

Forecasting Crashes with a Smile


Ian Martin Ran Shi

December 2025

What is the chance that Apple stock drops 20% over the next month?

- We derive bounds on this quantity using index options and individual stock options
- The bounds are observable in real time
- We argue that the lower bound should be expected to be closer to the truth
- And show that it forecasts well in and out of sample

Probabilities of a 20% decline over the next month



Today

- ① Theory
- ② Data
- ③ In-sample tests
- ④ Out-of-sample tests and the “crying wolf” problem
- ⑤ Industry crash risk series
- ⑥ Explaining crash probabilities

Theory

Information in market prices

- Market prices are often used to predict things
- We talk about
 - ▶ forward rates
 - ▶ CDS rates
 - ▶ implied volatility
 - ▶ breakeven inflation
 - ▶ ...
- These are almost continuously observable
- Don't need to rely on economists' models
- And they embody the collective views of market participants
- But they may be distorted by risk: people will pay more for insurance/hedge assets that pay off in scary states of the world

Information in market prices

- Market prices are often used to predict things
- We talk about
 - ▶ forward rates → risk-neutral expected future interest rates
 - ▶ CDS rates → risk-neutral default probabilities
 - ▶ implied volatility → risk-neutral volatility
 - ▶ breakeven inflation → risk-neutral expected future inflation
 - ▶ ...
- These are almost continuously observable
- Don't need to rely on economists' models
- And they embody the collective views of market participants
- But they may be distorted by risk: people will pay more for insurance/hedge assets that pay off in scary states of the world

We can infer risk-neutral probabilities directly from asset prices



Fed decision in December?



🏆 \$186,495,262 Vol. 🕒 Dec 10, 2025

Past ▾

Dec 10


Jan 28, 2026

Mar 18, 2026

Apr 29, 2026

● 25 bps decrease 88% ● No change 12% ● 50+ bps decrease 1.1% ● 25+ bps increase <1%

 Polymarket



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

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


option 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

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



Strengths and weaknesses of risk-neutral probabilities

- Risk-neutral probabilities perform quite well in forecasting crashes
- But they overstate the probability of a crash
- And the extent to which they overstate varies
- They overstate most in scary times and for scary (\approx high beta) stocks
- This is unfortunate! These are the situations, and stocks, for which a crash indicator is most useful

So 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

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

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

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

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

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

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

- ▶ γ is the investor's risk aversion
 - ▶ In the case $\gamma = 0$, we are back to the risk-neutral probabilities
- Good news: We avoid the standard, undesirable, assumption that historical measures are good proxies for the forward-looking risk measures that come out of theory

Calculating crash probabilities

- 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 risk-neutral 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
 - ▶ The case $i = m$ is easy. But *testing* the theory is hard because crashes are rare

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

Bounding crash probabilities

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

- 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
- Using these, the **Fréchet–Hoeffding theorem** provides upper and lower bounds on the right-hand side, as in the 2×2 example

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

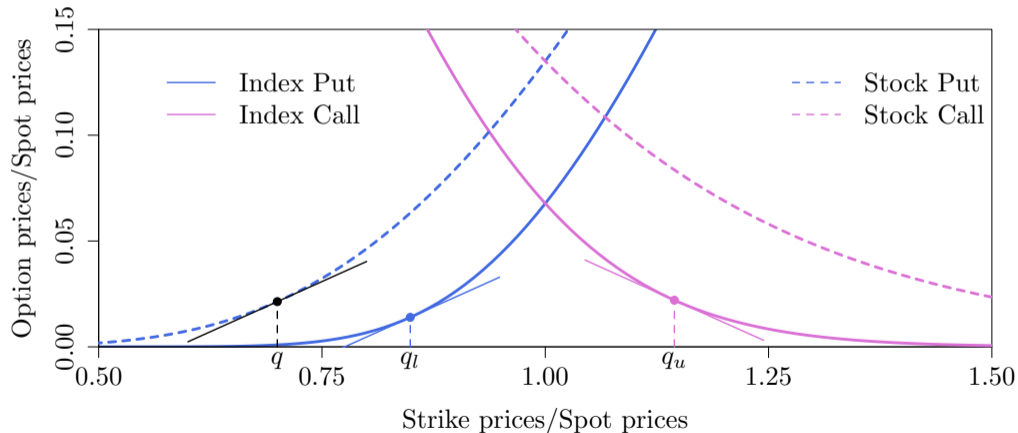
$$\mathbb{E}^* [R_m^\gamma] = R_f^\gamma + \gamma(\gamma - 1)R_f \left[\int_0^{R_f} K^{\gamma-2} \text{put}_m(K) \, dK + \int_{R_f}^\infty K^{\gamma-2} \text{call}_m(K) \, dK \right]$$

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

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

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

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



The upper and lower bounds 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**—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: summary

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

Data

- S&P 500 index and stock constituents from **Compustat**
- Risk-free rates and implied volatilities from **OptionMetrics**
 - ▶ Monthly from 1996/01 to 2022/12
 - ▶ On average around 492 firms each month
 - ▶ Options maturing in 1, 3, 6 and 12 months
 - ▶ 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
- We set risk aversion, γ , equal to 2

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
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
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

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
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
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

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
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
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

In-sample tests

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

In-sample tests (1)

with time fixed effects

Down by 30% ($q = 0.7$)

maturity	lower bound				risk-neutral				upper bound			
	1	3	6	12	1	3	6	12	1	3	6	12
β	0.93 (0.14) [0.16]	1.05 (0.10) [0.13]	1.11 (0.08) [0.12]	1.14 (0.08) [0.11]	0.68 (0.10) [0.13]	0.70 (0.07) [0.09]	0.74 (0.05) [0.10]	0.78 (0.05) [0.07]	0.55 (0.09) [0.09]	0.55 (0.05) [0.07]	0.58 (0.04) [0.06]	0.60 (0.04) [0.06]
R^2 -proj	3.27%	4.81%	5.06%	4.54%	3.21%	4.52%	4.87%	4.50%	3.16%	4.39%	4.74%	4.43%

In-sample tests (2)

Down by 20% ($q = 0.8$)

maturity	lower bound				risk-neutral				upper bound			
	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) [0.02]	0.00 (0.00) [0.00]	-0.01 (0.01) [0.01]	-0.02 (0.01) [0.01]	0.00 (0.01) [0.02]	0.00 (0.00) [0.00]	-0.01 (0.01) [0.01]	-0.01 (0.01) [0.02]	0.01 (0.02) [0.03]
β	0.92 (0.11) [0.11]	1.03 (0.09) [0.13]	1.15 (0.09) [0.15]	1.07 (0.08) [0.13]	0.68 (0.09) [0.09]	0.69 (0.07) [0.10]	0.73 (0.07) [0.11]	0.66 (0.07) [0.12]	0.56 (0.08) [0.07]	0.51 (0.06) [0.08]	0.49 (0.06) [0.10]	0.41 (0.06) [0.10]
R^2	5.65%	5.15%	4.76%	3.69%	5.48%	4.50%	3.89%	2.96%	5.32%	4.11%	3.22%	2.30%

In-sample tests (2)

with time fixed effects

Down by 20% ($q = 0.8$)

maturity	lower bound				risk-neutral				upper bound			
	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) [0.09]	1.10 (0.06) [0.09]	0.73 (0.07) [0.07]	0.80 (0.05) [0.07]	0.89 (0.05) [0.07]	0.87 (0.05) [0.06]	0.62 (0.06) [0.07]	0.67 (0.04) [0.07]	0.74 (0.04) [0.07]	0.71 (0.04) [0.06]
R^2 -proj	4.49%	4.65%	4.55%	4.01%	4.39%	4.53%	4.48%	4.00%	4.33%	4.45%	4.40%	3.98%

Intermission: Probability of a rise of at least 20%

maturity	lower bound				risk-neutral				upper bound			
	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]	0.04 (0.01) [0.02]	0.24 (0.03) [0.04]	0.00 (0.00) [0.00]	-0.01 (0.01) [0.01]	0.03 (0.01) [0.02]	0.21 (0.03) [0.04]
β	1.35 (0.13) [0.13]	1.58 (0.11) [0.14]	1.32 (0.11) [0.15]	0.12 (0.15) [0.21]	1.03 (0.10) [0.11]	1.17 (0.09) [0.13]	1.08 (0.09) [0.15]	0.46 (0.12) [0.17]	0.85 (0.09) [0.09]	0.91 (0.08) [0.10]	0.82 (0.07) [0.11]	0.42 (0.09) [0.13]
R^2	6.95%	5.78%	2.51%	0.01%	7.28%	6.66%	3.79%	0.38%	7.36%	6.81%	4.21%	0.72%

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

In-sample tests (3)


with time fixed effects

Down by 10% ($q = 0.9$)


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

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


20% return drop: 1 mo. ahead




20% return drop: 3 mo. ahead



20% return drop: 6 mo. ahead



20% return drop: 12 mo. ahead



This is *not* just “correlation must lie between plus and minus one”

- Here’s a different approach that does not work as well. Write

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

- Risk-neutral correlation, $\rho^*[R_m^\gamma, \mathbf{I}(R_i \leq q)]$, must lie between plus and minus one, so


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

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


- But $\rho^*[R_m^\gamma, \mathbf{I}(R_i \leq q)]$ **cannot** reach plus or minus one: one variable is continuous, the other discrete!
- So these correlation-based bounds are very weak: typically around three times wider for the 1 month/20% pair and around 10 times wider for the 1 month/30% pair

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 models that exploit 15 characteristics drawn from the literature
 - ▶ Usual suspects: 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

		$I(R_{t \rightarrow t+1} \leq 0.8)$				
$\mathbb{P}^L[R_{t \rightarrow t+1} \leq 0.8]$		3.40*	3.02*		4.41	2.72*
		(0.41)	(0.58)		(3.08)	(0.33)
$\mathbb{P}^*[R_{t \rightarrow t+1} \leq 0.8]$				2.81*	-1.39	
				(0.66)	(3.36)	
CHS-volatility	2.27*		0.31	0.44	0.32	0.50
	(0.31)		(0.37)	(0.44)	(0.39)	(0.18)
short int.	0.39*		0.34*	0.37*	0.33*	0.27*
	(0.09)		(0.08)	(0.08)	(0.08)	(0.06)
	⋮		⋮	⋮	⋮	⋮
R^2/R^2 -proj.	4.49%	5.65%	5.82%	5.69%	5.83%	4.72%

In-sample tests (4)



In-sample tests (4)



Out-of-sample tests


We compare OOS forecast performance of two models

- ① Competitor model uses 15 characteristics, lower bound, and risk-neutral probabilities
 - ▶ We train predictive models using expanding or rolling windows
 - ★ variable selection using elastic net
 - ★ tuning parameters for sparsity: 5-fold cross validation based on the training sample
 - ▶ Then make out-of-sample forecasts for the rest of the sample
- ② Our lower bound, directly used to forecast with fixed $\alpha = 0$ and $\beta = 1$
 - ▶ **Nothing** is estimated
- Performance measure: out-of-sample R^2
 - ▶ Diebold–Mariano tests reject the null of equal forecasting accuracy
- Similar results for a “kitchen sink” competitor that also uses interactions and squares of the 15 original characteristics (for a total of 137 variables)
- Also for a simpler competitor that attempts to rescale the risk-neutral probabilities optimally using backward-looking information


Out-of-sample forecasts

R^2 , **expanding** window, competing against in-sample mean crash probabilities (firm-specific)


20% crash in 1 months




20% crash in 3 months



20% crash in 6 months




20% crash in 12 months



Out-of-sample forecasts


$\hat{\beta}$, expanding window, competing against in-sample mean crash probabilities (firm-specific)




Out-of-sample forecasts

R^2 , 3yr rolling window, competing against in-sample mean crash probabilities (firm-specific)


20% crash in 1 months




20% crash in 3 months



20% crash in 6 months




20% crash in 12 months




Out-of-sample forecasts

$\hat{\beta}$, 3yr rolling window, competing against in-sample mean crash probabilities (firm-specific)



Industry crash risk indices

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

Explaining crash probabilities



Explaining crash probabilities


- If you accept the lower bound as a tolerable measure of crash risk, then we can use it to “de-noise” the realized crash event indicator
- This boosts power to detect variables that influence a stock’s likelihood of crashing: we find R^2 on the order of 70–75%
- Crash risk is higher for
 - ▶ stocks with high beta, CHS volatility, share turnover, and short interest (Chen, Hong and Stein, 2001; Hong and Stein, 2003)
 - ▶ penny stocks, with low log share price (Campbell, Hilscher and Szilagyi, 2008)
 - ▶ stocks with poor recent returns, either over the past month or from month -6 to -1
 - ▶ unprofitable stocks or stocks with low net income

Panel A: Regressions of the lower bound onto 15 characteristics



Panel B: Regressions onto characteristics, squared characteristics, and interaction terms





Interactions

(c) Interaction terms

	est.	s.e.	t-stat.
cash \times log.p	0.113	0.014	8.10
bm \times gprof	0.108	0.015	7.32
chs.vol \times log.p	-0.154	0.028	-5.52
s.int \times ni	-0.056	0.011	-5.26
chs.vol \times beta	0.173	0.036	4.76
rsize \times log.p	0.088	0.022	4.07
lev \times rlag1	-0.056	0.015	-3.81
beta \times s.int	0.053	0.015	3.60
	⋮		

- Crash risk is higher for

- ▶ penny stocks with low cash, and penny stocks that are volatile or small
- ▶ unprofitable growth stocks
- ▶ stocks with high short interest and low net income
- ▶ high beta stocks with high CHS volatility or short interest
- ▶ high leverage stocks with low lagged returns

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 10 times larger than the next most important competitor variable
- Risk-neutral probabilities perform well in sample, but overstate crash probabilities, and time variation in overstatement hurts OOS performance: “crying wolf” problem
- Our approach depends on one strong assumption: the form of the SDF
- Allows us to avoid the undesirable 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