Forecasting Crashes with a Smile

Ian Martin   Ran Shi

April, 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
- The bounds are observable in real time
- They perform well in and out of sample
- No distributional assumptions

Probabilities of a 20% decline over the next month

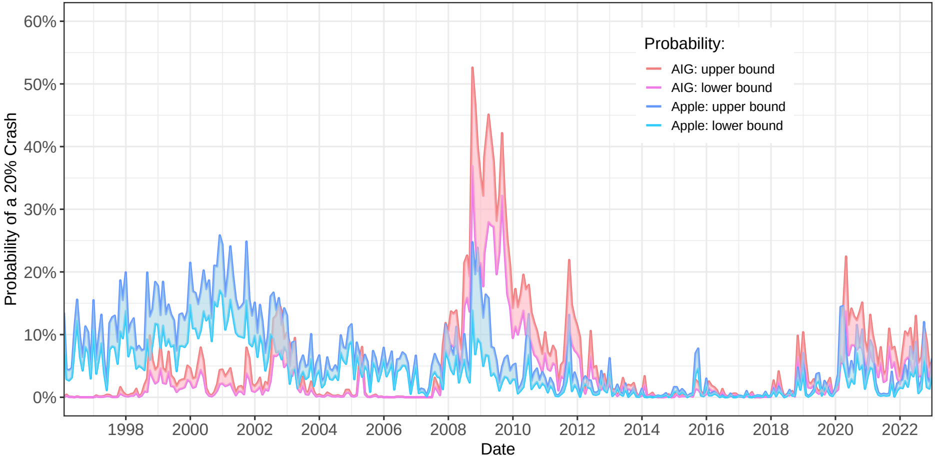

Probabilities of a 20% decline over the next year

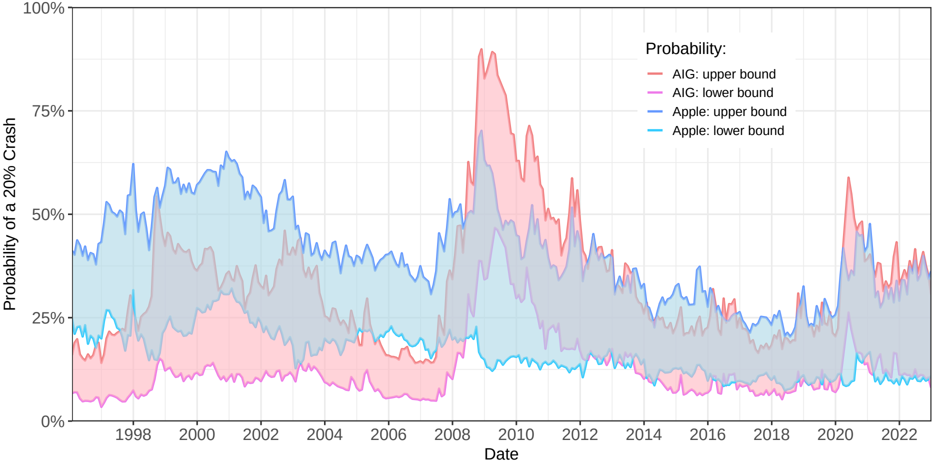

Probabilities of a 20% decline over the next month

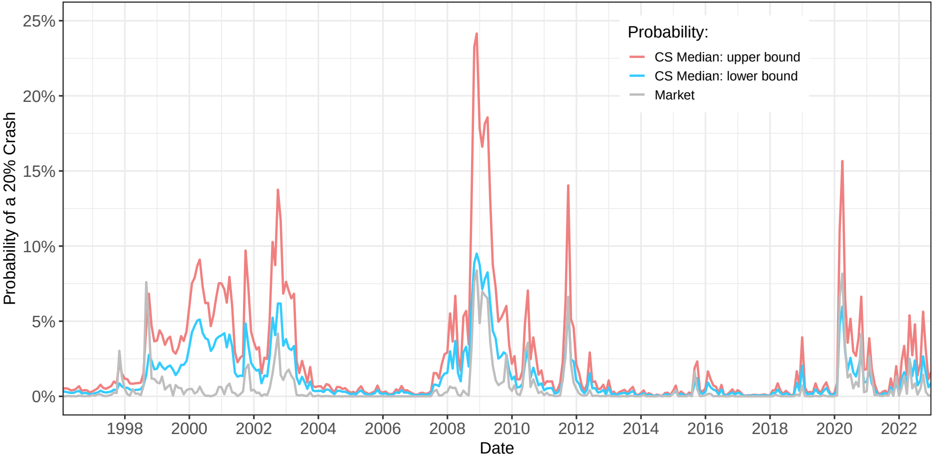

We can infer risk-neutral probabilities directly from asset prices

Elections | $102,532,490 Bet | Nov 5, 2024

Presidential Election Winner 2024

Polymarket

OUTCOME	% CHANCE		
Donald Trump $8,720,022 Bet	48%	Bet Yes 49¢	Bet No 52¢
Joe Biden $9,129,749 Bet	41%	Bet Yes 42¢	Bet No 59¢
Robert F. Kennedy Jr. $6,251,780 Bet	5%	Bet Yes 5.2¢	Bet No 95.5¢
Michelle Obama $7,033,244 Bet	2%	Bet Yes 2.1¢	Bet No 98.0¢
Kamala Harris $3,573,128 Bet	2%	Bet Yes 1.8¢	Bet No 98.5¢

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

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

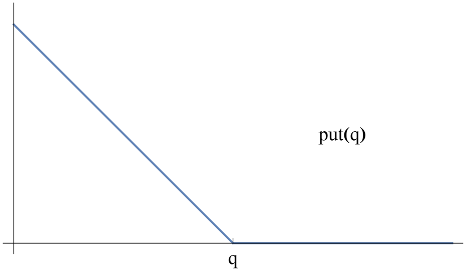

We can infer risk-neutral probabilities directly from asset prices

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

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

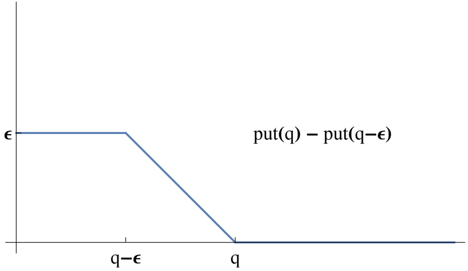

We can infer risk-neutral probabilities directly from asset prices

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

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

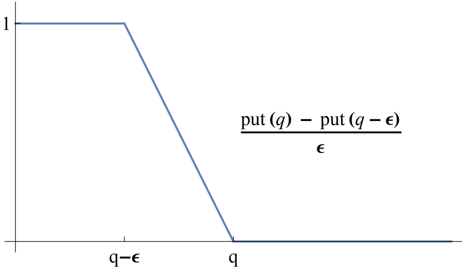

We can infer risk-neutral probabilities directly from asset prices

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

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

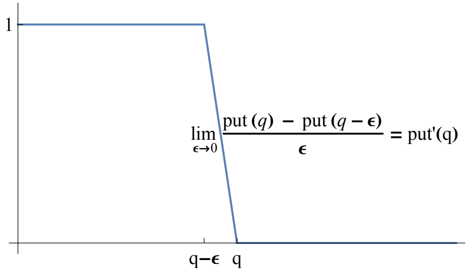

But we want true, not risk-neutral, probabilities

- We require an assumption (implicit or explicit) to link the true and risk-neutral probabilities—that is, about the stochastic discount factor
- Simple example: think from the perspective of a marginal investor with log utility who chooses to invest fully in the market
  - From this investor's perspective, $1/R_m$ must be a stochastic discount factor
  - This implies that
$$\mathbb{E}X = \mathbb{E}\left[\left(\frac{1}{R_m}\right) XR_m\right] = \frac{1}{R_f}\,\mathbb{E}^*\left(XR_m\right)$$
  - We can infer beliefs about X from the *price* of XR_m

Theory (1)

- We take the perspective of a one-period marginal investor with power utility who chooses to hold the market
- It follows that the SDF is of the form $M = R_m^{-\gamma}/\lambda$ for some constant λ
- γ is the coefficient of relative risk aversion

Theory (2)

- The true expectation of a random payoff X then satisfies

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

- Applied in the case $X = 1$, we must have $\lambda = \frac{R_f}{\mathbb{E}^*[R_m^\gamma]}$
- So,

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

Theory (3)

- Consider $X = \boldsymbol{I}(R_i \leq q)$, the crash probability of a stock

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

- To calculate $\mathbb{E}^* \left[R_m^\gamma\right]$, we need marginal 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 given assets that are traded in practice (Martin, 2018)
  - 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 (4)

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

- We do not observe the joint risk-neutral distribution, so cannot calculate 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 (5)

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 elements are

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

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

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

Theory (6)

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

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

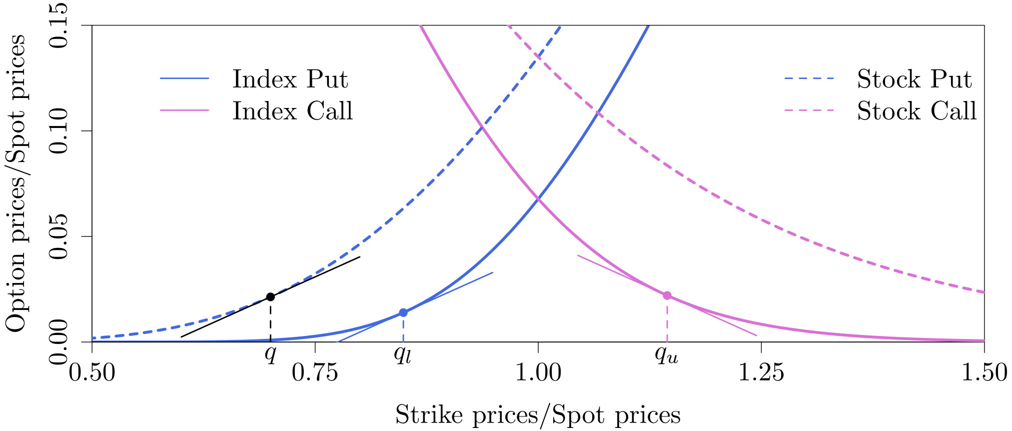

Theory (7)

- 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

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

Further theoretical results

- Both $\mathbb{P}[R_i \leq q]$ and $\mathbb{P}^*[R_i \leq q]$ lie in between the bounds
- 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 $\gamma \to \infty$, 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**
- Option 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
  - 1044 firms in total
  - On average around 492 firms each month
  - Over $155,000$ firm-month observations per maturity
- Risk-free rates from **OptionMetrics**
- Firm characteristics from **Compustat**
- Price, return, and volume data from **CRSP**

Empirical setup

- We set γ equal to 2
- Focus on "crashes" of $5\%, 10\%$ and 20%: $q = 0.95, 0.9$ and 0.8
- For stock i in month t, compute upper and lower bounds at horizons $\tau = 1, 3, 6$ and 12 months, $\mathbb{P}^U_{i,t}(\tau, q)$ and $\mathbb{P}^L_{i,t}(\tau, q)$
- I'll emphasize the case of 20% decline over one month

Calibrating risk aversion

We set $\gamma = 2$

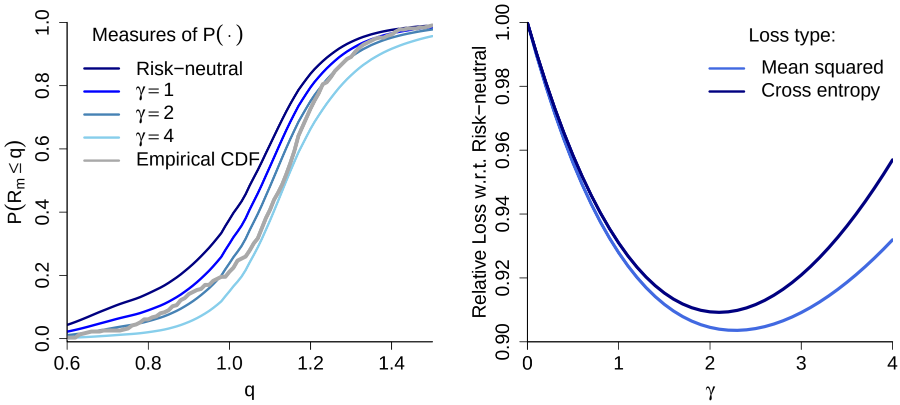

- Left: unconditional CDF of market return for various γ; and the empirical distribution
- Right: Loss functions when different values of γ are used to forecast market crashes

Summary statistics

		averaged across firms (number of obs. $T = 324$)				averaged across time (number of obs. $N = 1044$)			
	maturity	1	3	6	12	1	3	6	12
$q = 0.80$, down by over 20%									
realized	mean	0.021	0.069	0.111	0.152	0.029	0.084	0.130	0.173
	s.d.	0.048	0.107	0.141	0.160	0.059	0.092	0.129	0.166
lower bound	mean	0.022	0.073	0.102	0.123	0.027	0.079	0.110	0.133
	s.d.	0.020	0.029	0.028	0.027	0.029	0.046	0.052	0.056
upper bound	mean	0.038	0.144	0.233	0.339	0.044	0.152	0.242	0.350
	s.d.	0.040	0.071	0.082	0.098	0.042	0.069	0.079	0.089
risk-neutral	mean	0.031	0.113	0.173	0.236	0.037	0.120	0.181	0.245
	s.d.	0.031	0.050	0.053	0.059	0.036	0.058	0.065	0.072

Summary statistics

		averaged across firms (number of obs. $T = 324$)				averaged across time (number of obs. $N = 1044$)			
	maturity	1	3	6	12	1	3	6	12
$q = 0.90$, down by over 10%									
realized	mean	0.096	0.173	0.211	0.236	0.110	0.191	0.231	0.252
	s.d.	0.124	0.170	0.185	0.196	0.089	0.119	0.152	0.183
lower bound	mean	0.109	0.168	0.196	0.210	0.118	0.179	0.206	0.219
	s.d.	0.037	0.031	0.028	0.023	0.050	0.055	0.056	0.056
upper bound	mean	0.156	0.277	0.366	0.466	0.166	0.289	0.378	0.475
	s.d.	0.064	0.074	0.081	0.087	0.062	0.070	0.074	0.074
risk-neutral	mean	0.136	0.228	0.286	0.341	0.145	0.239	0.297	0.350
	s.d.	0.051	0.051	0.051	0.050	0.056	0.061	0.063	0.063

Summary statistics

		averaged across firms (number of obs. $T = 324$)				averaged across time (number of obs. $N = 1044$)			
	maturity	1	3	6	12	1	3	6	12
$q = 0.95$, down by over 5%									
realized	mean	0.216	0.271	0.288	0.289	0.230	0.287	0.306	0.306
	s.d.	0.187	0.200	0.204	0.210	0.101	0.122	0.155	0.185
lower bound	mean	0.215	0.264	0.277	0.271	0.228	0.275	0.286	0.279
	s.d.	0.036	0.024	0.020	0.020	0.052	0.049	0.047	0.048
upper bound	mean	0.281	0.393	0.465	0.541	0.294	0.404	0.474	0.548
	s.d.	0.064	0.064	0.066	0.074	0.059	0.058	0.056	0.057
risk-neutral	mean	0.251	0.332	0.375	0.408	0.264	0.343	0.383	0.415
	s.d.	0.049	0.041	0.038	0.040	0.055	0.052	0.049	0.050

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 threshold q

In-sample tests (1)

Down by 20% ($q = 0.80$)

	lower bound				upper bound				risk neutral			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
α	0.00	−0.01	−0.01	0.02	0.00	0.00	0.00	0.01	0.00	−0.01	−0.02	0.00
	(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]	[0.01]	[0.00]	[0.01]	[0.02]	[0.03]	[0.00]	[0.01]	[0.01]	[0.03]
β	0.92	1.03	1.15	1.08	0.55	0.51	0.50	0.41	0.68	0.69	0.73	0.66
	(0.11)	(0.09)	(0.09)	(0.08)	(0.08)	(0.06)	(0.06)	(0.06)	(0.09)	(0.07)	(0.07)	(0.08)
	[0.11]	[0.15]	[0.14]	[0.12]	[0.08]	[0.08]	[0.09]	[0.10]	[0.08]	[0.10]	[0.12]	[0.13]
R^2	5.66%	5.17%	4.78%	3.76%	5.29%	4.13%	3.26%	2.33%	5.45%	4.51%	3.91%	3.00%

In-sample tests (1)

with time fixed effects

Down by 20% ($q = 0.80$)

	lower bound				upper bound				risk neutral			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
β	0.93	1.04	1.13	1.11	0.62	0.68	0.74	0.72	0.73	0.81	0.89	0.87
	(0.09)	(0.07)	(0.06)	(0.06)	(0.06)	(0.04)	(0.04)	(0.04)	(0.07)	(0.05)	(0.05)	(0.05)
	[0.10]	[0.10]	[0.10]	[0.08]	[0.06]	[0.09]	[0.04]	[0.06]	[0.08]	[0.05]	[0.05]	[0.06]
R^2-proj	4.45%	4.66%	4.56%	4.11%	4.29%	4.46%	4.41%	4.08%	4.35%	4.54%	4.49%	4.10%

Intermission: Probability of a rise of at least 20%

	lower bound				upper bound				risk neutral			
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	0.03	0.20	0.00	0.00	0.04	0.23
	(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]	[0.03]
β	1.35	1.58	1.30	0.10	0.85	0.91	0.82	0.44	1.03	1.17	1.08	0.49
	(0.13)	(0.11)	(0.11)	(0.14)	(0.09)	(0.08)	(0.08)	(0.09)	(0.11)	(0.09)	(0.09)	(0.12)
	[0.13]	[0.16]	[0.19]	[0.19]	[0.10]	[0.11]	[0.10]	[0.13]	[0.11]	[0.12]	[0.14]	[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: sign flips because $\boldsymbol{I}(R_i \geq q)$ is an increasing function of R_i, whereas $\boldsymbol{I}(R_i \leq q)$ is decreasing
- 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 (2)

Down by 10% ($q = 0.90$)

	lower bound				upper bound				risk neutral			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
α	−0.02	−0.01	−0.01	0.02	−0.02	0.00	0.01	0.05	−0.02	−0.02	−0.02	0.00
	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.02)	(0.02)	(0.04)	(0.01)	(0.02)	(0.02)	(0.03)
	[0.01]	[0.01]	[0.02]	[0.03]	[0.01]	[0.03]	[0.03]	[0.06]	[0.01]	[0.02]	[0.03]	[0.04]
β	1.05	1.07	1.12	1.03	0.75	0.63	0.54	0.41	0.88	0.83	0.81	0.69
	(0.08)	(0.07)	(0.07)	(0.08)	(0.07)	(0.07)	(0.07)	(0.08)	(0.08)	(0.08)	(0.08)	(0.09)
	[0.08]	[0.10]	[0.11]	[0.13]	[0.07]	[0.11]	[0.10]	[0.13]	[0.08]	[0.11]	[0.12]	[0.14]
R^2	5.47%	3.73%	3.43%	2.54%	5.36%	3.06%	2.20%	1.25%	5.47%	3.41%	2.84%	1.88%

In-sample tests (2)

with time fixed effects

Down by 10% ($q = 0.90$)

	lower bound				upper bound				risk neutral			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
β	0.99	1.00	1.06	1.06	0.81	0.79	0.83	0.83	0.89	0.89	0.95	0.95
	(0.06)	(0.05)	(0.06)	(0.06)	(0.05)	(0.04)	(0.04)	(0.05)	(0.05)	(0.05)	(0.05)	(0.06)
	[0.07]	[0.08]	[0.06]	[0.06]	[0.05]	[0.05]	[0.06]	[0.07]	[0.06]	[0.05]	[0.08]	[0.09]
R^2-proj	4.02%	3.17%	3.18%	2.97%	3.96%	3.11%	3.14%	2.96%	3.98%	3.14%	3.17%	2.95%

In-sample tests (3)

Down by 5% ($q = 0.95$)

	lower bound				upper bound				risk neutral			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
α	0.00	0.02	−0.01	0.05	0.00	0.05	0.06	0.11	0.00	0.01	−0.01	0.04
	(0.01)	(0.02)	(0.02)	(0.03)	(0.02)	(0.03)	(0.04)	(0.05)	(0.02)	(0.03)	(0.03)	(0.05)
	[0.01]	[0.03]	[0.03]	[0.06]	[0.02]	[0.05]	[0.06]	[0.09]	[0.02]	[0.04]	[0.05]	[0.07]
β	0.98	0.95	1.06	0.88	0.76	0.56	0.49	0.33	0.88	0.77	0.80	0.61
	(0.07)	(0.07)	(0.08)	(0.10)	(0.08)	(0.09)	(0.09)	(0.10)	(0.08)	(0.09)	(0.10)	(0.12)
	[0.06]	[0.10]	[0.11]	[0.18]	[0.08]	[0.13]	[0.14]	[0.17]	[0.08]	[0.12]	[0.14]	[0.17]
R^2	3.01%	1.85%	1.86%	1.36%	3.02%	1.35%	0.94%	0.49%	3.08%	1.64%	1.45%	0.93%

In-sample tests (3)

with time fixed effects

Down by 5% ($q = 0.95$)

	lower bound				upper bound				risk neutral			
maturity	1	3	6	12	1	3	6	12	1	3	6	12
β	0.87	0.86	0.97	0.97	0.77	0.76	0.85	0.86	0.83	0.82	0.93	0.93
	(0.05)	(0.05)	(0.07)	(0.07)	(0.04)	(0.05)	(0.06)	(0.06)	(0.05)	(0.05)	(0.06)	(0.07)
	[0.04]	[0.06]	[0.12]	[0.07]	[0.02]	[0.08]	[0.07]	[0.06]	[0.03]	[0.09]	[0.10]	[0.09]
R^2-proj	2.21%	1.62%	1.75%	1.84%	2.19%	1.60%	1.74%	1.85%	2.20%	1.61%	1.74%	1.84%

Estimated β, by year: lower bound

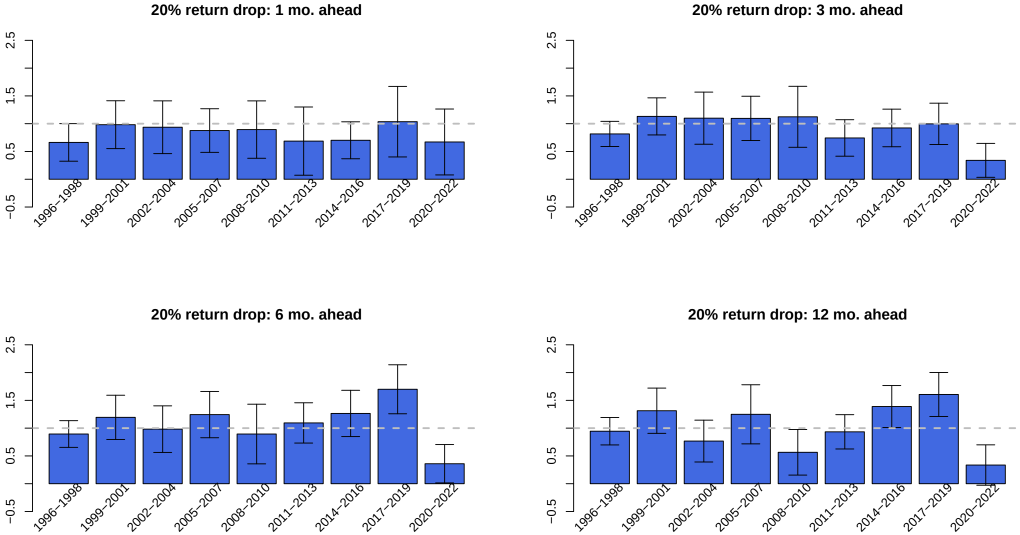

Estimated β, by year: risk-neutral probabilities

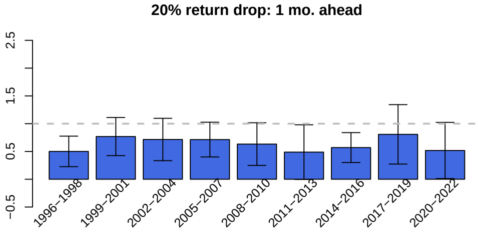

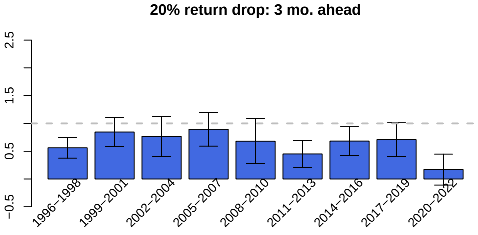

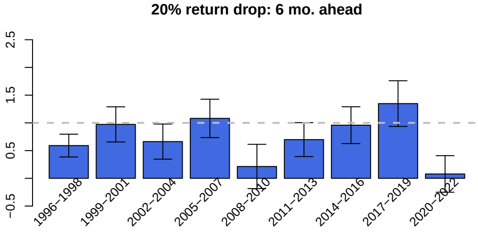

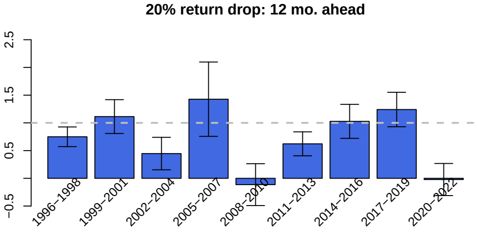

Lower bound vs. risk-neutral probabilities

- 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

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

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

$$\mathbb{P}^*\left[R_i \le q\right] - \frac{\sigma^*\left[R_m^\gamma\right]\sigma^*\left[\boldsymbol{I}(R_i \le q)\right]}{\mathbb{E}^*\left[R_m^\gamma\right]} \le \mathbb{P}\left[R_i \le q\right] \le \mathbb{P}^*\left[R_i \le q\right] + \frac{\sigma^*\left[R_m^\gamma\right]\sigma^*\left[\boldsymbol{I}(R_i \le q)\right]}{\mathbb{E}^*\left[R_m^\gamma\right]}$$

  where $\sigma^*\left[\cdot\right]$ 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	mo.	mean	sd	median	q25	q75	min	max
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
5%	1	0.615	0.083	0.630	0.578	0.668	0.000	0.849
5%	3	0.716	0.047	0.717	0.681	0.757	0.003	0.828
5%	6	0.751	0.033	0.761	0.735	0.775	0.053	0.812
5%	12	0.766	0.024	0.771	0.757	0.781	0.074	0.842

	20% crashes over one month		
Fréchet–Hoeffding	0.85	1.27	1.57
	(0.54)	(0.84)	(1.47)
Cauchy–Schwarz	0.09		−0.20
	(0.56)		(0.65)
risk-neutral		−0.27	−0.35
		(0.66)	(0.80)
constant	0.00	0.00	−0.00
	(0.00)	(0.00)	(0.00)

Estimated β, by industry: lower bound

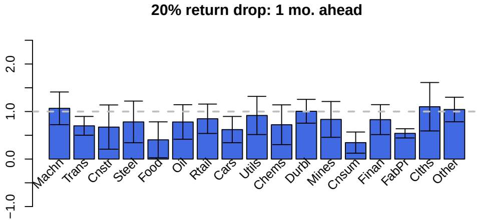

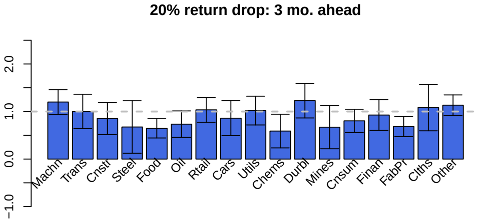

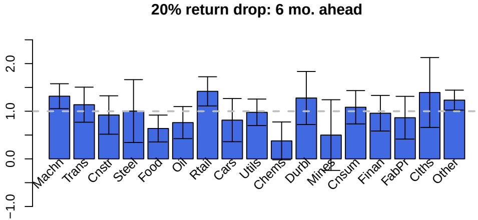

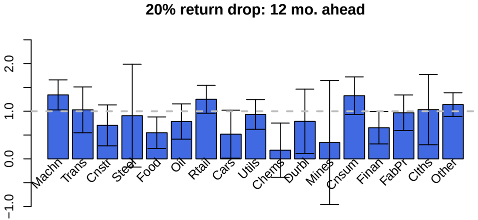

Competitor variables from the literature

- 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

			$\boldsymbol{I}(R_{t\to t+1} \leq 0.8)$			
$\mathbb{P}^L[R_{t\to t+1} \leq 0.8]$		3.41^*	3.05^*		4.44	2.74^*
		(0.41)	(0.59)		(3.08)	(0.33)
$\mathbb{P}^*[R_{t\to t+1} \leq 0.8]$				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)
	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
R^2/R^2-proj.	4.51%	5.66%	5.85%	5.72%	5.87%	4.74%

In-sample tests (4)

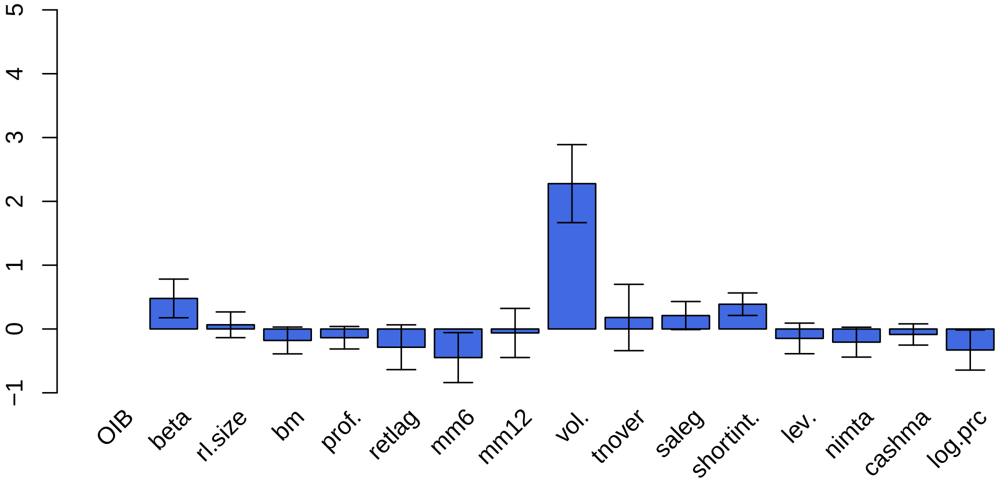

In-sample tests (4)

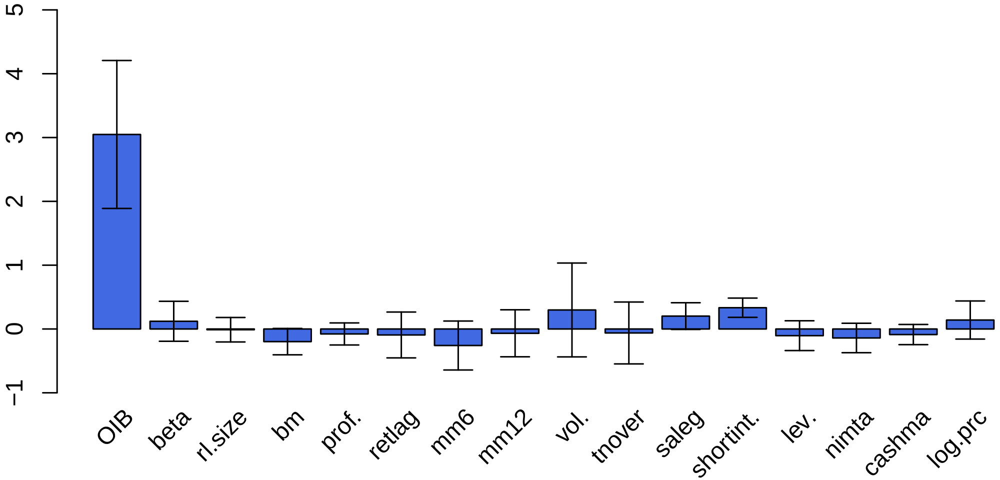

Back–Crotty–Kazempour (2022)

- 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

	lower bound				upper bound			
horizon	1	3	6	12	1	3	6	12
Panel A: $q = 0.80$, down by over 20%								
validity	0.452	0.381	0.621	0.487	0.769	1.000	0.754	1.000
tightness	0.352	0.022	0.043	0.164	0.011	0.000	0.000	0.018
Panel B: $q = 0.90$, down by over 10%								
validity	0.069	0.626	0.683	0.505	0.780	0.768	0.755	0.755
tightness	0.133	0.059	0.057	0.114	0.000	0.000	0.000	0.020
Panel C: $q = 0.95$, down by over 5%								
validity	0.552	0.629	0.563	0.486	1.000	0.779	0.760	1.000
tightness	0.176	0.043	0.048	0.096	0.001	0.000	0.000	0.019

Out-of-sample forecasts

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

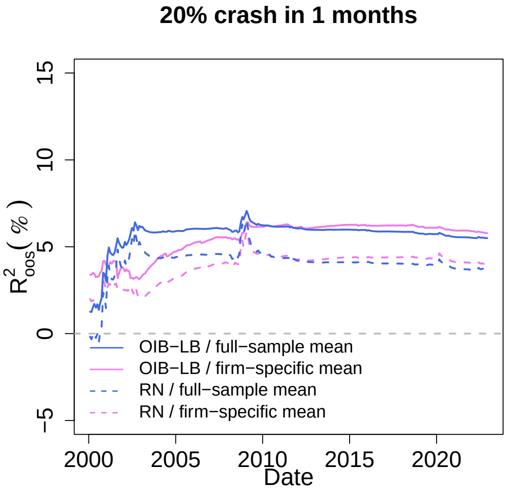

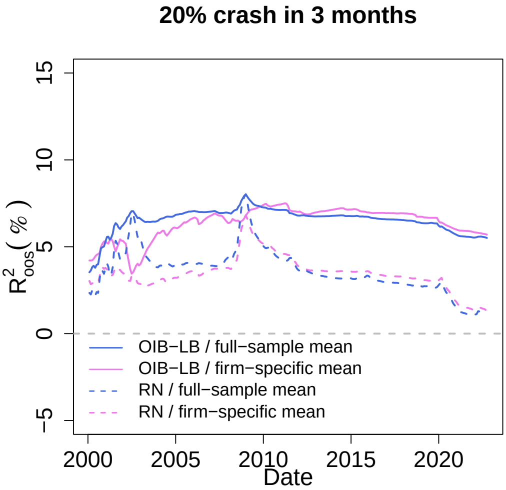

Out-of-sample forecasts

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

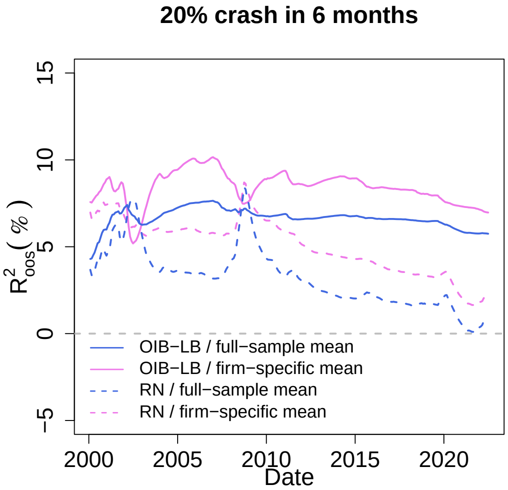

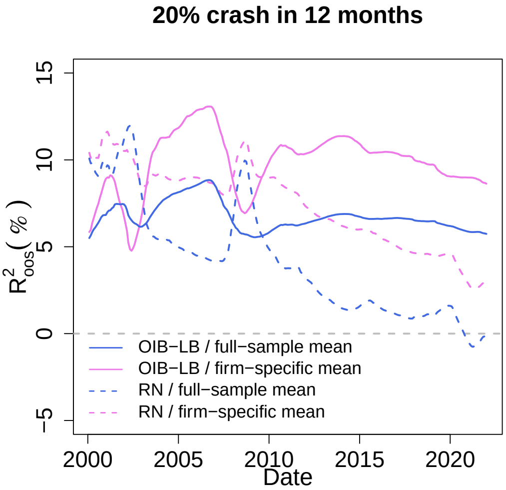

Out-of-sample forecasts

A more challenging competitor

- We include 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

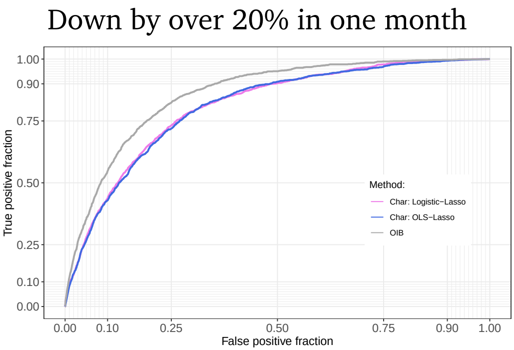

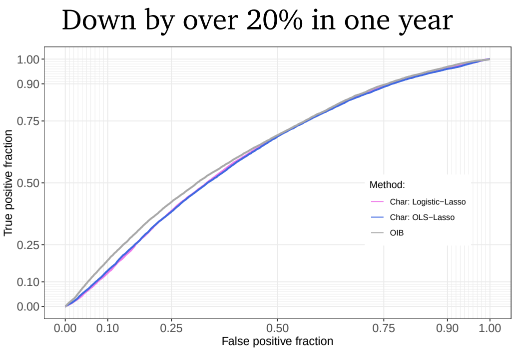

Out-of-sample forecasts: ROC curves

Training Sample 1996 – 2011

Down by over 20% in one month

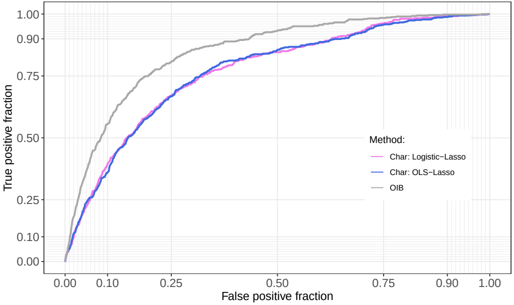

Down by over 20% in one year

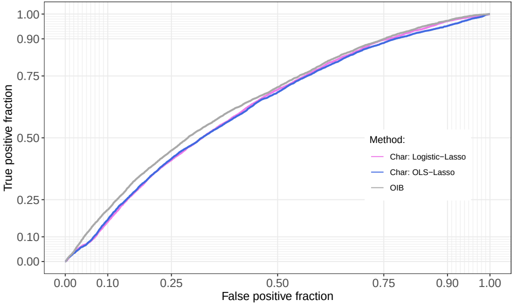

Out-of-sample forecasts: ROC curves

Training Sample 1996 – 2016

Down by over 20% in one month

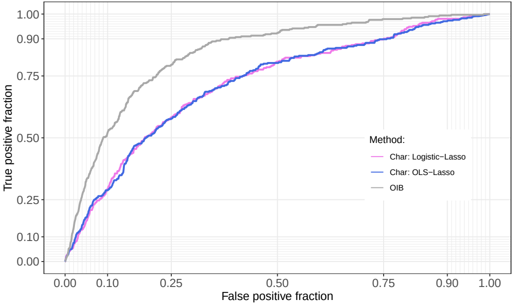

Down by over 20% in one year

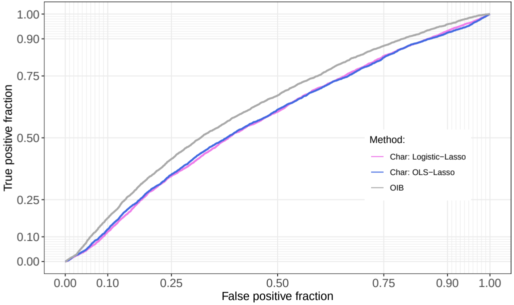

Industry average crash probabilities

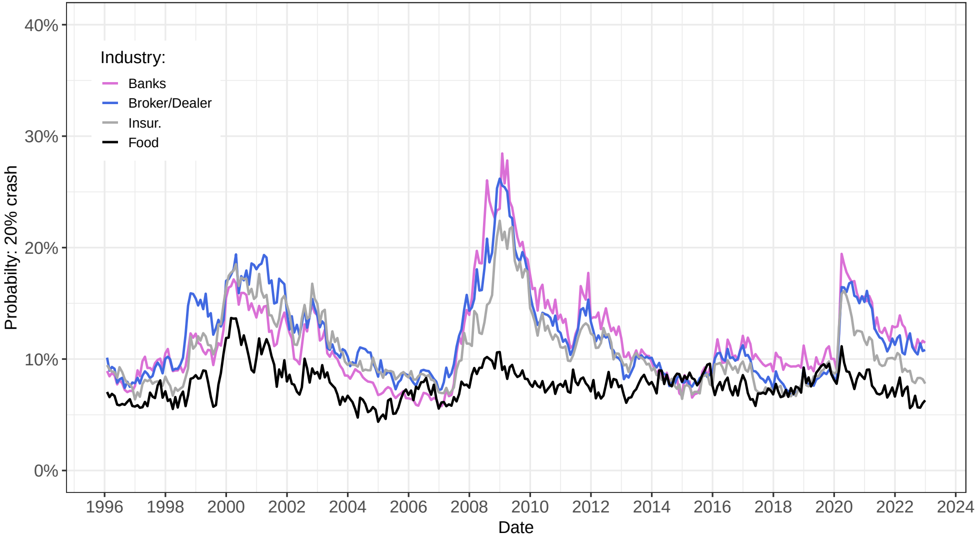

Industry average crash probabilities

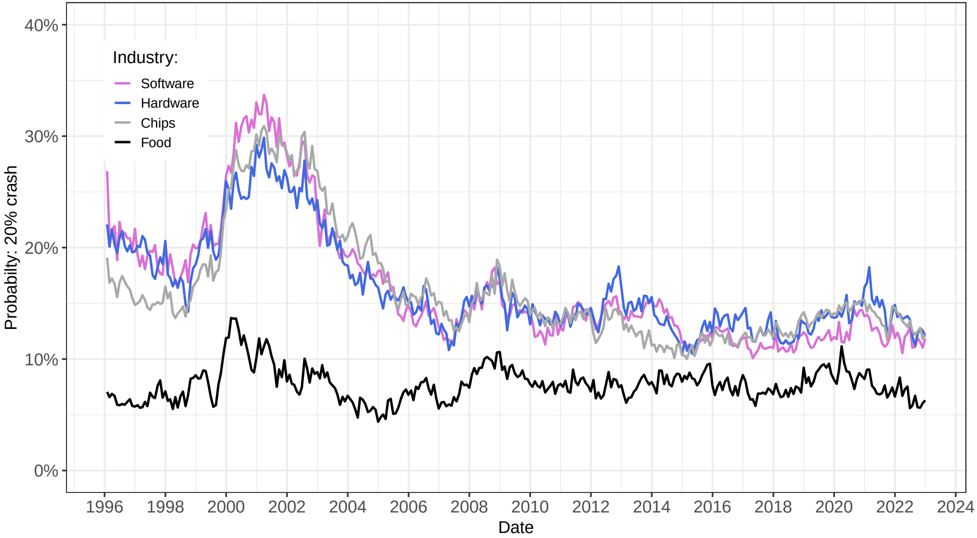

Industry average crash probabilities

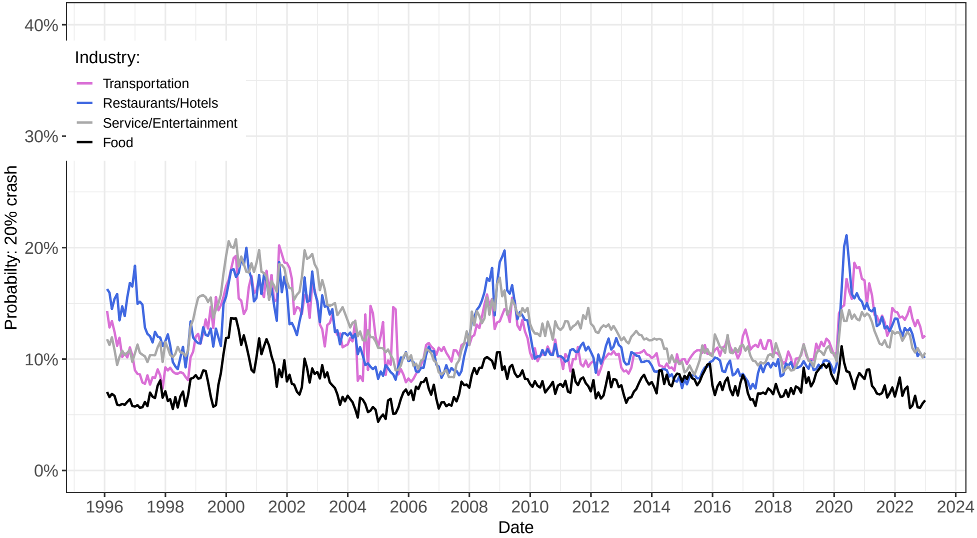

Summary

- 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 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