Internet Appendix

A. Volatility Smiles and Tail Risk

The Black-Scholes-Merton (BSM) model for valuing options has a crucial free parameter, the future return volatility of the underlying asset. One cannot observe future return volatilities, but for any given option, one can use the BSM model to estimate the return volatility that yields the observed option price. This is referred to as the option's *implied volatility* and can be interpreted as the market's expectation on the future return volatility of the underlying asset. If the BSM model described option prices accurately, the implied volatilities of all options written on a particular stock–and of equal time to expiration–should be the same, irrespective of their strike prices. Hence, plotting the implied volatility of different options against their corresponding strike price should produce a flat line. In reality, implied volatilities vary with strike prices, a phenomenon known as the *volatility smile*.[34]

This skewed shape has been partly attributed to empirical violations of the lognormal assumption for the distribution of stock prices embedded in the BSM model (see Derman and Miller, 2016). In practice, this assumption understates the actual probability of extreme downward moves.[35] In this regard, the risk-neutral density (RND) of stock prices has been shown to be more negatively skewed than the lognormal

[34]Other common names for this phenomenon include volatility smirk and volatility skew.

[35]Specifically, the BSM model assumes that stock log prices follow a constant volatility diffusion process where, over any finite time interval, log prices are normally distributed. In reality, stock return volatility is stochastic and correlated with price. This produces asymmetric and fat-tailed stock return distributions relative to a normal distribution (Corrado and Su, 1996).

density assumed in the BSM model (see Birru and Figlewski, 2012; Dennis and Mayhew, 2002).[36] As an example, Figure A1 presents the risk-neutral density–extracted from option prices–for Sterling Bancorp in December 2014, along with a lognormal density with the same mean and variance.[37] The visible left-skewness of this risk-neutral density makes the probability of a two standard deviations price drop almost three times what a lognormal density implies. A left-skewed RND suggests that investors perceive significant price drops as more likely compared to a lognormal distribution. Because of this, they are willing to pay higher prices for deep OTM put options, which in turn results in a downward sloping volatility smile.

[36] The risk-neutral density contains investors' beliefs about the true distribution of stock returns coupled with their own risk preferences (Figlewski, 2018).

[37] See Birru and Figlewski (2012) for a detailed procedure for constructing risk-neutral densities from option prices.

FIGURE A1

Risk-neutral density

This figure shows the Risk-Neutral Density (RND) for Sterling Bancorp in December 2014 (in blue), along with a lognormal density with the same mean and variance (in red). This RND is constructed using the procedure proposed by Birru and Figlewski (2012).

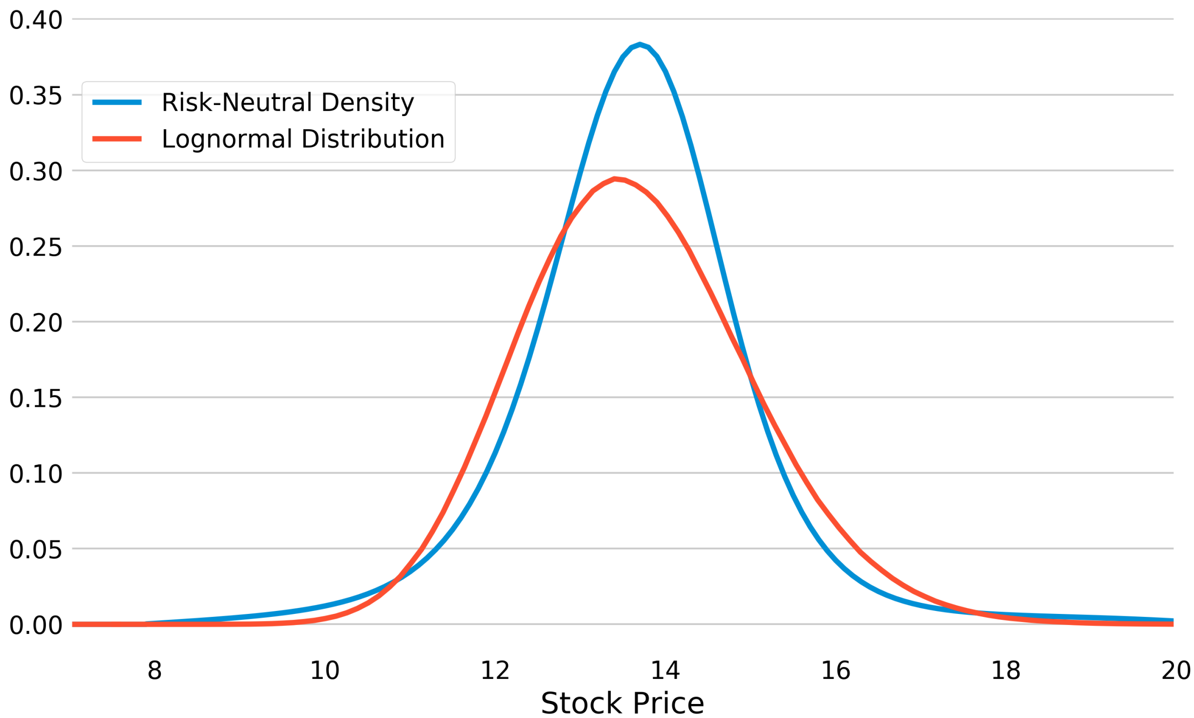

B. Tail Risk Around Crises

The Dot-Com Crash

After a long speculative period known as the dot-com bubble, the market for technology firms crashed in March 2000 and did not recover until late 2002.[38] Given its economic significance, we explore how this crash affected investor's perception of the technology industry's future downside risk (i.e. tail risk).

We use a sample of 165 technology firms listed on NASDAQ and with an active options market between 1996 and 2005. We define the pre-crash, crash, and post-crash periods as 1996-1999, 2000-2002, and 2003-2005, respectively. We show tail risk for technology firms spiked during the dot-com bubble and remained at higher levels compared to the pre-crash period. Specifically, Panel A in Table B1 shows technology firms experienced a 101.8% increase in average tail risk between the pre and post-crash periods. This substantial tail-risk surge represents a *structural* change in the shape of the implied volatility curve for these firms.

Next, we examine the tail-risk behaviour for large and small technology firms around the dot-com crash. We define firms with total assets in the top quartile, as of 2000Q1, as large, and all other firms as small. Panel A of Table B1 presents changes in tail risk for large and small technology firms. Both size groups show a substantial increase in tail risk in the post-crisis period of 132.6% and 45.5% for small and large firms, respectively.

[38] By October 2002, the NASDAQ Composite Index had fallen by 78% from its peak in March 2000.

The Global Financial Crisis

The more recent crisis in 2008-2009 presents another opportunity to study the dynamics of tail risk around crises. Using data for 619 non-financial firms with active options markets between 2001 and 2017, we define the pre-crisis, crisis and post-crisis periods as 2001-2007, 2008-2009, and 2010-2017, respectively.

As before, we calculate the average tail risk for for each period and present the results in Table B1. We find that tail risk significantly increases during the crisis for non-financial firms and remains elevated (12.6% higher) post-crisis compared to the pre-crisis period (see Panel B of Table B1).

We again investigate differences according to firm size. Non-financials are classified into two groups, small and large, based on their total assets as of 2009Q3.[39] The large group corresponds to firms in the top size quartile and the small group consists of all other non-financials.[40] Firm size (i.e. total assets) is obtained from Compustat.

Panel B of Table B1 presents average tail-risk changes for the pre and post-crisis period for small and large firms separately. The post-crisis tail risk increases significantly for both small and large firms by 13.6% and 6.6%, respectively. Thus, investors update their expectations of future crash-like events for both large and small non-financial companies in the wake of the GFC.

Next, using data for 85 U.S. bank holding companies with active options markets

[39]For comparability with the sample of banks, the non-financial firms sample includes non-financials with assets between $2 and $2,252 billion as of 2009Q3. This is the same size range observed for the sample of banks.

[40]This is consistent with the size distribution observed for banks where the above 50B group corresponds roughly to the top size quartile.

between 2001 and 2017 we document a similar but much more pronounced effect is observed for the U.S. banking industry as a whole: tail risk surges during the crisis and remains 69.9% higher post-crisis compared to the pre-crisis period (see Panel C of Table B1).

However, this rise is driven entirely by changes in below 50B banks' tail risk, which surges by 64.4% post-crisis. For above 50B banks, designated as systemically important, tail risk peaks during the crisis and falls back to pre-crisis levels in the post-crisis period.

TABLE B1

Tail Risk Around Crises

This table shows estimates of average quarterly tail risk for technology firms (Panel A), non-financials (Panel B), and banks (Panel C). For banks and non-financials, the sample consists of 85 and 619 firms, respectively, for which active options markets exist between the period 2001-2017. Pre-Crisis refers to the period 2001-2007, Crisis to the period 2008-2009, and Post-Crisis to the period 2010-2017. For technology firms, the sample consists of 165 companies listed on NASDAQ and with active option markets in the period 1996-2005. For these firms Pre-Crisis, Crisis, and Post-Crisis represent the time periods 1996-1999, 2000-2002, and 2003-2005, respectively. For banks, BELOW_50B corresponds to firms with assets lower than $50 billion as of 2009Q3, and ABOVE_50B is the group of firms with assets equal or greater than $50 billion. Non-financials with total assets in the top quartile, as of 2009Q3, are classified as Large and all others as Small. Similarly, technology firms are classified as Large (top quartile) and Small based on their total assets as of 2000Q1.

(A) Technology Firms

	Pre-Crisis	Crisis	Post-Crisis	Post-Pre	% Change
ALL_TECH_FIRMS	0.072	0.142	0.145	0.073***	101.8
SMALL	0.066	0.133	0.152	0.087***	132.6
LARGE	0.085	0.166	0.124	0.039***	45.5

(B) Non-Financials

	Pre-Crisis	Crisis	Post-Crisis	Post-Pre	% Change
ALL_NON_FINANCIALS	0.138	0.177	0.155	0.017***	12.6
SMALL	0.145	0.181	0.164	0.020***	13.6
LARGE	0.121	0.166	0.129	0.008***	6.6

(C) Banks

	Pre-Crisis	Crisis	Post-Crisis	Post-Pre	% Change
ALL_BANKS	0.165	0.288	0.281	0.116***	69.9
BELOW_50B	0.203	0.255	0.333	0.131***	64.4
ABOVE_50B	0.134	0.368	0.131	-0.003	-2.3

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

C. T-statistic for Event Study

The test statistic proposed by Kolari and Pynnönen (2010) has the following form:

$$t_{AR_g} = \frac{\overline{SAR}_g \sqrt{N_g}}{SD_g \sqrt{1 + (N_g - 1)\bar{\rho}_g}} \tag{5}$$

$\overline{SAR}_g$ is the average scaled abnormal return (SAR) for banks in group g on the event day. For each bank, scaled abnormal returns are calculated as $SAR_{i,t} = \frac{AR_{i,t}}{SD_i}$ where SD_i is bank's i sample standard deviation of abnormal returns over the estimation window. N_g corresponds to the number of banks in group g, and $\bar{\rho}_g$ is the average of the sample cross-correlations of scaled abnormal returns for banks in group g over the estimation window. That is:[41]

$$\bar{\rho}_g = \frac{1}{N_g(N_g - 1)/2} \sum_{i=2}^{N} \sum_{j=1}^{i-1} \mathbb{1}_{\{i:i \in g\}} \mathbb{1}_{\{j:j \in g\}} \frac{1}{T_1 - T_0} \sum_{t \in [T_0, T_1]} SAR_{i,t} SAR_{j,t} \tag{6}$$

Finally, SD_g corresponds to the adjusted cross-sectional sample standard deviation of scaled abnormal returns for banks in group g:

$$SD_g^2 = \frac{\frac{1}{N_g - 1} \sum_{i=1}^{N} \mathbb{1}_{\{i:i \in g\}} \left(SAR_i - \overline{SAR}_g \right)^2}{1 - \bar{\rho}_g} \tag{7}$$

For testing CARs, a robust test statistic is obtained by replacing the mean scaled

[41] $\mathbb{1}_{\{i:i \in g\}}$ is an the indicator function taking 1 for observations that are part of group g and zero otherwise.

abnormal return $\overline{SAR}_g$ with the mean scaled cumulative abnormal return ($SCAR$), and the standard deviation SD_g with the cross-sectional standard deviation of $SCAR$. Kolari and Pynnönen (2010) show their proposed test statistic outperforms other popular (parametric and non-parametric) tests, especially for longer CAR windows. For large estimation windows, this test statistic is approximately standard normal under the assumption of serially-independent jointly-normal abnormal returns, and an average (residual) cross-correlation $\bar{\rho}$ that goes to zero as the number of firms increases.

D. Parallel Trends

FIGURE D1
Tail risk for above and below $50 billion banks over time
This figure shows the eight-quarter moving average tail risk for banks with less than $50 billion in assets (BELOW_50B) and banks with assets equal or greater than $50 billion (ABOVE_50B) over the 2005-2017 period.