Forecasting Crashes with a Smile

Ian Martin Ran Shi

July, 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
- No distributional assumptions
- 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

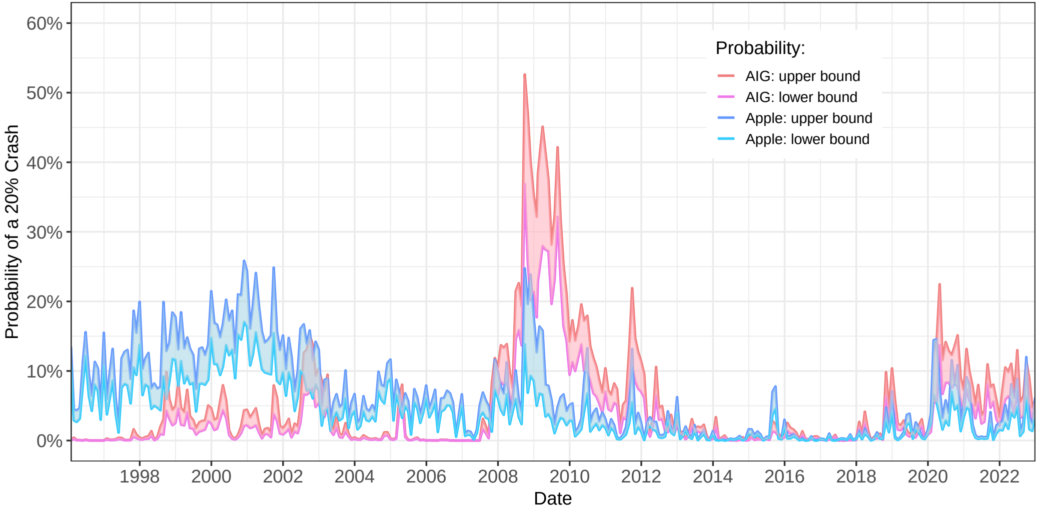

Today

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 for forecasting:
  - 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 for forecasting:
  - forward rates $\longrightarrow$ risk-neutral expected future interest rates
  - CDS rates $\longrightarrow$ risk-neutral default probabilities
  - implied volatility $\longrightarrow$ risk-neutral volatility
  - breakeven inflation $\longrightarrow$ 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

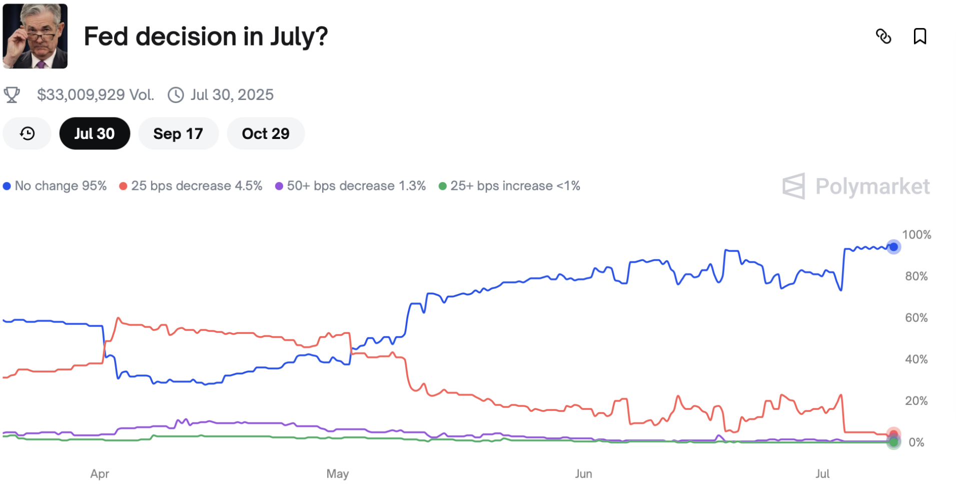

We can infer risk-neutral probabilities directly from asset prices

... almost in real time

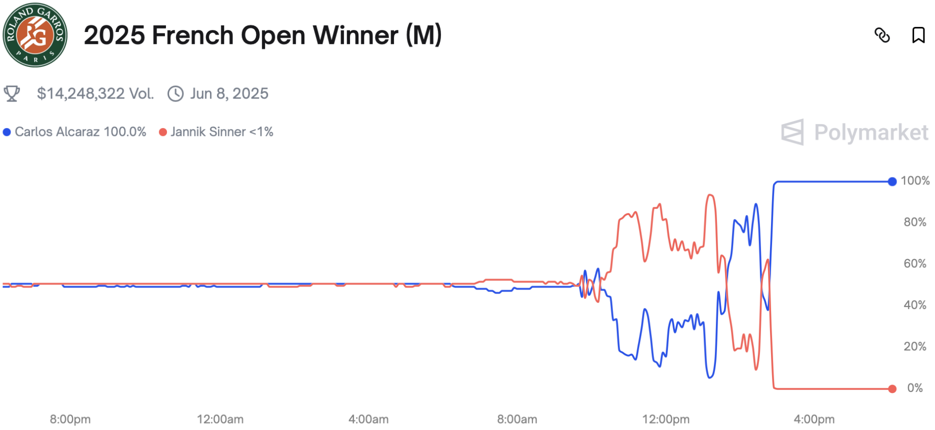

We can infer risk-neutral probabilities directly from asset prices

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

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

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

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

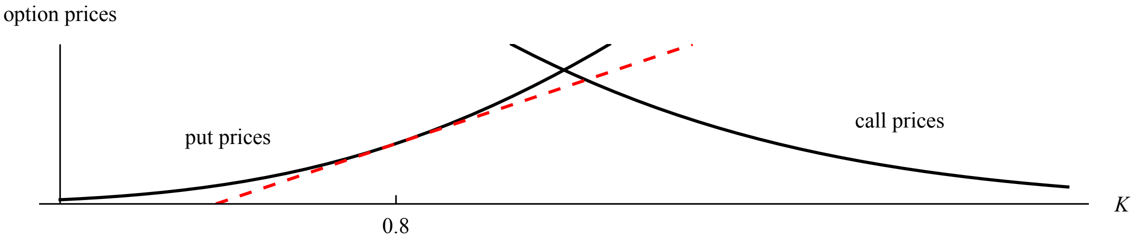

We can infer risk-neutral probabilities directly from asset prices

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

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

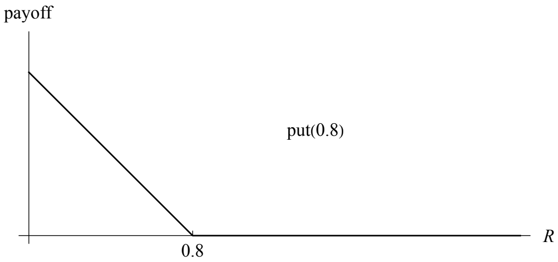

We can infer risk-neutral probabilities directly from asset prices

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

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

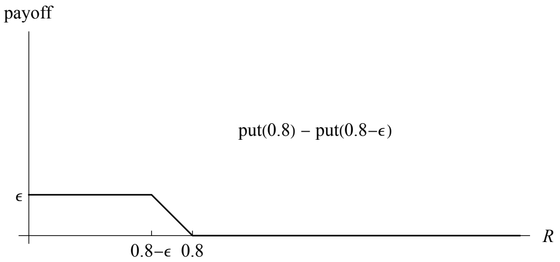

We can infer risk-neutral probabilities directly from asset prices

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

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

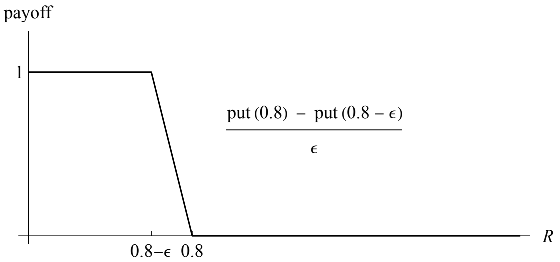

We can infer risk-neutral probabilities directly from asset prices

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

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

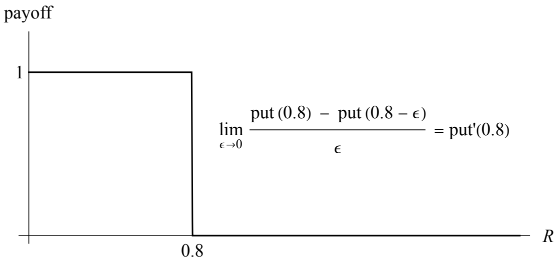

$$\lim_{\epsilon \to 0} \frac{\text{put}\,(0.8) - \text{put}\,(0.8 - \epsilon)}{\epsilon} = \text{put}'(0.8)$$

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

- Require an assumption to link the true and risk-neutral probabilities
  - that is, about the stochastic discount factor
- Take the perspective of a one-period marginal investor with power utility who chooses to hold the market. So the SDF must be $M = R_m^{-\gamma}/\lambda$
- 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 $X = \boldsymbol{I}(R_i \leq q)$, this implies that the crash probability of stock i is

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

  - 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

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

- To calculate $\mathbb{E}^* \left[R_m^\gamma\right]$, we need marginal risk-neutral distribution of R_m
  - Easy, using index option prices (Breeden and Litzenberger, 1978)
- To calculate $\mathbb{E}^* \left[R_m^\gamma \, \boldsymbol{I}(R_i \leq q)\right]$, we need the joint distribution of (R_m, R_i)
  - Problem: Joint risk-neutral distribution is not observable (from traded assets)
  - A general theme: we are often interested in covariances in financial economics
  - 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%
- 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

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

- We do not observe the joint (risk-neutral) distribution, so cannot calculate numerator
- 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

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

- The three (risk-neutral) expectations can be evaluated using index options
- The role of individual stock options?

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

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

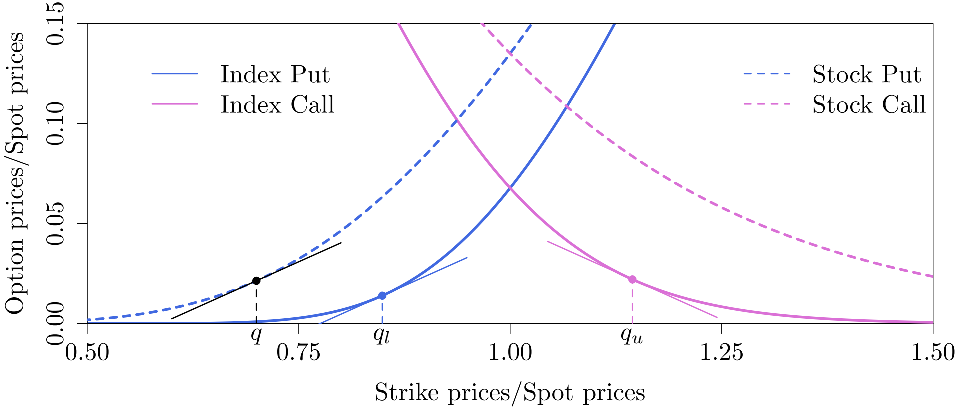

The upper and lower bounds are attainable in principle

- 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)

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

Further theoretical results

- Both $\mathbb{P}[R_i \leq q]$ and $\mathbb{P}^*[R_i \leq q]$ lie in between the bounds
- $\gamma = 0$: the lower and upper bounds both equal $\mathbb{P}^*[R_i \leq q]$, $\mathbb{P}^*$ and $\mathbb{P}$ coincide
- As γ increases, the interval widens monotonically
- As $\gamma \to \infty$, trivial: the lower bound $\to 0$ and the upper bound $\to 1$

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
- Firm characteristics from **Compustat**
- Price, return, and volume data from **CRSP**
- Focus on "crashes" of $10\%, 20\%$ and 30% at $\tau = 1, 3, 6$ 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

In-sample tests

Empirical tests

- $\boldsymbol{I}(R_i \leq q) = 0 + 1 \times \underbrace{\mathbb{E}[\ \boldsymbol{I}(R_i \leq q)\ ]}_{\mathbb{P}[R_i \leq q]} + \varepsilon$
- So a regression of the realized crash indicator $\boldsymbol{I}(R_i \leq q)$ 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

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

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% ($q = 0.7$)

	lower bound				risk-neutral				upper bound			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
α	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)
	[0.00]	[0.00]	[0.01]	[0.01]	[0.00]	[0.00]	[0.01]	[0.01]	[0.00]	[0.00]	[0.01]	[0.01]
β	0.95	1.03	1.09	1.05	0.66	0.60	0.59	0.56	0.51	0.43	0.39	0.35
	(0.15)	(0.12)	(0.11)	(0.10)	(0.11)	(0.08)	(0.07)	(0.07)	(0.09)	(0.06)	(0.05)	(0.05)
	[0.16]	[0.14]	[0.18]	[0.15]	[0.11]	[0.11]	[0.11]	[0.11]	[0.10]	[0.09]	[0.08]	[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% ($q = 0.7$)

	lower bound				risk-neutral				upper bound			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
β	0.93	1.05	1.11	1.14	0.68	0.70	0.74	0.78	0.55	0.55	0.58	0.60
	(0.14)	(0.10)	(0.08)	(0.08)	(0.10)	(0.07)	(0.05)	(0.05)	(0.09)	(0.05)	(0.04)	(0.04)
	[0.16]	[0.13]	[0.12]	[0.11]	[0.13]	[0.09]	[0.10]	[0.07]	[0.09]	[0.07]	[0.06]	[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% ($q = 0.8$)

	lower bound				risk-neutral				upper bound			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
α	0.00	−0.01	−0.01	0.02	0.00	−0.01	−0.02	0.00	0.00	−0.01	−0.01	0.01
	(0.00)	(0.01)	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.02)
	[0.00]	[0.01]	[0.01]	[0.02]	[0.00]	[0.01]	[0.01]	[0.02]	[0.00]	[0.01]	[0.02]	[0.03]
β	0.92	1.03	1.15	1.07	0.68	0.69	0.73	0.66	0.56	0.51	0.49	0.41
	(0.11)	(0.09)	(0.09)	(0.08)	(0.09)	(0.07)	(0.07)	(0.07)	(0.08)	(0.06)	(0.06)	(0.06)
	[0.11]	[0.13]	[0.15]	[0.13]	[0.09]	[0.10]	[0.11]	[0.12]	[0.07]	[0.08]	[0.10]	[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% ($q = 0.8$)

	lower bound				risk-neutral				upper bound			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
β	0.93	1.03	1.13	1.10	0.73	0.80	0.89	0.87	0.62	0.67	0.74	0.71
	(0.09)	(0.07)	(0.06)	(0.06)	(0.07)	(0.05)	(0.05)	(0.05)	(0.06)	(0.04)	(0.04)	(0.04)
	[0.10]	[0.10]	[0.09]	[0.09]	[0.07]	[0.07]	[0.07]	[0.06]	[0.07]	[0.07]	[0.07]	[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%

	lower bound				risk-neutral				upper bound			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
α	0.00	0.01	0.09	0.34	0.00	0.00	0.04	0.24	0.00	−0.01	0.03	0.21
	(0.00)	(0.00)	(0.01)	(0.02)	(0.00)	(0.01)	(0.01)	(0.03)	(0.00)	(0.01)	(0.01)	(0.03)
	[0.00]	[0.01]	[0.01]	[0.03]	[0.00]	[0.01]	[0.02]	[0.04]	[0.00]	[0.01]	[0.02]	[0.04]
β	1.35	1.58	1.32	0.12	1.03	1.17	1.08	0.46	0.85	0.91	0.82	0.42
	(0.13)	(0.11)	(0.11)	(0.15)	(0.10)	(0.09)	(0.09)	(0.12)	(0.09)	(0.08)	(0.07)	(0.09)
	[0.13]	[0.14]	[0.15]	[0.21]	[0.11]	[0.13]	[0.15]	[0.17]	[0.09]	[0.10]	[0.11]	[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 rallies than crashes (perhaps because rallies are more idiosyncratic so comonotonicity is less likely to hold)

In-sample tests (3)

Down by 10% ($q = 0.9$)

	lower bound				risk-neutral				upper bound			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
α	−0.02	−0.01	−0.01	0.03	−0.02	−0.02	−0.02	0.00	−0.02	0.00	0.01	0.05
	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.02)	(0.02)	(0.03)	(0.01)	(0.02)	(0.02)	(0.03)
	[0.01]	[0.02]	[0.02]	[0.03]	[0.01]	[0.02]	[0.03]	[0.04]	[0.01]	[0.03]	[0.04]	[0.05]
β	1.05	1.07	1.12	1.01	0.88	0.83	0.80	0.68	0.75	0.63	0.54	0.41
	(0.08)	(0.07)	(0.07)	(0.08)	(0.08)	(0.08)	(0.08)	(0.09)	(0.07)	(0.07)	(0.07)	(0.08)
	[0.08]	[0.11]	[0.12]	[0.12]	[0.07]	[0.11]	[0.12]	[0.13]	[0.08]	[0.12]	[0.12]	[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% ($q = 0.9$)

	lower bound				risk-neutral				upper bound			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
β	0.99	0.99	1.05	1.05	0.88	0.89	0.94	0.93	0.80	0.79	0.83	0.82
	(0.06)	(0.05)	(0.06)	(0.06)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.04)	(0.04)	(0.05)
	[0.06]	[0.07]	[0.08]	[0.08]	[0.05]	[0.07]	[0.07]	[0.08]	[0.05]	[0.06]	[0.06]	[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 slope β, by year: lower bound

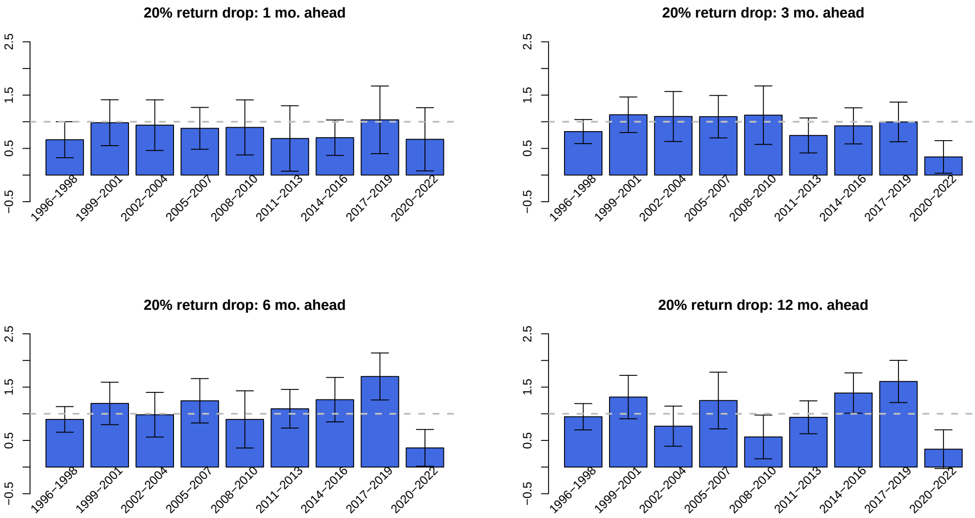

Estimated slope β, by year: risk-neutral probabilities

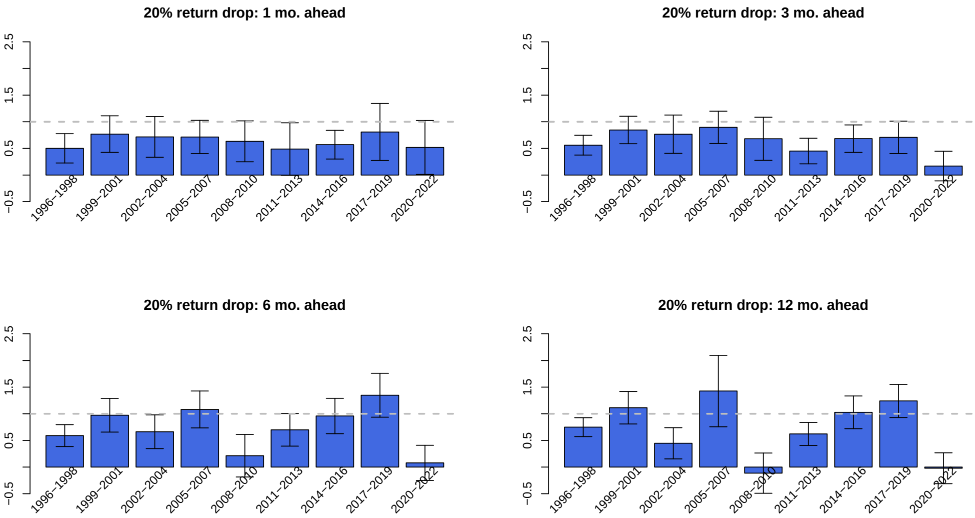

Estimated slope β, by industry: lower bound

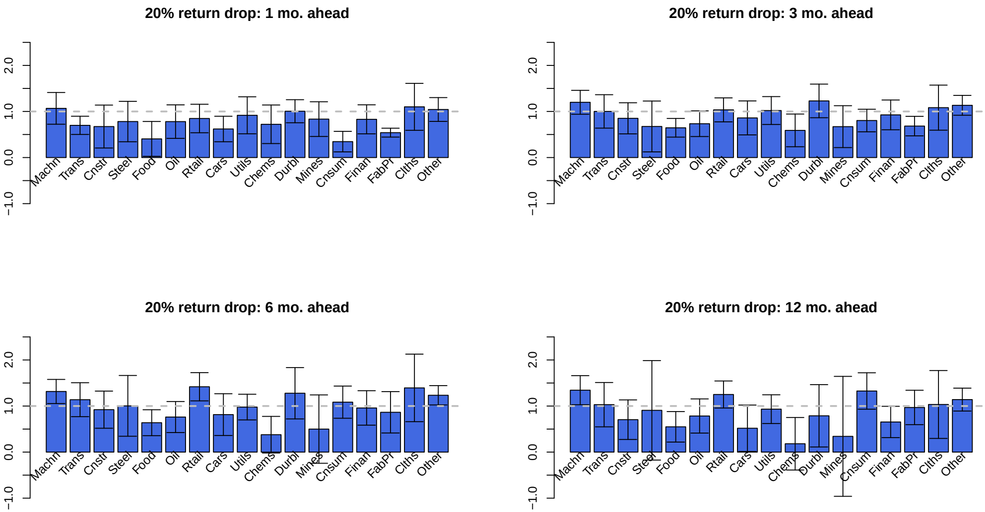

Competitor variables from the literature

- We compare against 15 variables drawn from the literature
  - Stock characteristics: CAPM β, (log) relative size, book-to-market, gross profitability, momentum (prior 2-6 and 2-12 month returns), lagged return
  - Chen–Hong–Stein, 2001: realized volatilities and monthly turnover
  - 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, earnings, cash, log price per share (winsorized from above at $15)
- All variables are standardized to unit standard deviation for comparability

In-sample tests (4)

	$\boldsymbol{I}(R_{t\to t+1} \leq 0.8)$					
$\mathbb{P}^L[R_{t\to t+1} \leq 0.8]$		3.40*	3.02*		4.41	2.72*
		(0.41)	(0.58)		(3.08)	(0.33)
$\mathbb{P}^*[R_{t\to t+1} \leq 0.8]$				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)

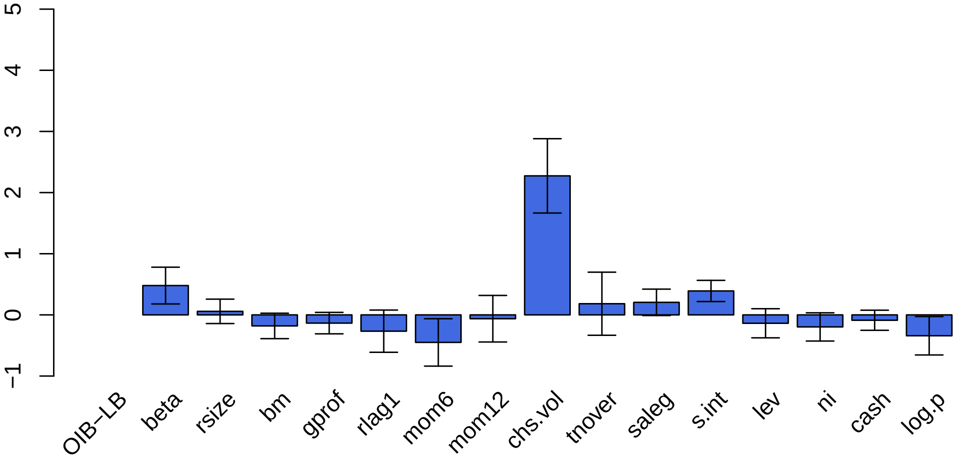

In-sample tests (4)

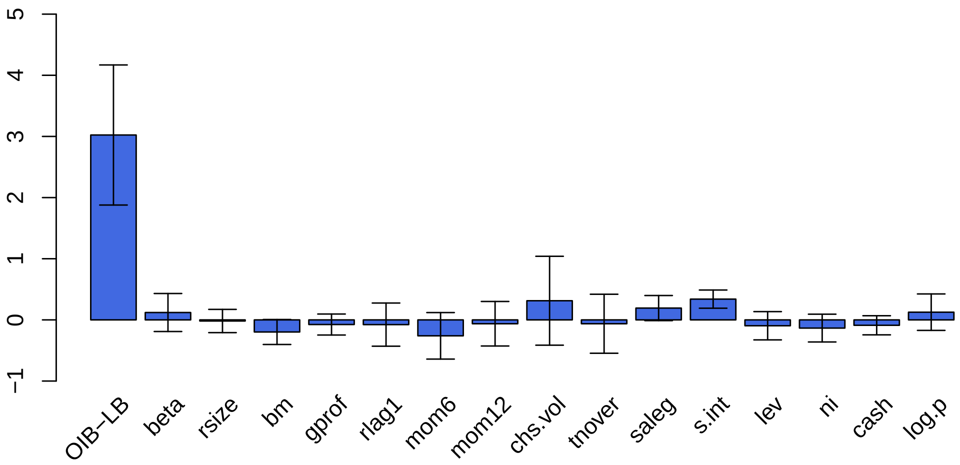

Out-of-sample tests

We compare OOS forecast performance of two models

1. Competitor model uses 15 char. + risk-neutral + lower bound
   - 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

Out-of-sample forecasts

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

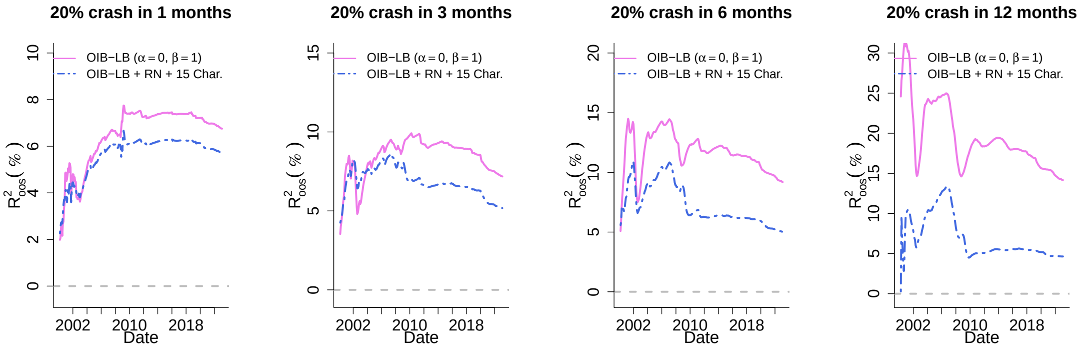

Out-of-sample forecasts

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

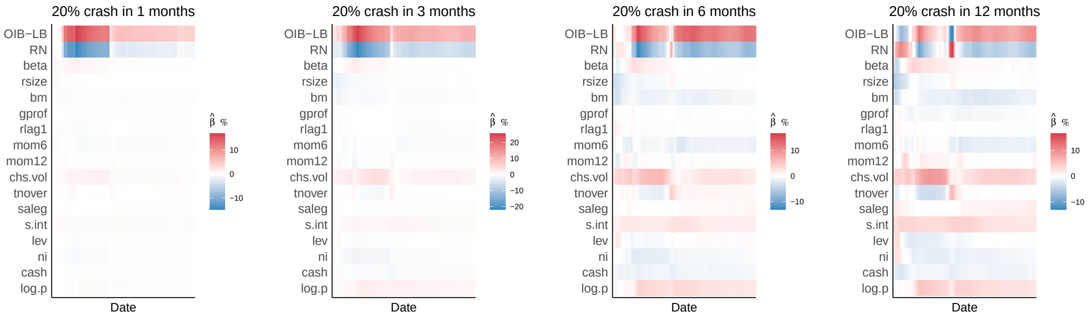

Out-of-sample forecasts

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

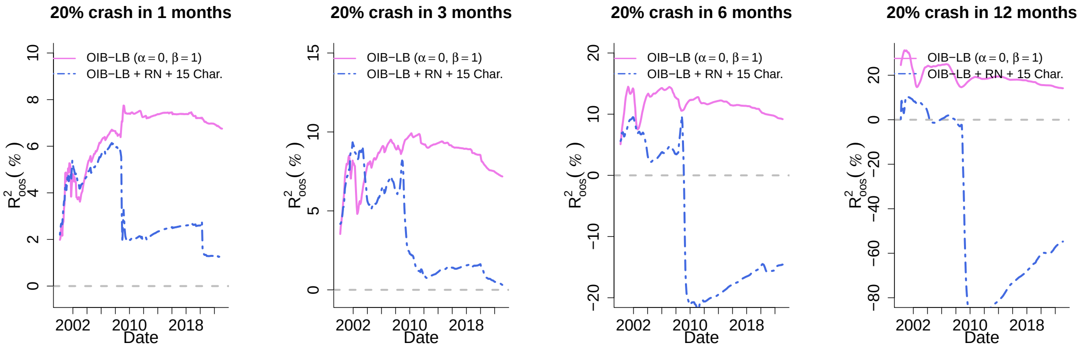

Out-of-sample forecasts

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

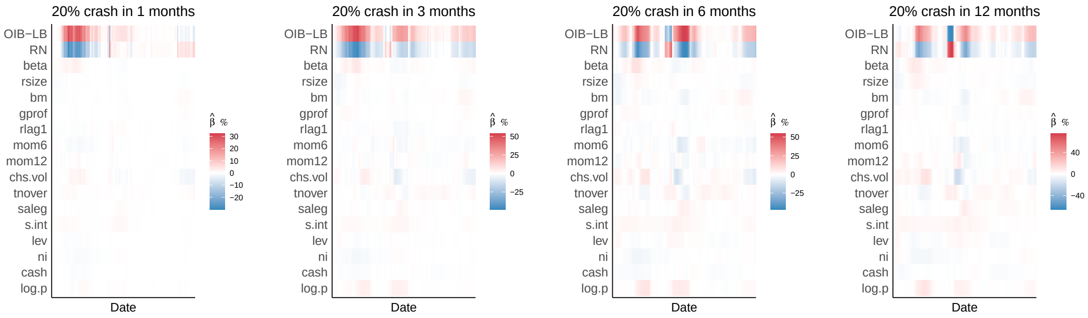

Industry crash risk

Industry average crash probabilities

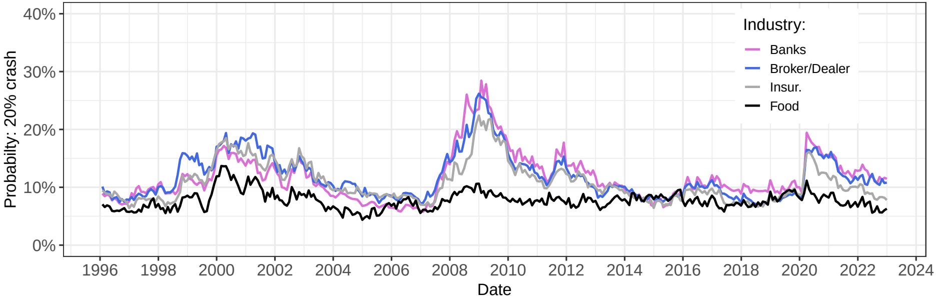

- Substantial variation in crash probability over time and across industries
- News about crash risk is not just idiosyncratic: related industries' probabilities comove

Industry average crash probabilities

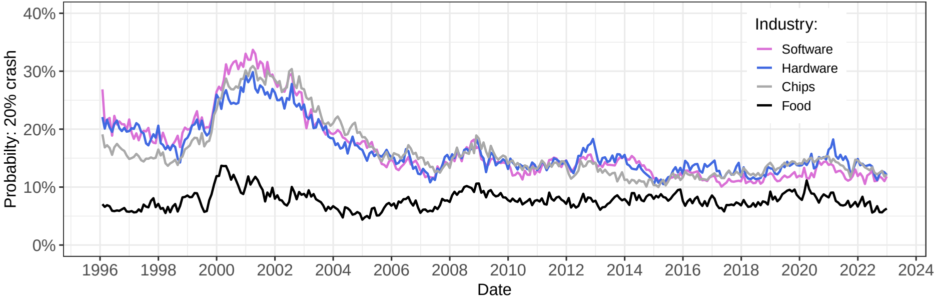

- Substantial variation in crash probability over time and across industries
- News about crash risk is not just idiosyncratic: related industries' probabilities comove

Industry average crash probabilities

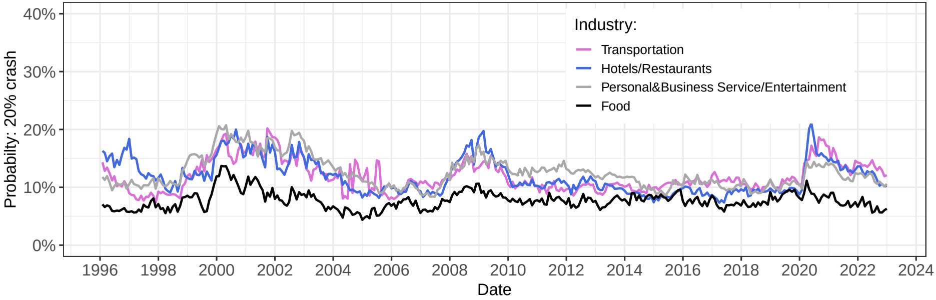

- Substantial variation in crash probability over time and across industries
- News about crash risk is not just idiosyncratic: related industries' probabilities comove

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 CHS volatility (Chen, Hong and Stein, 2001) and penny stocks (Campbell, Hilscher and Szilagyi, 2008)
  - for certain industries: high beta, share turnover, short interest (Hong and Stein, 2003); poor recent returns, profit, and earnings
- Realized crash event regressions cannot reveal these patterns

Regressions of the lower bound onto 15 characteristics

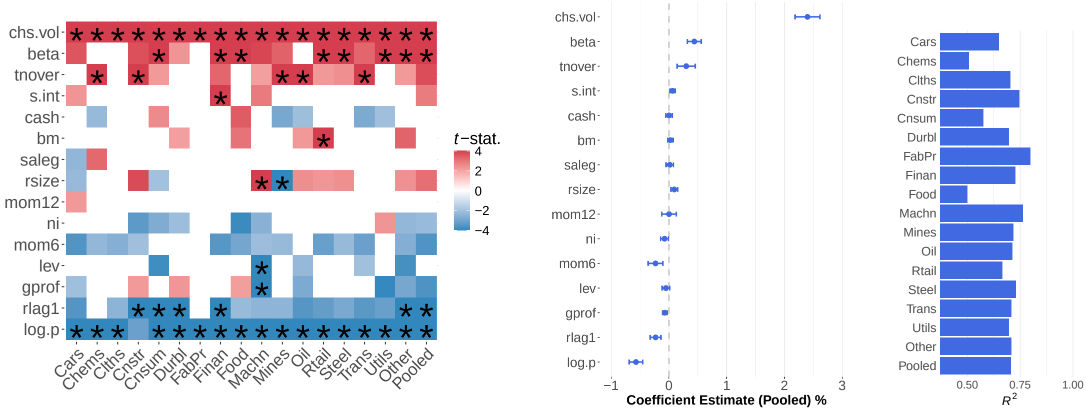

Summary

- 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
- 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 key assumption: the form of the SDF
    - it allows to avoid the costly (and commonly made) assumptions that trailing estimates are good proxies for the forward-looking measures backed by theory
- It seems the price of our assumption is worth paying