.00)
IC	-1.1817 (0.12)	-1.1102 (0.14)	0.1358 (0.93)	0.3942 (0.81)	-0.1236 (0.93)	0.0910 (0.95)
VRP	-0.9988 (0.52)	-1.3029 (0.38)	-3.1493 (0.49)	-2.9797 (0.47)	-6.3335 (0.30)	-5.4262 (0.26)
McFadden R^2 , %	7.95	9.81	11.85	15.96	19.29	17.49
<i>Panel B: Marginal effects</i>						
TLM	0.3320		0.6257		3.1273	
Δ TLM		0.3574		0.4751		0.2129
VIX	0.3936	0.6453	-0.2209	-0.0607	-0.3123	-0.1485
MFIS	-0.1467	-0.1613	-0.2664	-0.2022	-0.2324	-0.2999
MFIK	-0.1923	-0.2116	-0.3979	-0.2877	-0.3174	-0.3779
IC	-0.2373	-0.2144	0.0232	0.0514	-0.0155	0.0140
VRP	-0.0444	-0.0546	-0.0897	-0.0617	-0.1245	-0.1297

Table 4: Optimal Managed Portfolios with Option-Implied Variables

In this table, we report the statistics for the portfolios created by rolling-window optimization of the past average utility of a CRRA investor with relative risk aversion $\gamma = 5$. The portfolio consists of a risk-free asset and a managed portfolio, which is equal to the product of the instrumental variable and the market return. The investor chooses the optimal loading on the managed portfolio each day and holds this portfolio for a specified holding period (a week, or a month, or a quarter). As instrumental variables we use a “non-informative” variable (NI) equal to a vector of constants, and the variables that significantly predict future market returns in-sample: tail-loss measure (TLM), Implied Volatility Index (VIX), implied correlation (IC), and variance risk premium (VRP). The portfolio performance in *Panel A* is calculated from the overlapping holding period returns and annualized accordingly. The p-values are based on the empirical bootstrap distribution of a given performance metric. We resample the portfolio returns blockwise to preserve the autorrelation in the overlapping returns. The market weight (w) and managed portfolio loading (θ) statistics in *Panel B* are calculated from the sample weights and loadings, respectively, and the p-values are based on the Newey and West (1987) autocorrelation-adjusted standard errors with the lag equal to the number periods (250) in estimation.

Panel A: Portfolio Return Metrics

Instruments	1-week			1-month			3-month		
	mean	std	ce	mean	std	ce	mean	std	ce
NI	6.32 (1.00)	12.88 (1.00)	2.04 (1.00)	9.38 (1.00)	14.47 (1.00)	4.01 (1.00)	9.98 (1.00)	15.19 (1.00)	4.30 (1.00)
TLM	6.48 (0.81)	13.58 (0.00)	1.68 (0.67)	9.67 (0.66)	14.81 (0.15)	4.10 (0.87)	10.58 (0.40)	15.80 (0.25)	4.58 (0.64)
VIX	5.83 (0.72)	13.16 (0.40)	1.31 (0.60)	9.04 (0.80)	14.61 (0.75)	3.69 (0.79)	10.19 (0.86)	16.01 (0.26)	4.04 (0.83)
IC	7.24 (0.16)	13.48 (0.00)	2.50 (0.54)	10.24 (0.12)	14.78 (0.13)	4.69 (0.24)	10.95 (0.00)	15.64 (0.34)	5.09 (0.01)
VRP	5.99 (0.90)	13.89 (0.09)	0.98 (0.66)	5.32 (0.05)	13.37 (0.08)	0.36 (0.11)	8.61 (0.38)	15.53 (0.69)	2.08 (0.32)
TLM, VIX, IC, VRP	6.84 (0.49)	13.21 (0.19)	2.29 (0.75)	9.85 (0.48)	14.64 (0.45)	4.40 (0.55)	10.93 (0.01)	15.64 (0.41)	5.08 (0.02)

Panel B: Market Weights and Managed Portfolio Loadings

Instruments	1-week			1-month			3-month		
	w_{mean}	w_{std}	θ_{mean}	w_{mean}	w_{std}	θ_{mean}	w_{mean}	w_{std}	θ_{mean}
NI	0.48 (0.01)	0.78	0.48 (0.01)	0.58 (0.00)	0.84	0.58 (0.00)	0.67 (0.00)	0.86	0.67 (0.00)
TLM	0.56 (0.00)	0.76	61.89 (0.00)	0.64 (0.00)	0.82	16.60 (0.00)	0.71 (0.00)	0.86	6.08 (0.00)
VIX	0.64 (0.00)	0.66	26.67 (0.00)	0.70 (0.00)	0.76	29.95 (0.00)	0.75 (0.00)	0.85	31.98 (0.00)
IC	0.55 (0.00)	0.77	1.66 (0.00)	0.62 (0.00)	0.83	1.87 (0.00)	0.69 (0.00)	0.86	2.13 (0.00)
VRP	0.58 (0.00)	0.68	47.33 (0.00)	0.60 (0.00)	0.71	48.03 (0.00)	0.73 (0.00)	0.79	59.32 (0.00)
TLM, VIX, IC, VRP	0.55 (0.00)	0.75	1.66 (0.00)	0.61 (0.00)	0.82	1.85 (0.00)	0.68 (0.00)	0.85	2.01 (0.00)

Table 5: Cross-Section of Individual Stocks and TLM: Portfolio Sorting

In this table, we report the average monthly returns on portfolio sorted on stock-specific tail loss measure (TLM) over our sample period from January 1996 to October 2010. At the end of each month, we sort the stocks into terciles, quintiles, or deciles based on individual stock TLM, and create a long-short portfolio by taking a long (short) position in stocks within the top (bottom) quantile. This long-short portfolio is held until the end of the subsequent month and value-weighted (VW), prior gross-return weighted (RW), and equal-weighted (EW) returns of this portfolio are calculated. The table presents returns (Return) and alphas from the market model (CAPM alpha), from the Fama-French three-factor model (3F alpha), and from the Fama-French-Carhart four-factor model (4F alpha). The p-values for the significance are provided in parentheses.

	Terciles			Quintiles			Deciles		
	VW	RW	EW	VW	RW	EW	VW	RW	EW
Return	5.36 (0.42)	14.69 (0.01)	14.74 (0.02)	8.67 (0.29)	16.90 (0.03)	16.75 (0.04)	10.08 (0.29)	21.70 (0.02)	21.47 (0.02)
CAPM alpha	0.86 (0.86)	10.37 (0.01)	10.29 (0.02)	2.92 (0.62)	11.22 (0.04)	10.93 (0.05)	3.47 (0.61)	14.93 (0.02)	14.54 (0.02)
3F alpha	2.44 (0.52)	9.01 (0.01)	8.93 (0.01)	4.90 (0.29)	10.06 (0.02)	9.74 (0.03)	5.34 (0.34)	14.27 (0.01)	13.78 (0.01)
4F alpha	4.29 (0.25)	12.05 (0.00)	12.42 (0.00)	6.94 (0.13)	14.07 (0.00)	14.21 (0.00)	7.16 (0.20)	18.61 (0.00)	18.64 (0.00)


Table 6: Cross-Section of Individual Stocks and TLM: Fama-MacBeth Approach

In this table, we report the estimates from the two-stage Fama and MacBeth (1973) regressions of the next month stock returns on stock-specific tail-loss measure (TLM) and other stock characteristics: book-to-market ratio (BTM), size (SIZE), momentum (MOM), short-term reversal (REV), model-free implied volatility (MFIVol), model-free implied skewness (MFIS), and model-free implied kurtosis (MFIK). To reduce estimation errors, we use the average of the last five observations in each month as the monthly measures. To treat the returns of the stocks experiencing a low-probability crash distinctly, we include crash dummies (slope and interaction) in the regressions. The dummy is equal to 1 for a given stock when this stock return has a decline in the subsequent month larger than one ($D1_{const}$ for constant and $D1_{slope}$ for the slope) or two ($D2_{const}$ and $D2_{slope}$, respectively) sigmas, where sigma is defined as the average value of stock-specific MFIVol over the last quarter. We report the p-values for the significance in parentheses.

	(1)	(2)	(3)	(4)	(5)
$D1_{const}$		-0.0360 (0.00)	-0.0342 (0.00)		
$D2_{const}$				-0.0522 (0.42)	-0.0504 (0.44)
$TLM \times D1_{slope}$		-1.7057 (0.00)	-1.7238 (0.00)		
$TLM \times D2_{slope}$				-2.7814 (0.00)	-2.7925 (0.00)
TLM	0.0943 (0.06)	0.1471 (0.00)	0.0757 (0.01)	0.0384 (0.29)	-0.0112 (0.65)
BTM		-0.0031 (0.05)	-0.0026 (0.07)	-0.0028 (0.14)	-0.0026 (0.12)
SIZE		-0.0038 (0.00)	-0.0027 (0.00)	-0.0040 (0.00)	-0.0031 (0.00)
MOM		0.0000 (0.99)	0.0015 (0.63)	0.0007 (0.87)	0.0022 (0.51)
REV		-0.0082 (0.36)	-0.0065 (0.41)	-0.0035 (0.74)	-0.0028 (0.76)
MFIVol			0.1078 (0.01)		0.0719 (0.13)
MFIS			0.0107 (0.01)		0.0106 (0.01)
MFIK			0.0027 (0.09)		0.0037 (0.04)

Figure 1: Portfolio Weights vs. Market Prices

The figure shows the dynamics of the market prices (S&P 500) and of the portfolio market index weights for two particular instrumental variables—a non-informative constant variable, and a combination of four option-implied variables (TLM, VIX, IC, and VRP) in a form of their first principal component. In portfolio optimization we restrict the market weight to be in the range from -0.5 to 1.5 , so that short selling and borrowing are restricted.



References

- Asparouhova, E., H. Bessembinder, and I. Kalcheva, 2010, "Liquidity biases in asset pricing tests," *Journal of Financial Economics*, 96, 215 – 237.
- Asparouhova, E., H. Bessembinder, and I. Kalcheva, 2012, "Noisy Prices and Inference Regarding Returns," *Journal of Finance*.
- Bakshi, G. S., N. Kapadia, and D. B. Madan, 2003, "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options," *The Review of Financial Studies*, 16, 101–143.
- Balkema, A. A., and L. d. Haan, 1974, "Residual Life Time at Great Age," *The Annals of Probability*, 2, pp. 792–804.
- Barberis, N., and M. Huang, 2008, "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices," *American Economic Review*, 98, 2066–2100.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard, 2009, "Realised Kernels in Practice: Trades and Quotes," *Econometrics Journal*.
- Bates, D., 1991, "The Crash of '87: Was It Expected? The Evidence from Options Markets," *The Journal of Finance*, 46, 1009–1044.
- Bates, D., 2008, "The Market for Crash Risk," *Journal of Economic Dynamics and Control*, 32, 2291–2321.
- Black, F., and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637–654.
- Bollerslev, T., G. Tauchen, and H. Zhou, 2009, "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 22, 4463–4492.
- Bollerslev, T., and V. Todorov, 2011, "Tails, Fears, and Risk Premia," *The Journal of Finance*, 66, 2165–2211.
- Brandt, M. W., P. Santa-Clara, and R. Valkanov, 2009, "Parametric Portfolio Policies: Exploiting Characteristics in the Cross-Section of Equity Returns," *Review of Financial Studies*, 22, 3411–3447.
- Brunnermeier, M. K., C. Gollier, and J. A. Parker, 2007, "Optimal Beliefs, Asset Prices, and the Preference for Skewed Returns," *American Economic Review Papers and Proceedings*, 97, 159–165.
- Brunnermeier, M. K., and J. A. Parker, 2005, "Optimal expectations," *American Economic Review*, 95, 1092–1118.
- Carson, J. M., J. S. Doran, and D. R. Peterson, 2006, "Market Crash Risk and Implied Volatility Skewness: Evidence and Implications for Insurer Investments," Working Paper Florida State University.
- Chang, B. Y., P. Christoffersen, and K. Jacobs, 2012, "Market skewness risk and the cross section of stock returns," *Journal of Financial Economics*.
- Christie, A. A., 1982, "The stochastic behavior of common stock variances: Value, leverage and interest rate effects," *Journal of Financial Economics*, 10, 407 – 432.

- Christoffersen, P., K. Jacobs, and B. Y. Chang, 2011, “Forecasting with Option Implied Information,” CREATES Research Papers 2011-46, School of Economics and Management, University of Aarhus.
- Cochrane, J. H., 2005, *Asset Pricing: (Revised)*, Princeton University Press.
- Cochrane, J. H., 2008, “The Dog That Did Not Bark: A Defense of Return Predictability,” *Review of Financial Studies*, 21, 1533–1575.
- Cremers, K. J. M., M. Halling, and D. Weinbaum, 2012, “Aggregate Jump and Volatility Risk in the Cross-Section of Stock Returns,” Working Paper University of Notre Dame.
- Doran, J. S., D. R. Peterson, and B. C. Tarrant, 2007, “Is there information in the volatility skew?,” *Journal of Futures Markets*, 27, 921–959.
- Driessen, J., P. Maenhout, and G. Vilkov, 2012, “Option-Implied Correlations and the Price of Correlation Risk,” Working paper, INSEAD.
- Driessen, J., P. J. Maenhout, and G. Vilkov, 2009, “The Price of Correlation Risk: Evidence from Equity Options,” *The Journal of Finance*, 64, pp. 1377–1406.
- Efron, B., and R. Tibshirani, 1993, *An Introduction to the Bootstrap*, Chapman and Hall, Boca Raton, FL.
- Fama, E. F., and J. D. MacBeth, 1973, “Risk Return and Equilibrium: Empirical Tests,” *Journal of Financial Political Economy*, 71, 607–636.
- Ferreira, M. A., and P. Santa-Clara, 2011, “Forecasting stock market returns: The sum of the parts is more than the whole,” *Journal of Financial Economics*, 100, 514 – 537.
- Gabaix, X., 2012, “Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance,” *The Quarterly Journal of Economics*, 127, 645–700.
- Glick, R., R. Moreno, and M. Spiegel, 2011, *Financial Crises in Emerging Markets*, Cambridge University Press.
- Hamidieh, K., 2011, “Recovering the Tail Shape Parameter of the Risk Neutral Density from Option Prices,” Working Paper Rice University.
- Han, B., 2008, “Investor Sentiment and Option Prices,” *The Review of Financial Studies*, Vol. 21, No. 1, pp. 387-414.
- Huber, P. J., 2005, *Robust Statistics* . No. 579 in Wiley Series in Probability and Statistics, Wiley, New York.
- Kelly, B. T., 2011, “Tail Risk and Asset Prices,” Fama-Miller Working Paper; Chicago Booth Research Paper No. 11-17.
- McNeil, A. J., R. Frey, and P. Embrechts, 2005, *Quantitative Risk Management: Concepts, Techniques and Tools*, Princeton University Press.
- Merton, R. C., 1973, “An Intertemporal Asset Pricing Model,” *Econometrica*, 41, 867–888.
- Mosteller, F., and J. Tukey, 1977, *Data analysis and regression: a second course in statistics* . Addison-Wesley series in behavioral science, Addison-Wesley Pub. Co.

- Newey, W. K., and K. D. West, 1987, “A simple, positive-semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix,” *Econometrica*, 55, 703–708.
- Pickand, s. J. I., 1975, “Statistical Inference Using Extreme Order Statistics,” *The Annals of Statistics*, 3, pp. 119–131.
- Politis, D., and J. Romano, 1994, “The Stationary Bootstrap,” *Journal of the American Statistical Association*, 89, 1303–1313.
- Poon, S.-H., and C. W. Granger, 2003, “Forecasting Volatility in Financial Markets: A Review,” *Journal of Economic Literature*, 41, 478–539.
- Resnick, S. I., 1987, *Extreme values, regular variation, and point processes* . No. 4 in Applied Probability. A Series of the Applied Probability Trust, Springer-Verlag, New York.
- Zhang, L., P. A. Mykland, and Y. Ait-Sahalia, 2005, “A Tale of Two Time Scales: Determining Integrated Volatility With Noisy High-Frequency Data,” *Journal of the American Statistical Association*, 100, 1394–1411.