

Forecasting Crashes with a Smile


Ian Martin Ran Shi

May 2024

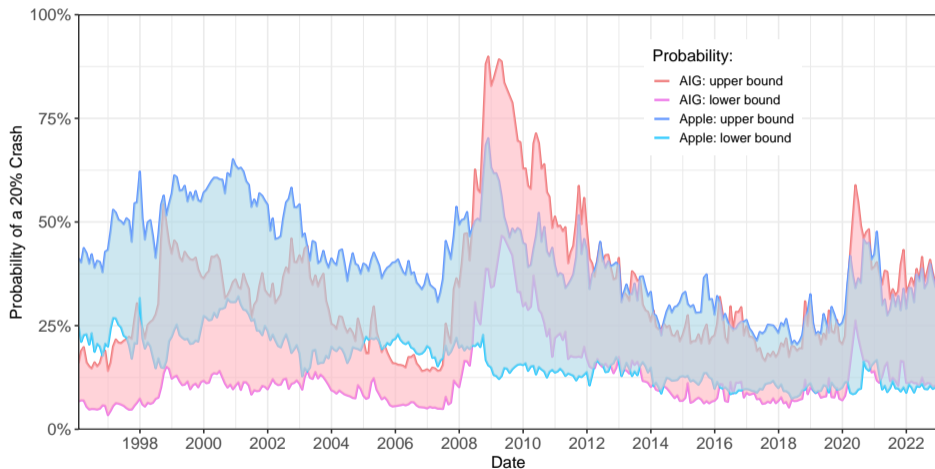
Introduction

- What is the probability that a given stock drops by 20% over the next month?
- We derive bounds on this quantity using index options and individual stock options
- No distributional assumptions
- The bounds are observable in real time
- They perform well in and out of sample


Probabilities of a 20% decline over the next month



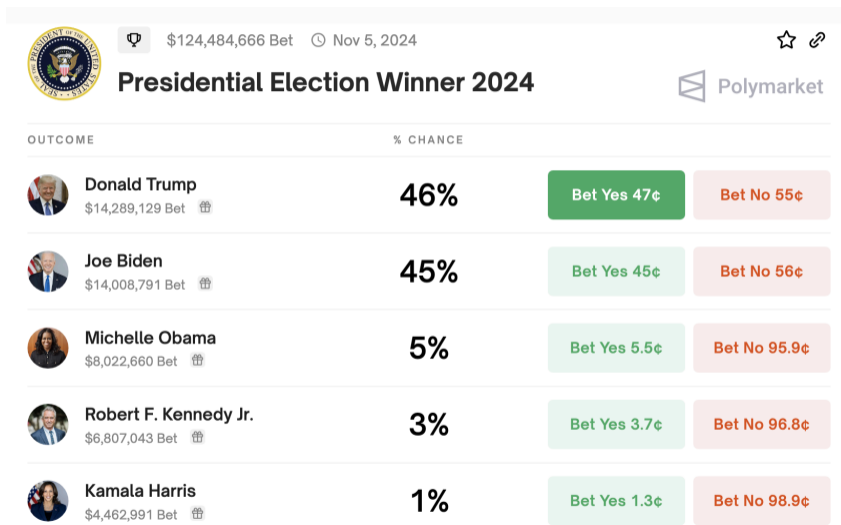
Probabilities of a 20% decline over the next year



Probabilities of a 20% decline over the next month



We can infer risk-neutral probabilities directly from asset prices



We can infer risk-neutral probabilities directly from asset prices


- The risk-neutral probability that the market declines by 20% over the next month can be calculated from index options expiring in a month

$$\mathbb{P}^*[R \leq 0.8] = R_f \times \underbrace{\frac{1}{R_f} \mathbb{E}^*[I(R \leq 0.8)]}_{\text{price of a binary option}} = R_f \times \underbrace{\text{put}'(0.8)}_{\text{slope of put prices}}$$

We can infer risk-neutral probabilities directly from asset prices

- The risk-neutral probability that the market declines by 20% over the next month can be calculated from index options expiring in a month


$$\mathbb{P}^*[R \leq 0.8] = R_f \times \underbrace{\frac{1}{R_f} \mathbb{E}^*[I(R \leq 0.8)]}_{\text{price of a binary option}} = R_f \times \underbrace{\text{put}'(0.8)}_{\text{slope of put prices}}$$



We can infer risk-neutral probabilities directly from asset prices

- The risk-neutral probability that the market declines by 20% over the next month can be calculated from index options expiring in a month


$$\mathbb{P}^*[R \leq 0.8] = R_f \times \underbrace{\frac{1}{R_f} \mathbb{E}^*[I(R \leq 0.8)]}_{\text{price of a binary option}} = R_f \times \underbrace{\text{put}'(0.8)}_{\text{slope of put prices}}$$



We can infer risk-neutral probabilities directly from asset prices

- The risk-neutral probability that the market declines by 20% over the next month can be calculated from index options expiring in a month


$$\mathbb{P}^*[R \leq 0.8] = R_f \times \underbrace{\frac{1}{R_f} \mathbb{E}^*[I(R \leq 0.8)]}_{\text{price of a binary option}} = R_f \times \underbrace{\text{put}'(0.8)}_{\text{slope of put prices}}$$



We can infer risk-neutral probabilities directly from asset prices

- The risk-neutral probability that the market declines by 20% over the next month can be calculated from index options expiring in a month


$$\mathbb{P}^*[R \leq 0.8] = R_f \times \underbrace{\frac{1}{R_f} \mathbb{E}^*[I(R \leq 0.8)]}_{\text{price of a binary option}} = R_f \times \underbrace{\text{put}'(0.8)}_{\text{slope of put prices}}$$



We can infer risk-neutral probabilities directly from asset prices

- The risk-neutral probability that the market declines by 20% over the next month can be calculated from index options expiring in a month

$$\mathbb{P}^*[R \leq 0.8] = R_f \times \underbrace{\frac{1}{R_f} \mathbb{E}^*[I(R \leq 0.8)]}_{\text{price of a binary option}} = R_f \times \underbrace{\text{put}'(0.8)}_{\text{slope of put prices}}$$



But we want **true**, not risk-neutral, probabilities

- We require an assumption (implicit or explicit) to link the true and risk-neutral probabilities—that is, about the stochastic discount factor
- We take the perspective of a one-period marginal investor with power utility who chooses to hold the market. So the SDF must be $M = R_m^{-\gamma} / \lambda$ for some constant λ
- The true expectation of a random payoff X then satisfies

$$\mathbb{E}[X] = \mathbb{E}[\underbrace{\lambda MR_m^\gamma}_{\equiv 1} X] = \lambda \mathbb{E}[M \times (R_m^\gamma X)] = \lambda \frac{\mathbb{E}^*[R_m^\gamma X]}{R_f}$$

- Eliminate λ by considering the case $X = 1$:

$$\mathbb{E}[X] = \frac{\mathbb{E}^*[R_m^\gamma X]}{\mathbb{E}^*[R_m^\gamma]}$$

But we want **true**, not risk-neutral, probabilities

- We have made a strong assumption on the form of the SDF
- In the case $\gamma = 0$, our approach simply forecasts using risk-neutral probabilities
 - ▶ Risk-neutral probabilities are widely used, and we will see that they forecast well
 - ★ But they overstate the true crash probabilities
 - ★ And the amount by which they overstate varies over time
- Bad news: Although our hypothetical investor understands market risk, he or she does not “know” about various anomalies demonstrated in the empirical asset pricing literature (. . . , momentum, value, profitability, . . .)
- Good news: We don't need to make the standard, undesirable, assumption that historical measures are good proxies for the forward-looking risk measures that come out of theory

Theory (1)

- Setting $X = I(R_i \leq q)$, this implies that the crash probability of stock i is

$$\mathbb{P}[R_i \leq q] = \frac{\mathbb{E}^* [R_m^\gamma I(R_i \leq q)]}{\mathbb{E}^* [R_m^\gamma]}$$

- To calculate $\mathbb{E}^* [R_m^\gamma]$, we need marginal distribution of R_m
 - ▶ Easy, using index option prices (Breedon and Litzenberger, 1978)
- To calculate $\mathbb{E}^* [R_m^\gamma I(R_i \leq q)]$, we need the joint distribution of (R_m, R_i)
 - ▶ **Problem:** Joint risk-neutral distribution is not observable given assets that are traded in practice (Martin, 2018, “Options and the Gamma Knife”)
 - ▶ This is a general theme: we are often interested in covariances and other features of the joint distribution

A 2×2 example

- Suppose the risk-neutral probability of a crash in Apple is 5%
- Suppose the risk-neutral probability of a crash in the market is also 5%
- These numbers can be calculated from options on Apple and options on the market
- But they are consistent with different joint distributions, eg,

| | | Apple | |
|---------|----------|-------|----------|
| | | Crash | No crash |
| S&P 500 | Crash | 5% | 0% |
| | No crash | 0% | 95% |

| | | Apple | |
|---------|----------|-------|----------|
| | | Crash | No crash |
| S&P 500 | Crash | 0% | 5% |
| | No crash | 5% | 90% |

A 2×2 example

| | | Apple | |
|---------|----------|-------|----------|
| | | Crash | No crash |
| S&P 500 | Crash | 5% | 0% |
| | No crash | 0% | 95% |

| | | Apple | |
|---------|----------|-------|----------|
| | | Crash | No crash |
| S&P 500 | Crash | 0% | 5% |
| | No crash | 5% | 90% |

- In the left-hand world, AAPL is risky
 - ▶ Risk-neutral probability of a crash will overstate the true probability of a crash
- In the right-hand world, AAPL is a hedge
 - ▶ Risk-neutral probability will understate the true probability of a crash
- Moral: Even if we can't observe the joint distribution, we may be able to derive bounds on the true crash probability

Theory (2)

$$\mathbb{P}[R_i \leq q] = \frac{\mathbb{E}^* [R_m^\gamma \mathbf{I}(R_i \leq q)]}{\mathbb{E}^* [R_m^\gamma]}$$

- We do not observe the joint risk-neutral distribution, so cannot calculate the right-hand side
- But we do observe the individual (marginal) risk-neutral distributions of R_m and R_i , from options on the market and on stock i
- The **Fréchet–Hoeffding theorem** provides upper and lower bounds on the right-hand side, as in the 2×2 example

Theory (3)

Result (Bounds on the probability of a crash)

We have

$$\frac{\mathbb{E}^* [R_m^\gamma \mathbf{I}(R_m \leq q_l)]}{\mathbb{E}^* [R_m^\gamma]} \leq \mathbb{P}[R_i \leq q] \leq \frac{\mathbb{E}^* [R_m^\gamma \mathbf{I}(R_m \geq q_u)]}{\mathbb{E}^* [R_m^\gamma]}$$

- The three elements are

$$\mathbb{E}^* [R_m^\gamma] = R_f^\gamma + \gamma(\gamma - 1)R_f \left[\int_0^{R_f} R^{\gamma-2} \text{put}_m(R) \, dR + \int_{R_f}^\infty R^{\gamma-2} \text{call}_m(R) \, dR \right]$$


$$\mathbb{E}^* [R_m^\gamma \mathbf{I}(R_m \leq q_l)] = R_f q_l^\gamma \left[\text{put}'_m(q_l) - \gamma \frac{\text{put}_m(q_l)}{q_l} \right] + \gamma(\gamma - 1)R_f \int_0^{q_l} R^{\gamma-2} \text{put}_m(R) \, dR$$

$$\mathbb{E}^* [R_m^\gamma \mathbf{I}(R_m \geq q_u)] = R_f q_u^\gamma \left[\gamma \frac{\text{call}_m(q_u)}{q_u} - \text{call}'(q_u) \right] + \gamma(\gamma - 1)R_f \int_{q_u}^\infty R^{\gamma-2} \text{call}_m(R) \, dR$$

Theory (4)

- The stock- i -specific quantiles q_l and q_u are such that

$$\mathbb{P}^*[R_m \leq q_l] = \mathbb{P}^*[R_i \leq q] = \mathbb{P}^*[R_m \geq q_u]$$



Theory (5)

- Bounds from the Fréchet–Hoeffding theorem are attainable in principle
 - ▶ Lower bound achieved for a stock that is **comonotonic** with the market—i.e., whose return is a (potentially nonlinear) increasing function of the market return
 - ▶ Upper bound achieved for a stock that is **countermonotonic** with the market—i.e., whose return is a (potentially nonlinear) decreasing function of the market return
- Intuitively, asset prices will tend to overstate crash probabilities if crashes are scary; or understate crash probabilities if crashes occur in good times
- A priori, we expect that the scary case is the relevant one, and hence that the lower bound should be closer to the truth in practice

Theory (8)

Result (Bounds on the probability of a crash)

We have

$$\frac{\mathbb{E}^* [R_m^\gamma \mathbf{I}(R_m \leq q_l)]}{\mathbb{E}^* [R_m^\gamma]} \leq \mathbb{P}[R_i \leq q] \leq \frac{\mathbb{E}^* [R_m^\gamma \mathbf{I}(R_m \geq q_u)]}{\mathbb{E}^* [R_m^\gamma]}$$

Further theoretical results

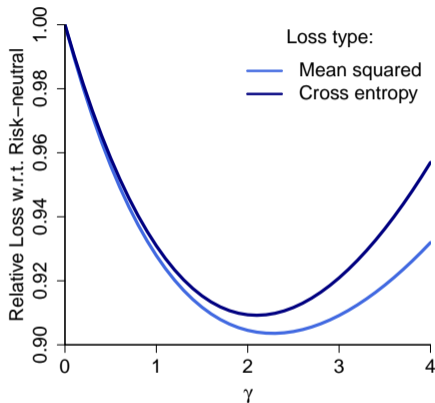
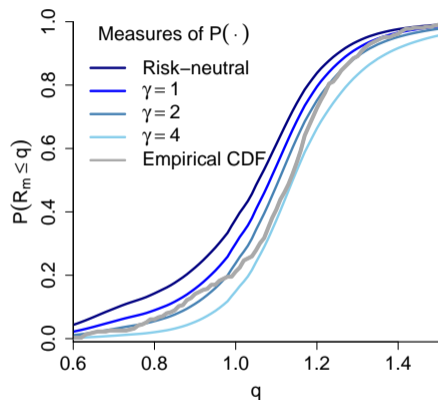
- Both $\mathbb{P}[R_i \leq q]$ and $\mathbb{P}^*[R_i \leq q]$ lie in between the bounds
- When $\gamma = 0$, the lower and upper bounds both equal $\mathbb{P}^*[R_i \leq q]$, and \mathbb{P}^* and \mathbb{P} coincide
- As γ increases, the bounds widen monotonically, so higher γ is more conservative
- As $\gamma \rightarrow \infty$, the bounds become trivial: the lower bound tends to zero and the upper bound tends to one

Data

- S&P 500 index and stock constituents from **Compustat**
- Risk-free rates and implied volatilities from **OptionMetrics**
 - ▶ Underlying stocks belonging to the S&P 500 index
 - ▶ Monthly from 1996/01 to 2022/12
 - ▶ Maturing in 1, 3, 6 and 12 months
 - ▶ On average around 492 firms each month
 - ▶ Over 155,000 firm-month observations per maturity
- Firm characteristics from **Compustat**
- Price, return, and volume data from **CRSP**
- Focus on “crashes” of 10%, 20% and 30% at horizons $\tau = 1, 3, 6$ and 12 months
- I'll often focus on the case of a 20% decline over one month

Calibrating risk aversion

We set $\gamma = 2$



- Left: unconditional CDF of market return for various γ ; and the empirical distribution
- Right: Loss functions when different values of γ are used to forecast market crashes

Summary statistics

| | | averaged across firms (number of obs. $T = 324$) | | | | averaged across time (number of obs. $N = 1044$) | | | |
|------------------------------|------|--|-------|-------|-------|--|-------|-------|-------|
| | | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 |
| maturity | | | | | | | | | |
| $q = 0.7$, down by over 30% | | | | | | | | | |
| realized | mean | 0.006 | 0.029 | 0.057 | 0.093 | 0.009 | 0.038 | 0.073 | 0.115 |
| | s.d. | 0.019 | 0.064 | 0.100 | 0.120 | 0.025 | 0.067 | 0.103 | 0.147 |
| lower bound | mean | 0.004 | 0.025 | 0.051 | 0.076 | 0.006 | 0.030 | 0.056 | 0.082 |
| | s.d. | 0.007 | 0.019 | 0.023 | 0.023 | 0.013 | 0.032 | 0.042 | 0.049 |
| upper bound | mean | 0.009 | 0.060 | 0.139 | 0.253 | 0.011 | 0.066 | 0.146 | 0.259 |
| | s.d. | 0.016 | 0.053 | 0.077 | 0.094 | 0.020 | 0.056 | 0.078 | 0.093 |
| risk-neutral | mean | 0.007 | 0.044 | 0.098 | 0.167 | 0.009 | 0.050 | 0.104 | 0.173 |
| | s.d. | 0.012 | 0.037 | 0.050 | 0.056 | 0.017 | 0.045 | 0.061 | 0.071 |

Summary statistics

| | | averaged across firms (number of obs. $T = 324$) | | | | averaged across time (number of obs. $N = 1044$) | | | |
|------------------------------|------|--|-------|-------|-------|--|-------|-------|-------|
| | | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 |
| maturity | | | | | | | | | |
| $q = 0.8$, down by over 20% | | | | | | | | | |
| realized | mean | 0.021 | 0.069 | 0.110 | 0.152 | 0.029 | 0.084 | 0.130 | 0.173 |
| | s.d. | 0.048 | 0.107 | 0.140 | 0.158 | 0.059 | 0.092 | 0.129 | 0.165 |
| lower bound | mean | 0.022 | 0.073 | 0.102 | 0.123 | 0.027 | 0.079 | 0.110 | 0.133 |
| | s.d. | 0.020 | 0.028 | 0.027 | 0.027 | 0.029 | 0.046 | 0.052 | 0.056 |
| upper bound | mean | 0.038 | 0.144 | 0.234 | 0.340 | 0.044 | 0.152 | 0.243 | 0.352 |
| | s.d. | 0.040 | 0.071 | 0.082 | 0.097 | 0.042 | 0.069 | 0.079 | 0.089 |
| risk-neutral | mean | 0.031 | 0.113 | 0.174 | 0.236 | 0.037 | 0.120 | 0.182 | 0.246 |
| | s.d. | 0.031 | 0.050 | 0.053 | 0.058 | 0.036 | 0.058 | 0.065 | 0.072 |

Summary statistics

| | | averaged across firms (number of obs. $T = 324$) | | | | averaged across time (number of obs. $N = 1044$) | | | |
|------------------------------|------|--|-------|-------|-------|--|-------|-------|-------|
| | | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 |
| maturity | | | | | | | | | |
| $q = 0.9$, down by over 10% | | | | | | | | | |
| realized | mean | 0.096 | 0.172 | 0.211 | 0.238 | 0.110 | 0.190 | 0.231 | 0.254 |
| | s.d. | 0.123 | 0.170 | 0.184 | 0.193 | 0.089 | 0.119 | 0.152 | 0.182 |
| lower bound | mean | 0.109 | 0.168 | 0.195 | 0.209 | 0.118 | 0.179 | 0.206 | 0.218 |
| | s.d. | 0.036 | 0.031 | 0.027 | 0.023 | 0.050 | 0.055 | 0.056 | 0.056 |
| upper bound | mean | 0.156 | 0.277 | 0.367 | 0.466 | 0.166 | 0.290 | 0.378 | 0.476 |
| | s.d. | 0.064 | 0.074 | 0.080 | 0.085 | 0.062 | 0.070 | 0.073 | 0.073 |
| risk-neutral | mean | 0.136 | 0.228 | 0.286 | 0.341 | 0.145 | 0.239 | 0.297 | 0.350 |
| | s.d. | 0.050 | 0.051 | 0.051 | 0.049 | 0.056 | 0.061 | 0.063 | 0.063 |

Empirical tests

- $I(R_i \leq q) = 0 + 1 \times \underbrace{\mathbb{E}[I(R_i \leq q)]}_{\mathbb{P}[R_i \leq q]} + \varepsilon$
- So a regression of the realized crash indicator $I(R_i \leq q)$ onto an ideal crash probability measure $\mathbb{P}[R_i \leq q]$ would yield zero constant term and a unit regression coefficient
- If the lower bound is close to the truth, then in a regression

$$I[R_{i,t \rightarrow t+\tau} \leq q] = \alpha^L + \beta^L \mathbb{P}_{i,t}^L(\tau, q) + \varepsilon_{i,t+\tau},$$

we should find $\alpha^L \approx 0$ and $\beta^L \approx 1$ at any horizon τ and for any crash size q

In-sample tests (1)

Down by 30% ($q = 0.7$)

| maturity | lower bound | | | | upper bound | | | | risk neutral | | | |
|----------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 |
| α | 0.00 (0.00) [0.00] | 0.00 (0.00) [0.00] | 0.00 (0.00) [0.01] | 0.01 (0.01) [0.01] | 0.00 (0.00) [0.00] | 0.00 (0.00) [0.00] | 0.00 (0.01) [0.01] | 0.01 (0.01) [0.02] | 0.00 (0.00) [0.00] | 0.00 (0.00) [0.00] | 0.00 (0.00) [0.01] | 0.00 (0.01) [0.01] |
| β | 0.95 (0.15) [0.16] | 1.03 (0.12) [0.17] | 1.09 (0.11) [0.18] | 1.05 (0.10) [0.14] | 0.51 (0.09) [0.09] | 0.43 (0.06) [0.09] | 0.39 (0.05) [0.09] | 0.35 (0.05) [0.09] | 0.66 (0.11) [0.11] | 0.60 (0.08) [0.11] | 0.59 (0.07) [0.12] | 0.56 (0.07) [0.12] |
| R^2 | 3.90% | 5.37% | 5.17% | 3.91% | 3.63% | 4.16% | 3.41% | 2.47% | 3.77% | 4.56% | 4.01% | 3.06% |

In-sample tests (1)

with time fixed effects

Down by 30% ($q = 0.7$)

| maturity | lower bound | | | | upper bound | | | | risk neutral | | | |
|-------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 |
| β | 0.93 (0.14) [0.15] | 1.05 (0.10) [0.14] | 1.11 (0.08) [0.12] | 1.14 (0.08) [0.11] | 0.55 (0.09) [0.10] | 0.55 (0.05) [0.08] | 0.58 (0.04) [0.06] | 0.60 (0.04) [0.06] | 0.68 (0.10) [0.11] | 0.70 (0.07) [0.10] | 0.74 (0.05) [0.07] | 0.78 (0.05) [0.07] |
| R^2 -proj | 3.27% | 4.81% | 5.06% | 4.54% | 3.16% | 4.39% | 4.74% | 4.43% | 3.21% | 4.52% | 4.87% | 4.50% |

In-sample tests (2)

Down by 20% ($q = 0.8$)

| maturity | lower bound | | | | upper bound | | | | risk neutral | | | |
|----------|--------------------------|---------------------------|---------------------------|--------------------------|--------------------------|---------------------------|---------------------------|--------------------------|--------------------------|---------------------------|---------------------------|--------------------------|
| | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 |
| α | 0.00 (0.00) [0.00] | -0.01 (0.01) [0.01] | -0.01 (0.01) [0.01] | 0.02 (0.01) [0.01] | 0.00 (0.00) [0.00] | -0.01 (0.01) [0.01] | -0.01 (0.01) [0.02] | 0.01 (0.02) [0.03] | 0.00 (0.00) [0.00] | -0.01 (0.01) [0.01] | -0.02 (0.01) [0.01] | 0.00 (0.01) [0.02] |
| β | 0.92 (0.11) [0.12] | 1.03 (0.09) [0.14] | 1.15 (0.09) [0.12] | 1.07 (0.08) [0.13] | 0.56 (0.08) [0.08] | 0.51 (0.06) [0.08] | 0.49 (0.06) [0.10] | 0.41 (0.06) [0.10] | 0.68 (0.09) [0.08] | 0.69 (0.07) [0.11] | 0.73 (0.07) [0.10] | 0.66 (0.07) [0.12] |
| R^2 | 5.65% | 5.15% | 4.76% | 3.69% | 5.32% | 4.11% | 3.22% | 2.30% | 5.48% | 4.50% | 3.89% | 2.96% |

In-sample tests (2)

with time fixed effects

Down by 20% ($q = 0.8$)

| maturity | lower bound | | | | upper bound | | | | risk neutral | | | |
|-------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 |
| β | 0.93 (0.09) [0.10] | 1.03 (0.07) [0.10] | 1.13 (0.06) [0.10] | 1.10 (0.06) [0.09] | 0.62 (0.06) [0.06] | 0.67 (0.04) [0.07] | 0.74 (0.04) [0.06] | 0.71 (0.04) [0.06] | 0.73 (0.07) [0.08] | 0.80 (0.05) [0.07] | 0.89 (0.05) [0.07] | 0.87 (0.05) [0.07] |
| R^2 -proj | 4.49% | 4.65% | 4.55% | 4.01% | 4.33% | 4.45% | 4.40% | 3.98% | 4.39% | 4.53% | 4.48% | 4.00% |

Intermission: Probability of a rise of at least 20%

| maturity | lower bound | | | | upper bound | | | | risk neutral | | | |
|----------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 |
| α | 0.00 (0.00) [0.00] | 0.01 (0.00) [0.01] | 0.09 (0.01) [0.01] | 0.34 (0.02) [0.03] | 0.00 (0.00) [0.00] | 0.00 (0.01) [0.01] | 0.03 (0.01) [0.02] | 0.20 (0.03) [0.04] | 0.00 (0.00) [0.00] | 0.00 (0.01) [0.01] | 0.04 (0.01) [0.02] | 0.23 (0.03) [0.03] |
| β | 1.35 (0.13) [0.13] | 1.58 (0.11) [0.16] | 1.30 (0.11) [0.19] | 0.10 (0.14) [0.19] | 0.85 (0.09) [0.10] | 0.91 (0.08) [0.11] | 0.82 (0.08) [0.10] | 0.44 (0.09) [0.13] | 1.03 (0.11) [0.11] | 1.17 (0.09) [0.12] | 1.08 (0.09) [0.14] | 0.49 (0.12) [0.17] |
| R^2 | 7.01% | 5.78% | 2.44% | 0.01% | 7.42% | 6.86% | 4.24% | 0.80% | 7.35% | 6.70% | 3.80% | 0.43% |

- For rises, the **upper** bound would be tight in the comonotonic case
- At the one year horizon, it is harder to predict booms than crashes (perhaps because booms are more idiosyncratic so comonotonicity is further from the truth)

In-sample tests (3)

Down by 10% ($q = 0.9$)

| maturity | lower bound | | | | upper bound | | | | risk neutral | | | |
|----------|---------------------------|---------------------------|---------------------------|--------------------------|---------------------------|--------------------------|--------------------------|--------------------------|---------------------------|---------------------------|---------------------------|--------------------------|
| | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 |
| α | -0.02 (0.01) [0.01] | -0.01 (0.01) [0.01] | -0.01 (0.01) [0.02] | 0.03 (0.02) [0.03] | -0.02 (0.01) [0.01] | 0.00 (0.02) [0.03] | 0.01 (0.02) [0.03] | 0.05 (0.03) [0.05] | -0.02 (0.01) [0.01] | -0.02 (0.02) [0.02] | -0.02 (0.02) [0.03] | 0.00 (0.03) [0.04] |
| β | 1.05 (0.08) [0.08] | 1.07 (0.07) [0.10] | 1.12 (0.07) [0.11] | 1.01 (0.08) [0.11] | 0.75 (0.07) [0.07] | 0.63 (0.07) [0.10] | 0.54 (0.07) [0.10] | 0.41 (0.08) [0.12] | 0.88 (0.08) [0.08] | 0.83 (0.08) [0.12] | 0.80 (0.08) [0.12] | 0.68 (0.09) [0.13] |
| R^2 | 5.46% | 3.71% | 3.38% | 2.41% | 5.35% | 3.03% | 2.16% | 1.23% | 5.46% | 3.39% | 2.80% | 1.83% |

In-sample tests (3)


with time fixed effects

Down by 10% ($q = 0.9$)


| maturity | lower bound | | | | upper bound | | | | risk neutral | | | |
|-------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 |
| β | 0.99 (0.06) [0.06] | 0.99 (0.05) [0.07] | 1.05 (0.06) [0.08] | 1.05 (0.06) [0.08] | 0.80 (0.05) [0.06] | 0.79 (0.04) [0.06] | 0.83 (0.04) [0.06] | 0.82 (0.05) [0.06] | 0.88 (0.05) [0.05] | 0.89 (0.05) [0.06] | 0.94 (0.05) [0.07] | 0.93 (0.05) [0.08] |
| R^2 -proj | 4.02% | 3.15% | 3.14% | 2.85% | 3.96% | 3.08% | 3.09% | 2.82% | 3.99% | 3.12% | 3.12% | 2.83% |

Estimated β , by year: lower bound


20% return drop: 1 mo. ahead




20% return drop: 3 mo. ahead







20% return drop: 6 mo. ahead



20% return drop: 12 mo. ahead



Estimated β , by year: risk-neutral probabilities



Lower bound vs. risk-neutral probabilities

- Risk-neutral probabilities overstate true crash probabilities
- The extent to which they overstate varies over time
- We should expect risk-neutral probabilities to overstate most—hence estimated β coefficients to be lowest—in scary times
- The lower bound adjusts for scariness, so estimated β coefficients are more stable
- This gives the lower bound an advantage when we look at OOS performance

Fréchet–Hoeffding vs. Cauchy–Schwarz

- Here's another approach that also does not work as well. Write

$$\mathbb{P}[R_i \leq q] = \mathbb{P}^*[R_i \leq q] + \frac{\text{cov}^*[R_m^\gamma, \mathbf{I}(R_i \leq q)]}{\mathbb{E}^*[R_m^\gamma]}$$

- As correlation must lie between plus and minus one, it follows that

$$\mathbb{P}^*[R_i \leq q] - \frac{\sigma^*[R_m^\gamma] \sigma^*[\mathbf{I}(R_i \leq q)]}{\mathbb{E}^*[R_m^\gamma]} \leq \mathbb{P}[R_i \leq q] \leq \mathbb{P}^*[R_i \leq q] + \frac{\sigma^*[R_m^\gamma] \sigma^*[\mathbf{I}(R_i \leq q)]}{\mathbb{E}^*[R_m^\gamma]}$$

where $\sigma^*[\cdot]$ denotes risk-neutral volatility

- These bounds depend only on univariate risk-neutral expectations, so can be calculated from observable option prices. But they are much wider than our bounds

Fréchet–Hoeffding vs. Cauchy–Schwarz

- If, say, returns were jointly lognormal, then it *could* in principle be the case that log returns were perfectly positively or negatively correlated
- But observed option prices rule out lognormality
- They also bound correlations away from ± 1
- Are we just using the fact that correlations lie in $[-1, 1]$? No!


Width of FH bounds relative to CS bounds

| crash size | horizon | mean | sd | median | q25 | q75 | min | max |
|------------|---------|-------|-------|--------|-------|-------|-------|-------|
| 30% | 1 | 0.098 | 0.138 | 0.035 | 0.004 | 0.137 | 0.000 | 0.799 |
| 30% | 3 | 0.354 | 0.190 | 0.345 | 0.198 | 0.503 | 0.000 | 0.815 |
| 30% | 6 | 0.531 | 0.148 | 0.558 | 0.438 | 0.645 | 0.000 | 0.812 |
| 30% | 12 | 0.642 | 0.094 | 0.656 | 0.607 | 0.704 | 0.000 | 0.816 |
| 20% | 1 | 0.271 | 0.184 | 0.247 | 0.113 | 0.410 | 0.000 | 0.800 |
| 20% | 3 | 0.561 | 0.127 | 0.592 | 0.490 | 0.648 | 0.000 | 0.813 |
| 20% | 6 | 0.658 | 0.075 | 0.662 | 0.623 | 0.706 | 0.000 | 0.811 |
| 20% | 12 | 0.704 | 0.057 | 0.711 | 0.672 | 0.745 | 0.001 | 0.842 |
| 10% | 1 | 0.544 | 0.108 | 0.565 | 0.487 | 0.618 | 0.000 | 0.848 |
| 10% | 3 | 0.679 | 0.059 | 0.678 | 0.642 | 0.723 | 0.000 | 0.828 |
| 10% | 6 | 0.727 | 0.043 | 0.733 | 0.698 | 0.761 | 0.000 | 0.812 |
| 10% | 12 | 0.751 | 0.032 | 0.758 | 0.736 | 0.772 | 0.002 | 0.842 |


| | 20% crashes over one month | | |
|-------------------|----------------------------|-----------------|-----------------|
| Fréchet–Hoeffding | 0.84 (0.54) | 1.27 (0.84) | 1.55 (1.47) |
| Cauchy–Schwarz | 0.09 (0.56) | | −0.20 (0.65) |
| risk-neutral | | −0.26 (0.66) | −0.35 (0.80) |
| constant | 0.00 (0.00) | 0.00 (0.00) | −0.00 (0.00) |

Estimated β , by industry: lower bound


20% return drop: 1 mo. ahead




20% return drop: 3 mo. ahead



20% return drop: 6 mo. ahead



20% return drop: 12 mo. ahead



Competitor variables from the literature


- We compare against 15 variables drawn from the literature
 - ▶ Stock characteristics: CAPM beta, (log) relative size, book-to-market, profitability (gross profit/assets), momentum (prior 2-6 and 2-12 month returns), lagged return
 - ▶ Chen–Hong–Stein, 2001: realized volatilities (standard deviations of daily market-adjusted returns over the last six months) and monthly turnover (shares traded scaled by shares outstanding)
 - ▶ Greenwood–Shleifer–You, 2019: sales growth
 - ▶ Asquith–Pathak–Ritter, 2005; Nagel, 2005: short interest (shares shorted/shares held by institutions)
 - ▶ Campbell–Hilscher–Szilagyi, 2008: leverage (debt/total assets), net income/total assets, cash/total assets, log price per share (winsorized from above at \$15)
- All variables are standardized to unit standard deviation for comparability

In-sample tests (4)


Asterisks indicate t -statistics above 4

| | | $I(R_{t \rightarrow t+1} \leq 0.8)$ | | | | |
|--|--------|-------------------------------------|--------|--------|--------|--------|
| $\mathbb{P}^L[R_{t \rightarrow t+1} \leq 0.8]$ | | 3.41* | 3.05* | | 4.44 | 2.74* |
| | | (0.41) | (0.59) | | (3.08) | (0.33) |
| $\mathbb{P}^*[R_{t \rightarrow t+1} \leq 0.8]$ | | | | 2.83* | -1.40 | |
| | | | | (0.67) | (3.37) | |
| CHS-volatility | 2.28* | | 0.30 | 0.43 | 0.31 | 0.50 |
| | (0.31) | | (0.38) | (0.45) | (0.39) | (0.18) |
| short int. | 0.39* | | 0.33* | 0.36* | 0.32* | 0.27* |
| | (0.09) | | (0.08) | (0.08) | (0.08) | (0.06) |
| | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ |
| R^2/R^2 -proj. | 4.51% | 5.66% | 5.85% | 5.72% | 5.87% | 4.74% |

In-sample tests (4)



In-sample tests (4)



Back–Crotty–Kazempour (2022)

- GMM-based tests for the validity and tightness of bounds, applied to Martin (2017), Martin–Wagner (2019), Kadan–Tang (2020), Chabi-Yo–Loudis (2020)
- Conclusions:
 - ▶ Our upper and lower bounds are valid
 - ▶ Our upper bound is (with very high confidence) not tight
 - ▶ Mixed evidence on tightness of the lower bound

BCK tests

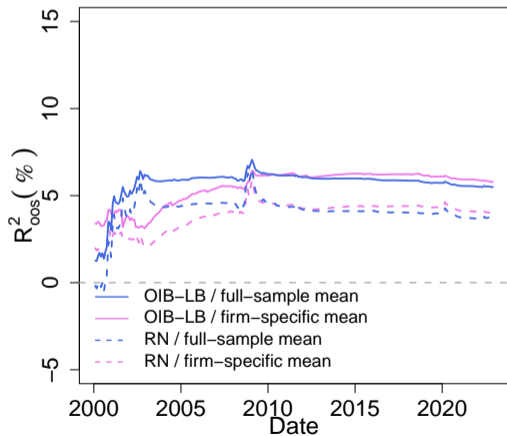
p-values for tests of validity and tightness

| horizon | lower bound | | | | upper bound | | | |
|---------------------------------------|-------------|-------|-------|-------|-------------|-------|-------|-------|
| | 1 | 3 | 6 | 12 | 1 | 3 | 6 | 12 |
| Panel A: $q = 0.7$, down by over 30% | | | | | | | | |
| Validity | 0.691 | 1.000 | 0.512 | 0.430 | 0.763 | 0.781 | 0.774 | 0.752 |
| Tightness | 0.414 | 0.118 | 0.039 | 0.157 | 0.316 | 0.009 | 0.000 | 0.016 |
| Panel B: $q = 0.8$, down by over 20% | | | | | | | | |
| Validity | 0.462 | 0.375 | 0.621 | 0.502 | 1.000 | 1.000 | 0.751 | 0.754 |
| Tightness | 0.348 | 0.022 | 0.043 | 0.161 | 0.011 | 0.000 | 0.000 | 0.018 |
| Panel C: $q = 0.9$, down by over 10% | | | | | | | | |
| Validity | 0.068 | 0.634 | 0.686 | 0.490 | 1.000 | 1.000 | 0.760 | 0.753 |
| Tightness | 0.134 | 0.059 | 0.058 | 0.116 | 0.000 | 0.000 | 0.000 | 0.019 |


Out-of-sample forecasts

R^2 , expanding window, competing against in-sample mean crash probabilities

20% crash in 1 months




20% crash in 3 months



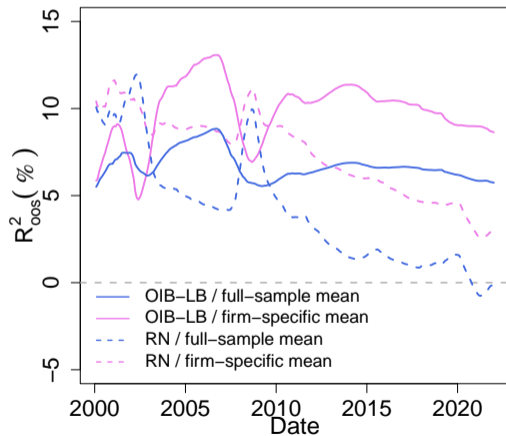
Out-of-sample forecasts

R^2 , expanding window, competing against in-sample mean crash probabilities

20% crash in 6 months



20% crash in 12 months



Out-of-sample forecasts


A more challenging competitor

- Competitor uses all 15 additional variables together with risk-neutral probabilities
- We train predictive models using data from 1996 to 2006/2011/2016
 - ▶ variable selection using Lasso
 - ▶ tuning parameters for sparsity: 10-fold cross validation based on the training sample
- Then make out-of-sample forecasts for the rest of the sample
- Option-implied bounds are directly used to forecast with fixed $\alpha = 0$ and $\beta = 1$
- Performance measure: (area under) ROC curves


Out-of-sample forecasts: ROC curves

Training Sample 1996 – 2006

Down by over 20% in one month




Down by over 20% in one year




Out-of-sample forecasts: ROC curves

Training Sample 1996 – 2011

Down by over 20% in one month




Down by over 20% in one year




Out-of-sample forecasts: ROC curves

Training Sample 1996 – 2016

Down by over 20% in one month



Down by over 20% in one year



Industry average crash probabilities




- Substantial variation in crash probability over time and across industries
- News about crash risk is not just idiosyncratic: related industries' probabilities comove

Industry average crash probabilities



- Substantial variation in crash probability over time and across industries
- News about crash risk is not just idiosyncratic: related industries' probabilities comove

Industry average crash probabilities



- Substantial variation in crash probability over time and across industries
- News about crash risk is not just idiosyncratic: related industries' probabilities comove

Summary

- We derive bounds on crash probabilities and show that the lower bound successfully forecasts crashes in and out of sample
- For one month forecasts of 20% crashes, we find
 - ▶ t -stats in the range 5 to 13
 - ▶ estimated coefficient is 10 times larger than the next most important competitor variable
- Risk-neutral probabilities also perform well in sample, but overstate crash probabilities—and time variation in overstatement hurts out-of-sample performance
- Our results depend on one key assumption: the form of the SDF
- This is a strong assumption, but it allows us to avoid the undesirable (and commonly made) assumption that backward-looking historical measures are good proxies for the forward-looking measures that come out of theory
- It seems the price of our assumption is worth paying

Earnings months have little effect on our results

| | $I(R_{t \rightarrow t+1} \leq q)$ | | |
|--|-----------------------------------|-----------------|-----------------|
| | $q = 0.70$ | $q = 0.80$ | $q = 0.90$ |
| OIB-LB | 0.93 (0.17) | 0.93 (0.12) | 1.10 (0.09) |
| $I(\text{earnings in mo. } t + 1)$ | 0.00 (0.00) | 0.00 (0.00) | 0.02 (0.01) |
| OIB-LB \times $I(\text{earnings in mo. } t + 1)$ | 0.08 (0.21) | -0.01 (0.13) | -0.13 (0.09) |
| constant | 0.00 (0.00) | -0.00 (0.00) | -0.02 (0.01) |