Forecasting Crashes with a Smile

Ian W. R. Martin Ran Shi*

April 2025

Abstract

We derive option-implied bounds on the probability of a crash in an individual stock, and argue that the lower bound should be close to the truth a priori. Empirically, the lower bound successfully forecasts crashes both in and out of sample; and it outperforms models based on stock characteristics previously studied in the literature. In a multivariate regression, a one standard deviation increase in the bound raises the predicted crash probability by 3 percentage points, whereas a one standard deviation increase in the next most important predictor (a measure of short interest) raises the predicted probability by only 0.3 percentage points.

*Martin: London School of Economics, i.w.martin@lse.ac.uk; Shi: University of Colorado Boulder, ran.shi@colorado.edu. We are grateful to John Campbell, Mike Chernov, Florian Heider, Mike Johannes, Christopher Jones, Dasol Kim, Lukas Kremens, Dong Lou, Stefan Nagel, Scott Richardson, Andrei Shleifer, and to participants in seminars at the 7th Marshall PhD Conference in Finance, LBS, LSE, HKU, the Office of Financial Research, Goethe University Frankfurt, Columbia Business School, Harvard University, the Federal Reserve Board, SITE 2024, BIS, UNSW Asset Pricing Workshop, and Schroders for helpful comments.

In this paper, we propose a new way of estimating the probability of a crash in an individual stock. Our approach performs well in and out of sample, across industries and over time, and it outperforms competitor models that exploit characteristics studied in prior related literature. As our forecasts are based solely on asset prices—namely, the prices of options on the stock in question, and of options on a broad stock index—they are, in principle, available in real time.

Aside from the intrinsic interest of forecasting crashes, we would like to highlight two further sources of motivation.

First, we can use our predictor variable to generate industry-level crash probability measures. Such series have many potential applications, but a particularly important one is suggested by the work of Baron, Verner, and Xiong (2021), who document a link between large declines in bank equity and macroeconomic downturns. Instead of defining bank equity declines using ex post returns, our approach supplies a measure of banks' equity capital crash risk (plotted in the top panel of Figure 10) that is forward-looking and observable in real time.

Second, our results relate to the issue of forecasting crashes at the aggregate level. Martin (2017) shows how to use option prices to calculate the probability of a crash in the market from the perspective of a log investor who holds the market. Barro and Liao (2021) derive an option-pricing formula within an equilibrium model and hence use index options to infer the probabilities of disasters in the sense of Barro (2006). But a fundamental challenge for these papers is that it is hard to test the predictive success of the resulting series directly, given the relatively short available time series and the fact that, by definition, crashes and disasters only occur rarely.[1] Expanding to the international data does not fully solve this problem, in part because few countries have long option prices series, and in part because, as Barro and Liao show, crashes are highly correlated across countries (as theory would predict even if there is little or no correlation in fundamentals, as in Martin (2013)). By exploiting the cross-section of stock prices, we gain statistical power that allows us to document the empirical success of option-based measures convincingly, thereby supporting the approach of these earlier papers.

From a theoretical point of view, forecasting crashes represents an interesting challenge

[1] Backus et al. (2011) also study the relationship between option prices and disaster probabilities, but they focus on the unconditional distribution whereas our goal is to quantify *conditional* crash probabilities.

for two reasons. First, there is an obvious and widely used competitor for our approach, namely, the risk-neutral probability of a crash, which can be calculated from asset prices without any assumptions other than the absence of arbitrage. And yet it is natural to worry that the risk-neutral probabilities, which put more weight on bad states of the world, may overstate the true probabilities of crashes. Furthermore, theory predicts that the risk-neutral probabilities will overstate most dramatically—that is, will be prone to "crying wolf"—at times of high risk or high risk aversion. This is unfortunate, given that policymakers are likely to be particularly interested in accurate measurement of crash probabilities in times of crisis.

Second, any attempt to forecast crashes in individual stocks using option prices seems to run into the problem that the inferred crash probability ought to reflect the correlation structure: the conclusions one would draw from a fixed set of prices should depend strongly on whether the stock in question has, for example, a positive or negative beta. But the prices of options on individual stocks and on the market reveal information only about the marginal risk-neutral distributions of those stocks and of the market, and not about their joint distribution.

We address these issues in two steps. To connect risk-neutral and true probabilities, we take the perspective of a myopic investor with power utility who chooses to invest his or her wealth fully in the S&P 500 index, which we treat as a proxy for "the market."[2] (While we think of this investor as a sensible benchmark, we do not need to assume that all investors look the same: our investor may coexist with others who have different beliefs, for rational or irrational reasons, and/or different preferences.) This implies that the stochastic discount factor (SDF) is proportional to a power of the return on the S&P 500 index. In the special case in which risk aversion equals zero, the predictive variable reduces to the risk-neutral probability of a crash, which can be inferred from out-of-the-money put option prices, following Breeden and Litzenberger (1978): this is a widely used indicator of crash probabilities but, as we will show, allowing for positive risk aversion improves predictive performance.

Evidently, the power utility assumption is restrictive. In an ideal world we would

[2]Related approaches have been adopted in the context of the stock market (Martin, 2017; Chabi-Yo and Loudis, 2020; Martin, 2021; Gao and Martin, 2021; Gandhi, Gormsen, and Lazarus, 2022), individual stocks (Martin and Wagner, 2019; Kadan and Tang, 2020; Chabi-Yo, Dim, and Vilkov, 2023), and currencies (Kremens and Martin, 2019; Della Corte, Gao, and Jeanneret, 2023).

allow the SDF to depend on broader measures of wealth and potentially other state variables. But option prices on the S&P 500 and on individual stocks are *observable*; and they are forward-looking. The great strength of our approach is that it allows us to avoid the alternative undesirable assumption, commonly made in the literature, that backward-looking historical measures are good proxies for the forward-looking measures that come out of theory. The empirical success of our approach suggests that the price of our assumption is worth paying.

Having made the assumption, it is straightforward to infer the true distribution of market returns from the risk-neutral distribution of market returns, as in Martin (2017). To calculate the true distribution of a given *stock's* returns, however, we would need to observe the *joint* risk-neutral distribution of that stock's and the market's returns. The problem is that observable option prices only allow us to infer the individual (that is, marginal) risk-neutral distributions of the stock and of the market, without giving us any control on the correlation structure. This is the central theoretical challenge.

We handle it by exploiting the theory of copulas and, more specifically, the Fréchet–Hoeffding bounds. These allow us to derive upper and lower bounds on the true probabilities of a crash that apply, under our maintained assumption on the form of the SDF, for *any* correlation structure. As the bounds fully exploit information in the two marginal distributions, they are tighter than naive bounds that only exploit the fact that correlation must lie between plus and minus one. (This paper might more accurately be titled "Forecasting Crashes with Two Smiles.")

The resulting bounds vary substantially across firms and over time. Figure 1 illustrates by plotting upper and lower bounds on the probability of a crash of at least 20% over a one-month horizon for Apple and AIG. Figure 2 plots the time-series of the cross-sectional median of the upper and lower bounds on crash probabilities, together with the probability of a crash in the market (with the latter calculated based on the approach in Martin (2017)). The market crash probability tends to be lower and less volatile than the individual stock probabilities.

As we will show, the lower bound is tight if the stock's return is a monotonic—and potentially nonlinear—increasing function of the market return, while the upper bound is tight if the stock's return is a monotonic decreasing function of the market return. The former case is more plausible, so we expect, of the two bounds, the lower bound to be a

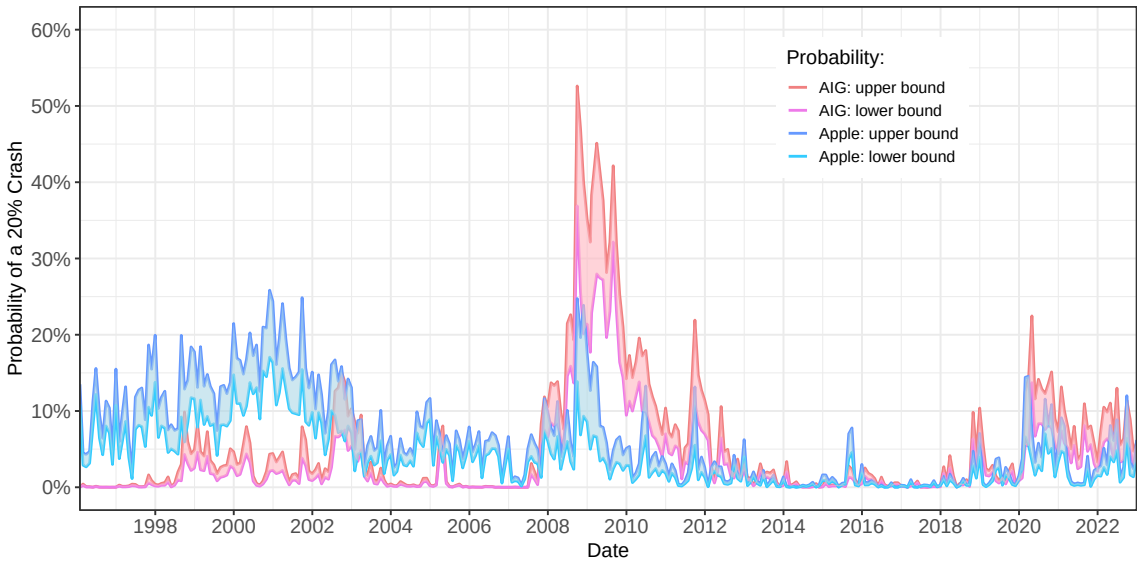


Figure 1: Bounds on forward-looking probabilities of a crash (one month return less than -20%) for Apple and AIG.

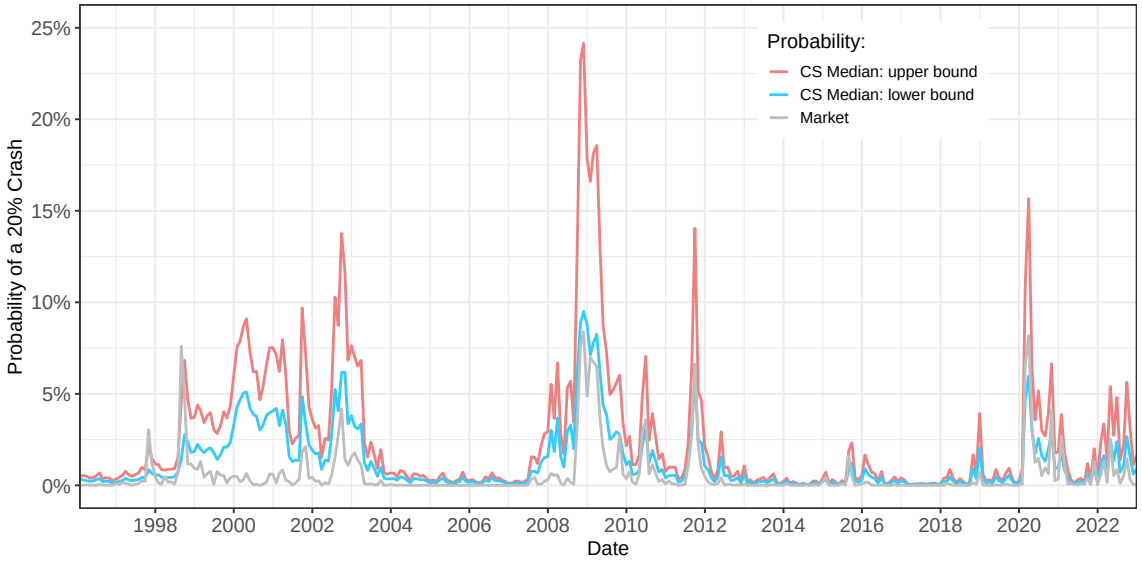


Figure 2: Time series of cross-sectional medians of upper and lower bounds on crash probabilities; and the crash probability of the S&P 500 index.

better measure of the true crash probability.

To assess this prediction, we regress realized crash indicators onto the upper bound and onto the lower bound. We consider crashes of size 10%, 20% and 30% over horizons of one, three, six, and 12 months, and find that both bounds are statistically significant predictors of crashes at all horizons and for all crash sizes. The same is true for the risk-neutral probability of a crash (which, as we show, must always lie between the upper and lower bounds).

If the lower bound were a perfect measure of the crash probability, we would find an intercept equal to zero and slope coefficient equal to one in the associated regression. And indeed we do find, for all 12 horizon/crash-size pairs, intercepts that are not significantly different from zero and slope coefficients that are significantly positive and close to one. The lower bound also outperforms the upper bound and the risk-neutral probability in an R^2 sense for all 12 horizon/crash-size pairs.

The lower bound remains significant when we include 15 stock characteristics that the prior literature has found useful in accounting for stock return variation, forecasting crashes, or predicting bankruptcies: CAPM beta, firm size, the book-to-market ratio, gross profits divided by total assets, three measures of trailing returns, realized volatility of market-adjusted returns, turnover, one-year sales growth, short interest scaled by institutional ownership, leverage, net income to assets, cash to assets, and log price per share. At horizons of one month and one quarter, the lower bound on its own achieves a higher R^2 than all 15 stock characteristics do together.

We also test the validity and tightness of our bounds for stock crash probabilities, following the approach of Back, Crotty, and Kazempour (2022). At all horizons and for all crash sizes, we do not reject the null that the bounds are valid (that is, the lower bounds are smaller and upper bounds are larger than the true crash probabilities). As expected, we strongly reject the hypothesis that the upper bound is tight (with p-values on the null hypothesis of tightness below 0.02 for 11 of the 12 horizon/crash-size pairs), while the evidence is mixed on whether the lower bound is tight: we do not reject tightness at the one-month and one-year horizons, but at the three-month and six-month horizons we can reject tightness (with p-values between 0.02 and 0.12).

The lower bound successfully predicts crashes out of sample. We use it as a crash forecaster with coefficient constrained to equal one, so that there are no parameters to es-

timate, and find that it outperforms the unadjusted risk-neutral probability at all horizons. It also outperforms an approach that attempts to salvage the risk-neutral probabilities by rescaling them, at each point in time t, by a stock- and time-specific parameter tuned to match the historical relationship between that stock's risk-neutral crash probability and its realized crash probability. We carry out this exercise in two ways, using either an expanding window or a rolling 3-year window to calibrate the historical relationship. In both cases, the lower bound outperforms at all horizons.

We next compare the lower bound to a "kitchen sink" model that optimally combines the 15 stock characteristics together with the risk-neutral probabilities *and the lower bound itself*, using ridge, Lasso (Tibshirani, 1996) or elastic net (Zou and Hastie, 2005) regularization with cross-validation to select models with the best out-of-sample performance. The lower bound on its own outperforms the kitchen sink model at all horizons for all three regularization methodologies, and for both rolling and expanding windows.

The problem with the kitchen sink model is that it chases performance by trying to fit the recent historical experience. Following the financial crisis of 2008–9, for example, the kitchen sink model puts extra weight on the risk-neutral probabilities, because their excessively pessimistic predictions happen to have been correct during that period. But this severely degrades the kitchen sink model's subsequent performance. This highlights the advantages of a well-specified structural approach relative to a black-box approach that chases performance by trying to fit the recent historical experience.

Related Literature. A large literature proposes methods to recover risk-neutral return densities from option prices. An incomplete list includes Breeden and Litzenberger (1978); Rubinstein (1994); Jackwerth and Rubinstein (1996); Aït-Sahalia and Lo (1998); Carr and Madan (2001). Christoffersen, Jacobs, and Chang (2013) provide a survey. While the starting point of our derivation relies on the insights of Breeden and Litzenberger (1978), the major challenge of bounding the physical, as opposed to risk-neutral, expectations is addressed by the new approach introduced in this paper.

Our work builds on a variety of papers that have studied the predictability of crashes. At the level of market indices, Bates (1991) explores the behavior of S&P 500 futures options prior to the stock market crash of 1987; more recently, Goetzmann, Kim, and Shiller (2022) link "crash narratives" to volatility and investor expectations about crashes.

At the level of individual stocks, Chen, Hong, and Stein (2001) show that character-

istics such as turnover and past returns forecast negative return skewness in individual stocks, Greenwood, Shleifer, and You (2019) use characteristics to forecast crashes at the industry level conditional on past price rises, and Daniel, Klos, and Rottke (2023) document that price run-ups combined with high short interest and low institutional ownership forecast lower stock returns. These papers focus on returns over horizons of at least several months: a distinctive feature of our approach is that it is empirically successful at forecasting crashes at horizons as short as 1 month.

There is also a literature that focuses on how measures of skewness and tail- or downside-risk are priced in the cross section of stock returns (see, for example, Ang, Chen, and Xing (2006); Boyer, Mitton, and Vorkink (2009); Vilkov and Xiao (2013); Kelly and Jiang (2014); Pederzoli (2021)).

Elsewhere in the economics literature, the Fréchet–Hoeffding inequalities have been applied by Heckman, Smith, and Clements (1997) and Manski (1997) in the context of programme evaluation; and Dybvig (1988) uses similar ideas to derive bounds on contingent claim prices in an economy with finitely many states.

Organization of the paper. Section 1 introduces our approach and establishes various theoretical properties of the bounds. Section 2 provides details of our data sample. Sections 3 and 4 present our empirical results, in and out of sample. Section 5 constructs industry-level crash probability measures. Section 6 studies the relationship between our measure of the crash probability and other characteristics. Section 7 concludes. All proofs are in the Appendix.

1 Theory

We adopt the perspective of an investor ("the investor") with power utility over next-period wealth who is marginal in all markets, including option markets, but who chooses to invest her wealth fully in the market, by which we mean the S&P 500 index. At time t, the investor chooses portfolio weights $\boldsymbol{w} = [w_1, \ldots, w_n]^\top$ to solve the problem[3]

$$\underset{\boldsymbol{w}}{\text{maximize}} \quad \mathbb{E}\left[u\left(\boldsymbol{w}^\top \boldsymbol{R}\right)\right] \quad \text{s.t.} \sum_{i=1}^{n} w_i = 1,$$

[3]All expectations are conditional on current, time t, information. We suppress time subscripts to streamline the notation.

where $u(x) = x^{1-\gamma}/(1-\gamma)$, risk aversion equals γ, and we write $\boldsymbol{R} = [R_1, \ldots, R_n]^\top$ for the vector of gross returns on the n assets from time t to time $t+1$. The first-order conditions for this problem are

$$\mathbb{E}\left[\left(\boldsymbol{w}^\top \boldsymbol{R}\right)^{-\gamma} R_i\right] = \lambda \quad \text{for all } i,$$

where λ is a Lagrange multiplier. By assumption, the investor chooses to invest fully in the market, thus the market return, R_m, satisfies $R_m = \boldsymbol{w}^\top \boldsymbol{R}$. It follows that $M = R_m^{-\gamma}/\lambda$ is a stochastic discount factor (SDF).

For any tradable payoff X, the risk-neutral expectation of X (which we denote with an asterisk) satisfies, by definition,

$$\frac{1}{R_f}\mathbb{E}^*[X] = \mathbb{E}[MX]$$

where R_f is the gross risk-free rate: the two sides of the above equation represent different notational conventions for expressing the price at time t of a claim to the payoff X paid at time $t+1$. As $M\lambda R_m^\gamma \equiv 1$, it follows that

$$\mathbb{E}[X] = \mathbb{E}[M\lambda R_m^\gamma X] = \lambda \mathbb{E}[M(R_m^\gamma X)] = \frac{\lambda}{R_f}\mathbb{E}^*[R_m^\gamma X]\,. \tag{1}$$

Setting $X = 1$, we must have $R_f = \lambda \mathbb{E}^*[R_m^\gamma]$, which allows us to eliminate λ from (1):

$$\mathbb{E}[X] = \frac{\mathbb{E}^*[R_m^\gamma X]}{\mathbb{E}^*[R_m^\gamma]}. \tag{2}$$

Hence we can infer the investor's expectation of X if we can *price* a claim to $R_m^\gamma X$.

For the rest of the paper we will assume that the payoff $X = h(R_i)$ is a well-behaved function of the return on a particular asset i, where $h : \mathbb{R}_+ \mapsto \mathbb{R}$ is continuous almost everywhere. We write Q_{mi} for the joint risk-neutral cumulative distribution function (CDF) of the market and individual stock return (R_m, R_i), and Q_m and Q_i for the marginal CDFs of R_m and R_i. Equation (2) can then be rewritten

$$\mathbb{E}[h(R_i)] = \frac{\int x^\gamma h(y)\, \mathrm{d}Q_{mi}(x,\, y)}{\int x^\gamma \, \mathrm{d}Q_m(x)}. \tag{3}$$

For example, if $X = \boldsymbol{I}(R_i \leq q)$ is the indicator function for the event that stock i's gross

return is less than q, then equation (2) implies that

$$\mathbb{P}[R_i \leq q] = \frac{\mathbb{E}^*[R_m^\gamma \boldsymbol{I}(R_i \leq q)]}{\mathbb{E}^*[R_m^\gamma]}, \tag{4}$$

because $\mathbb{E}[\boldsymbol{I}(R_i \leq q)] = \mathbb{P}[R_i \leq q]$. Equation (4) shows that we can in principle infer the true probability distribution of a particular stock return, as perceived by the power utility investor who is holding the market, from risk-neutral distributions.

If stock i's return were independent of the market return (under the joint risk-neutral measure), then equation (4) would imply that the true and risk-neutral probabilities of a crash in stock i would be equal. In practice, however, we do not expect independence to hold, as stocks are generally exposed to systematic risk. We can rewrite equation (4) as

$$\mathbb{P}\left[R_i \leq q\right] = \mathbb{P}^*\left[R_i \leq q\right] + \frac{\mathrm{cov}^*\left[R_m^\gamma, \boldsymbol{I}(R_i \leq q)\right]}{\mathbb{E}^*\left[R_m^\gamma\right]}. \tag{5}$$

As crashes are more likely to occur at times of market-wide bad news, the risk-neutral covariance term in equation (5) will typically be negative, so that the risk-neutral probabilities will overestimate actual crash probabilities.

Moreover, we should expect—and we will confirm, below—that the risk-neutral crash probabilities overstate the truth most dramatically at scary moments in time (i.e., when market-wide risk is high) and for scary stocks. This is unfortunate: it means that using risk-neutral probabilities as forecasters of crashes is likely to be most problematic at precisely the times (and for precisely those stocks for which) an accurate forecast would be most valuable. We therefore view equation (4) as a disciplined way of adjusting the risk-neutral probabilities to account for systematic (i.e., market) risk.

The challenge of implementing the equation, however, is that while index options and individual stock options reveal risk-neutral expectations of *univariate* functions of index or stock returns, they do not reveal risk-neutral expectations of two- (or higher-) dimensional functions of index *and* stock returns simultaneously, as would be needed to calculate the numerators in (3) or (4).[4] Options on the market and on large-cap individual stocks are liquid, but because they are written on a single underlying asset they reveal

[4] Ross (1976) showed in a finite-state setting that options on portfolios of assets could in principle be used to recover risk-neutral joint densities. Martin (2018) points out that this result fails with continuous states, and even with finite states given the assets that are traded in practice.

		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%

Figure 3: Two joint distributions that are each consistent with the marginal risk-neutral distributions in a 2×2 example.

only the marginal risk-neutral distributions, and not the correlation structure. To recover the risk-neutral joint distribution, one would need to observe the prices of derivatives whose payoffs are functions both of the stock index *and* of the stock of interest. But such prices are not observable in practice. (By contrast the probability of a crash in the market itself, as plotted in gray in Figure 2, is relatively easy to handle: when $i = m$ the right-hand side of (4) is a ratio of risk-neutral expectations of functions of the single random variable R_m, which can be calculated from index option prices in the usual way.)

An example. From the price of S&P 500 index options and options on Apple stock, we can calculate the risk-neutral probability of a crash in the market and in Apple, respectively. We might, for example, find that both Apple and the S&P 500 each have a risk-neutral crash probability equal to 5%. But these probabilities are consistent with a continuum of joint distributions.

The two panels of Figure 3 indicate two possible scenarios that are in a sense polar opposites. The left-hand panel describes a world in which Apple is risky, crashing if—and only if—the market crashes. In this case Apple crashes are "scary", so the risk-neutral crash probability (which, in standard models, distorts the true probability by a factor related to the marginal value of a dollar) should be expected to *overstate* the true probability of a crash. In the right-hand panel, by contrast, Apple crashes if and only if the market does not crash. In this case, Apple crashes are relatively benign, so the standard logic predicts that the risk-neutral probability should *understate* the true probability of an Apple crash.

More generally, although the observable one-dimensional risk-neutral distributions do not make it possible to pin down the true crash probability precisely, we can nonetheless derive bounds on the right-hand sides of (3) or (4): as in the 2×2 example, the marginals

place restrictions on the joint distribution.[5] To do so, we decompose the joint distribution into two parts: the marginals and the dependence structure. The marginals can be inferred from index and stock options, using the Breeden–Litzenberger approach. Roughly speaking, we can then bound the integral in the numerator of (3) by minimizing and maximizing over all possible dependence structures—more precisely, over all *copulas*.

Definition 1. *A (two-dimensional) copula is a function* $C : [0,1]^2 \mapsto [0,1]$ *such that*

1. C *is grounded:* $C(x,0) = C(0,y) = 0$ *for any* (x,y) *in its domain;*
2. $C(x,1) = x$ *and* $C(1,y) = y$ *for any* (x,y) *in its domain;*
3. C *is two-increasing: for all rectangles* $B = [x_1, y_1] \times [x_2, y_2] \subset [0,1]^2$*, the "volume" of* B*, which is defined by* $V_H(B) = C(x_2,y_2) - C(x_2,y_1) - C(x_1,y_2) + C(x_1,y_1)$ *is non-negative.*

The following theorem of Sklar (1959) shows that any joint distribution can be associated with a copula that "glues together" its marginals.

Theorem 1 (Sklar)**.** *Let* Q *be the joint CDF for the random vector* (X,Y) *with marginal CDFs* F_X *and* F_Y*. Then there exists a copula* C*, such that for all* $x, y \in \mathbb{R}$*,*

$$Q(x,y) = C(F_X(x),\, F_Y(y)).$$

We can therefore express the joint risk-neutral distribution of the market and stock return as $Q_{mi}(x,y) = C(Q_m(x), Q_i(y))$, where the risk-neutral index and individual stock CDFs, Q_m and Q_i, can be calculated from index and individual stock option prices. Although $C(\cdot,\cdot)$ is unknown, the following theorem supplies pointwise bounds that apply to any copula.

Theorem 2 (Fréchet–Hoeffding)**.** *If* $C(u,v)$ *is a copula, then*

$$\max(u+v-1, 0) \leq C(u,v) \leq \min(u,v), \quad (u,v) \in [0,1]^2.$$

Using the Fréchet–Hoeffding theorem, together with the work of Tchen (1980), we have the following result, whose proof is in the Appendix.

[5] Moreover, the diagonal and off-diagonal patterns visible in the two panels of Figure 3 are echoed in the form of the lower and upper bounds: see the proof of Result 1 in Appendix A.1.

Result 1. *For a continuous and two-increasing*[6] *function* g *defined on* $[0,\infty)\times[0,\infty)$*, we have the bounds*

$$\mathbb{E}^*\left[g\left(R_m,\, Q_i^{-1}\left(1-Q_m(R_m)\right)\right)\right] \leq \mathbb{E}^*[g(R_m,\, R_i)] \leq \mathbb{E}^*\left[g\left(R_m,\, Q_i^{-1}\left(Q_m(R_m)\right)\right)\right]\,. \quad (6)$$

Result 1 provides bounds on the price of an asset whose payoff $g(R_m, R_i)$ can depend in an arbitrary way on the correlation structure of R_m and R_i. As the one-dimensional risk-neutral distributions are observable from index and individual stock option prices, we can treat Q_i and Q_m as *observable functions*. Thus the upper and lower bounds in (6) are risk-neutral expectations of known functions of the *single* variable R_m. They can therefore be calculated given observable index option prices.

Result 1 exhibits bounds that relate risk-neutral expectations of different random variables to one another. It does not rely on any assumptions about the form of the SDF. But, under our assumption on the power utility form of the SDF, we can set $g(x,y) = x^\gamma h(y)$, as in equation (3), to derive the following result.[7]

Result 2. *Let* h *be a continuous increasing function defined on* $[0,\infty)$ *that does not cross the* x*-axis (that is,* $h(x)h(y)\geq 0$ *for any* $x,y\geq 0$*), and suppose the SDF is proportional to* $R_m^{-\gamma}$*. Then*

$$\frac{\mathbb{E}^*\left[R_m^\gamma h\left(Q_i^{-1}\left(1-Q_m(R_m)\right)\right)\right]}{\mathbb{E}^*\left[R_m^\gamma\right]} \leq \mathbb{E}[h(R_i)] \leq \frac{\mathbb{E}^*\left[R_m^\gamma h\left(Q_i^{-1}\left(Q_m(R_m)\right)\right)\right]}{\mathbb{E}^*\left[R_m^\gamma\right]}.$$

These bounds are sharp, in the sense that the lower bound is achieved if R_i *and* R_m *are countermonotonic, and the upper bound is achieved if* R_i *and* R_m *are comonotonic.*[8]

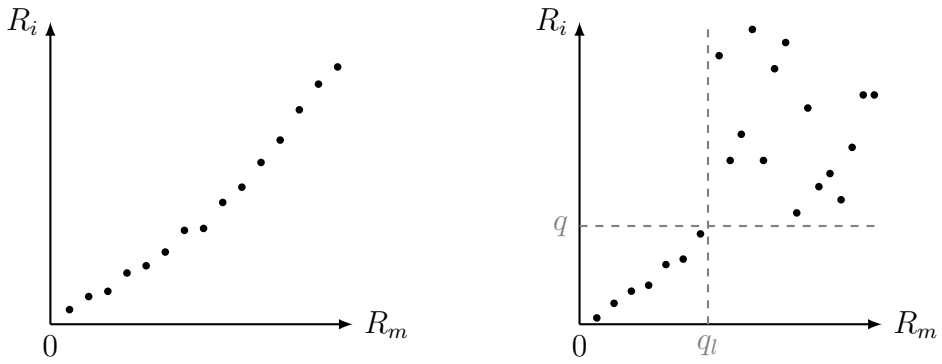
Note that the middle expectation above is a true—not a risk-neutral—expectation. In our application to crash probabilities, we set $h(x) = -\boldsymbol{I}(x\leq q)$ in Result 2. This delivers the following variant of Result 2 on which our empirical work is based. As the bounds depend only on tail outcomes for the stock, the lower bound is tight under the weaker condition that the stock return and market return are comonotonic *in the left tail*.

More precisely, the condition for tightness of the lower bound is expressed in terms of

[6] See Definition 1.

[7] In Appendix B, we show how to adapt Result 2 if h is Lipschitz continuous but not monotonic.

[8] Two random variables are said to be countermonotonic if one is a monotonically decreasing transformation of the other, and comonotonic if one is a monotonically increasing transformation of the other.



Comonotonic

Lower comonotonic with thresholds (q_l, q)

Figure 4: Left: an example of a comonotonic pair of random variables. Right: an example of a lower comonotonic pair of random variables with thresholds (q_l, q).

lower comonotonicity (a weaker assumption than comonotonicity). Specifically, (R_m, R_i) is *lower comonotonic* with threshold (q_l, q) if, in the terminology of Cheung (2009), $(-R_m, -R_i)$ is upper comonotonic with threshold $(-q_l, -q)$. Figure 4 illustrates.

Result 3. *The probability of a crash in stock i, $\mathbb{P}[R_i \leq q]$, satisfies the bounds*

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

where $q_l = Q_m^{-1}(Q_i(q))$ and $q_u = Q_m^{-1}(1 - Q_i(q))$.

The lower bound is attained if the return on the market R_m and return on the stock R_i are lower comonotonic with thresholds (q_l, q); this holds, in particular, if R_m and R_i are comonotonic. The upper bound is attained if the two returns are countermonotonic.

The risk-neutral probability of a crash, $\mathbb{P}^*[R_i \leq q]$, lies between the two bounds. It is equal to the true crash probability if stock i's return is independent of the market return.

As most stocks typically move with, rather than against, the market, we anticipate that comonotonicity is closer to the truth than countermonotonicity. Hence, a priori, we expect the lower bound to be tighter—closer to the true crash probability—than the upper bound. As we will see below, our empirical results strongly support this expectation.

Result 3 takes a particularly simple form in the special case in which asset i is the market index. As the market is comonotonic with itself, the lower bound holds with

equality, so that the probability of a market crash is

$$\mathbb{P}[R_m \leq q] = \frac{\mathbb{E}^* \left[R_m^\gamma \boldsymbol{I} \left(R_m \leq q \right) \right]}{\mathbb{E}^* \left[R_m^\gamma \right]}. \tag{7}$$

Equation (7) generalizes one of the results in Martin (2017), and we use it to calibrate the level of risk aversion, γ. We choose to calibrate risk aversion using market returns—rather than by targeting forecasting performance for individual stocks—to minimize the effect of in-sample information. As we show in Appendix C, setting γ equal to two approximately optimizes the predictive performance of our framework for movements in the market. For the rest of the paper, we therefore set $\gamma = 2$.

Our next result shows that the bounds in Result 3 widen as risk aversion rises.

Result 4. *The lower bound is decreasing in γ and the upper bound is increasing in γ.*

When $\gamma = 0$, the lower and the upper bounds are both equal to $\mathbb{P}^*[R_i \leq q]$: this is the case in which the true and risk-neutral expectations coincide, so that crash probabilities can be inferred perfectly from option prices.

As $\gamma \to \infty$, the bounds become trivial: for any q such that $0 < Q_i(q) < 1$, the lower bound tends to 0 and the upper bound tends to 1.

It follows that higher risk aversion leads to more conservative bounds: increasing risk aversion drives the lower bound down and the upper bound up.

It only remains to show how we calculate the bounds that appear in Result 3. Given a chosen value of q, and hence of q_l and q_u, the risk-neutral expectations that appear in the bounds can be calculated from index option prices. The only point at which the prices of options *on stock i itself* are used is, therefore, in the calculation of q_l and q_u, which are determined by the prices of index and of individual stock options via the risk-neutral marginals $Q_m(\cdot)$ and $Q_i(\cdot)$.

Result 5. *For any $\gamma > 0$, we can calculate the risk-neutral expectations in Result 3 using observable option prices:*

$$\mathbb{E}^* \left[R_m^\gamma \right] = R_f^\gamma + \frac{R_f}{S_0^\gamma} \left[\int_0^F \gamma(\gamma-1) K^{\gamma-2} \text{put}(K) \, \mathrm{d}K + \int_F^\infty \gamma(\gamma-1) K^{\gamma-2} \text{call}(K) \, \mathrm{d}K \right],$$

$$\mathbb{E}^* \left[R_m^\gamma \boldsymbol{I} \left(R_m \leq q_l \right) \right] = R_f q_l^\gamma \left[\text{put}'(K_l) - \gamma \frac{\text{put}\left(K_l\right)}{K_l} \right] + \frac{R_f}{S_0^\gamma} \int_0^{K_l} \gamma(\gamma-1) K^{\gamma-2} \text{put}\left(K\right) \, \mathrm{d}K,$$

$$\mathbb{E}^*\left[R_m^\gamma \boldsymbol{I}\left(R_m \geq q_u\right)\right] = R_f q_u^\gamma \left[\gamma \frac{\text{call}(K_u)}{K_u} - \text{call}'(K_u)\right] + \frac{R_f}{S_0^\gamma} \int_{K_u}^{\infty} \gamma(\gamma-1) K^{\gamma-2} \text{call}\left(K\right) \, \mathrm{d}K,$$

where S_0 *is the spot price of the market index;* $F = R_f S_0$ *is the forward price;* $\text{put}(K)$ *and* $\text{call}(K)$ *are the prices of index put and call options; and* $K_l = q_l S_0$ *and* $K_u = q_u S_0$.

1.1 Fréchet–Hoeffding vs. Cauchy–Schwarz

The bounds in Result 3 are stronger than the bounds that follow from the fact that correlation lies between plus and minus one (that is, from the Cauchy–Schwarz inequality). To compare the two approaches, note that equation (5) implies that

$$\mathbb{P}^*\left[R_i \leq q\right] - \frac{\sigma^*\left[R_m^\gamma\right] \sigma^*\left[\boldsymbol{I}(R_i \leq q)\right]}{\mathbb{E}^*\left[R_m^\gamma\right]} \leq \mathbb{P}\left[R_i \leq q\right] \leq \mathbb{P}^*\left[R_i \leq q\right] + \frac{\sigma^*\left[R_m^\gamma\right] \sigma^*\left[\boldsymbol{I}(R_i \leq q)\right]}{\mathbb{E}^*\left[R_m^\gamma\right]}, \tag{8}$$

where $\sigma^*\left[\cdot\right] = \sqrt{\text{var}^*\left[\cdot\right]}$ denotes risk-neutral volatility. These bounds depend only on univariate risk-neutral expectations, so can be calculated from observable option prices.

But this approach is less efficient than the bounds derived above because, in general, comonotonic random variables are not perfectly positively correlated and countermonotonic random variables are not perfectly negatively correlated. This is easily seen in our application, because the stock crash indicator (a binary variable) *cannot* be a linear function of the market return (a continuous variable).[9] It is thus *impossible* for the risk-neutral correlation between the two to reach plus or minus one. We report Fréchet–Hoeffding-implied bounds on risk-neutral correlations for stocks in our sample in Table A4 of the appendix.

It therefore follows that bounds obtained by "setting correlation equal to one" (or to minus one) will be looser than the bounds supplied by Result 3. Indeed, the upper and lower bounds implied by the Cauchy–Schwarz inequality need not even lie between zero and one. Table A5, in the appendix, reports the relative widths of our bounds compared with the Cauchy–Schwarz bounds across firms. For all crash sizes and horizons, the Cauchy–Schwarz bounds are substantially wider than the bounds based on the Fréchet–Hoeffding theorem. In the case of the 1 month/20% pair, the Cauchy–Schwarz bounds

[9] For an example in which both variables are continuous, suppose that Z is Normal. Then e^Z and $e^{\sigma Z}$ are comonotonic if $\sigma > 0$ and countermonotonic if $\sigma < 0$. But as σ tends to plus or minus infinity, the correlation between the two tends to zero.

are, on average, more than three times wider than the Fréchet–Hoeffding bounds, while for the more extreme 1 month/30% pair, the Cauchy–Schwarz bounds are on the order of ten times wider on average.

2 Data

We focus on firms included in the S&P 500 index, using index constituent information from CRSP. Our sample runs from January 1996 to December 2022. On the last trading day of each month t, we obtain, from OptionMetrics, the volatility surfaces of the S&P 500 index and of all firms that are S&P 500 constituents during month t, together with risk-free rates. We then obtain stock prices, returns, trading volumes and shares outstanding from CRSP to construct a firm-month panel. We emphasize that, unlike much of the literature, we do *not* drop financial firms from our sample, as our theory applies equally well to financial firms as to other industries

We face the issue that individual stock options are American style rather than European style. We deal with this issue, following the related literature (Carr and Wu, 2009; Kelly, Lustig, and Van Nieuwerburgh, 2016; Christoffersen, Fournier, and Jacobs, 2018; Martin and Wagner, 2019), by using volatility surfaces reported by OptionMetrics, who use proprietary multinomial tree models to account for early exercise premia. This is not a perfect solution, but we believe that the distinction is likely to be relatively minor for our applications, as the calculations required by Results 3 and 5 depend on the prices of out-of-the-money options.

When calculating the integrals in Result 5, we extrapolate a flat volatility smile outside the range of observed strikes, as is also standard in the literature. We provide additional computational details on the construction of our bounds in Section D of the Appendix.

We write $R_{i,t\to t+\tau}$ for the gross return on stock i from time t to time $t+\tau$, $\mathbb{P}^L_{i,t}(\tau, q)$ and $\mathbb{P}^U_{i,t}(\tau, q)$ for the lower and upper bounds on the probability that $R_{i,t\to t+\tau}$ is less than or equal to q, and $\mathbb{P}^*_{i,t}(\tau, q)$ for the corresponding risk-neutral probability.

Table 1 reports summary statistics for these measures with $q = 70\%$, 80% and 90%, at 1, 3, 6, and 12 month horizons. For comparison, we also report the realized frequencies of crashes. Specifically, for each month from January 1996 to December 2022, we calculate cross-sectional averages of the realized crash indicator $\boldsymbol{I}(R_{i,t\to t+\tau} \leq q)$ (which equals one

Table 1: Summary statistics

This table presents summary statistics of realized crash events, our crash probability bounds, and risk-neutral crash probabilities. The sample data are monthly from January 1996 to December 2022. The crash events (realized crashes) under consideration are $\boldsymbol{I}(R_{i,t\to t+\tau} \leq q)$ for $q = 0.7$, 0.8, 0.9 and $\tau = 1$, 3, 6, 12 months. The bounds and risk-neutral probabilities are measures of the conditional probabilities of crash events.

		averaged across firms (number of obs. $T = 324$)				averaged across time (number of obs. $N = 1044$)			
	horizon	1	3	6	12	1	3	6	12
Panel A: $q = 0.7$, 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
Panel B: $q = 0.8$, 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
Panel C: $q = 0.9$, 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

if the realized return is less than or equal to q, and zero otherwise), the upper and lower bounds, and the risk-neutral crash probabilities. The first four columns of the table report the means and standard deviations of these $T = 324$ observations at each of the four horizons. Similarly, we calculate time-series averages of the same quantities for each of the $N = 1044$ firms in our sample. The last four columns of Table 1 report the means and standard deviations of these time-series averages. The sample means of cross-sectional and time-series averages differ slightly because we have an unbalanced panel.

Consistent with the predictions of the theory and the discussion following Result 3, the time-series and cross-sectional means of the lower bounds are close to the corresponding mean realized crash frequencies, whereas the risk-neutral probabilities and (even more so) the upper bounds overestimate the likelihood of crashes.

3 In-sample results

3.1 Regression tests

To test whether the option-implied bounds successfully measure the probability of a crash, we run the regression

$$\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) = \alpha + \beta X_{i,t}(\tau, q) + \varepsilon_{i,t+\tau} \tag{9}$$

for a range of crash sizes q and forecasting horizons τ (measured in months). Here $X_{i,t}(\tau, q)$ is the lower or upper bound on the crash probability (that is, $\mathbb{P}^L_{i,t}(\tau, q)$ or $\mathbb{P}^U_{i,t}(\tau, q)$), or the risk-neutral crash probability, $\mathbb{P}^*_{i,t}(\tau, q)$. Result 3 showed, under our maintained assumptions, that the inequality

$$\mathbb{P}^L_{i,t}(\tau, q) \leq \mathbb{P}_t[R_{i,t\to t+\tau} \leq q] \leq \mathbb{P}^U_{i,t}(\tau, q)$$

holds for any stock i, forecasting horizon τ, and crash size q. If, moreover, one of the bounds is close to the true crash probability, we should find α close to zero and β close to one in the corresponding regression.

The regression results are shown in Table 2, which reports two-way clustered standard errors in parentheses, following Thompson (2011), and block bootstrapped standard errors

Table 2: Regression tests of the option-implied crash probability measures

This table reports the results of running linear regressions

$$\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) = \alpha + \beta X_{it}(\tau, q) + \varepsilon_{i,t+\tau},$$

in which $q = 0.7, 0.8$ and 0.9, and X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^U$ (the upper bounds), or $\mathbb{P}^*$ (the risk-neutral probabilities). The data are monthly from January 1996 to December 2022. The stocks under consideration are S&P 500 constituents. The return horizon τ is one month, three months, six months, or one year. Values in parentheses are firm-month two-way clustered standard errors following Thompson (2011). Values in square brackets are standard errors following the panel bootstrap procedures of Martin and Wagner (2019) using 2500 bootstrap samples.

	lower bound				risk-neutral				upper bound			
horizon	1	3	6	12	1	3	6	12	1	3	6	12
Panel A: $q = 0.70$, down by over 30%												
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%
Panel B: $q = 0.80$, down by over 20%												
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%
Panel C: $q = 0.90$, down by over 10%												
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%

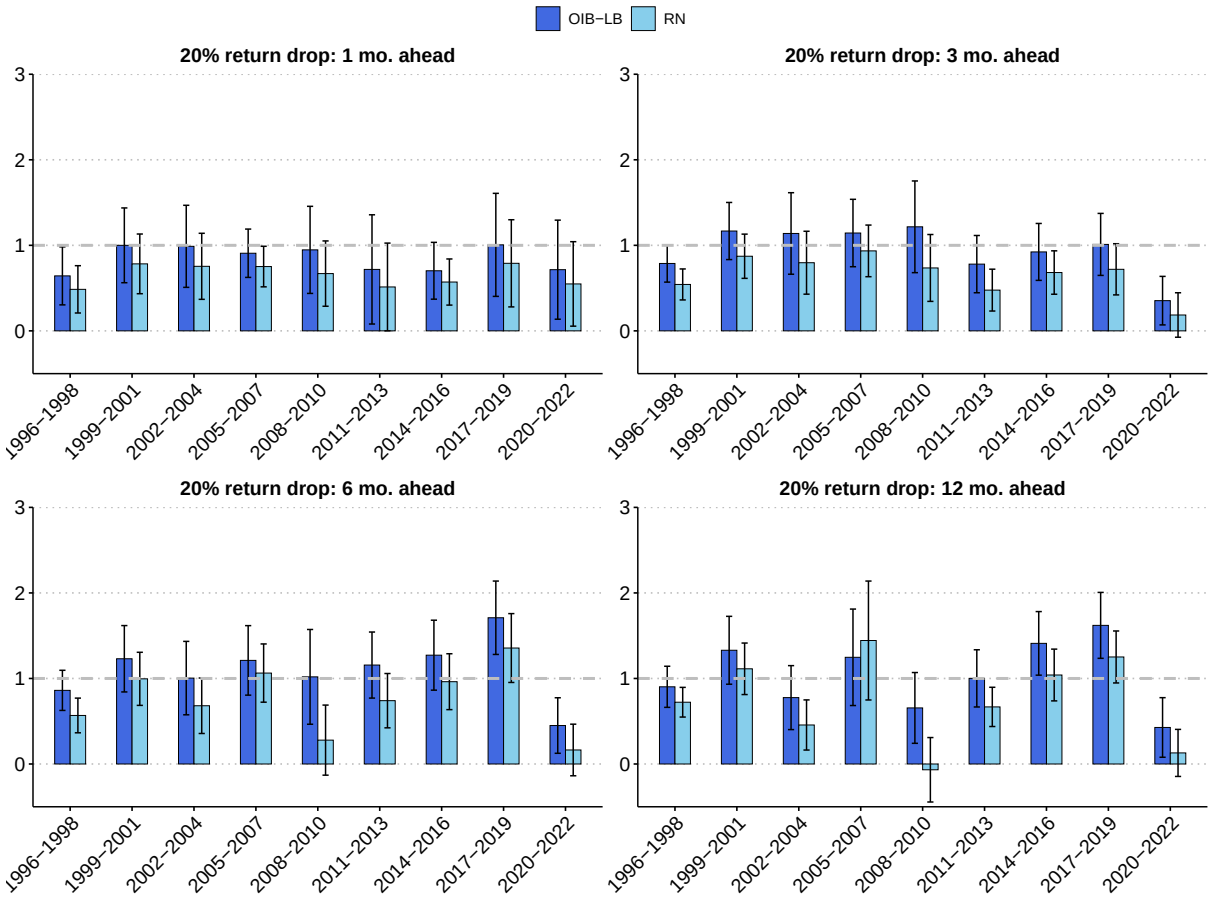


Figure 5: Regression coefficients β for the lower bounds (OIB-LB) and the risk-neutral (RN) probabilities in subsamples.

This figure presents the regression coefficients of equation (9) in nine completely different subsamples, when the independent variable is the lower bound, forecasting horizons are $\tau = 1, 3, 6, 12$ months, and the crash size $q = 0.8$. The height of a blue bar represents the β estimate; error bars represent the 95% confidence intervals calculated based on firm-month clustered standard errors.

in square brackets, using the procedure of Martin and Wagner (2019). Across crash sizes and forecasting horizons—and for all three right-hand side variables—the estimated intercepts are close to zero, while the estimated slope coefficients are positive and strongly significant.

The estimated slope coefficients exhibit a clear monotonic pattern[10] that is consistent with the theory. The estimated coefficients on the lower bound are largest (averaging around 1.05 across crash sizes and horizons); the estimated coefficients on the risk-neutral probability are significantly below one (averaging around 0.70); and the estimated coeffi-

[10]Recall from Result 3 that the risk-neutral probability must lie between the upper and lower bounds.

cients on the upper bound are smallest (averaging around 0.50).

In the case of the lower bound, the estimated coefficients are insignificantly different from one at all horizons and for all crash sizes. The lower bound also outperforms the other two variables in an R^2 sense for almost all horizons and crash sizes. These results are consistent with the discussion following Result 3.

Tables A6, A7 and A8, in the appendix, report the same regressions with time fixed effects, firm fixed effects, and time *and* firm fixed effects, respectively. (Although such specifications are not useful for prediction without prior knowledge of the values of the fixed effects, they help us to understand where the success of the predictor variables comes from.) Table A6 shows that the slope coefficients are little changed by the introduction of time fixed effects: thus our measures successfully explain cross-sectional variation in crash probabilities. Tables A7 and A8 show that the slope coefficients remain highly significant at short horizons when firm fixed effects are included, either on their own or even jointly with time fixed effects (but not at the 12 month horizon). For example, at the one-month horizon with both time and firm fixed effects included, the coefficient on the lower bound is 0.77 for the largest crashes ($q = 0.7$), 0.73 for intermediate crashes ($q = 0.8$), and 0.65 for the smallest crashes ($q = 0.9$), with standard errors in the range 0.05 to 0.13.

The good performance of the lower bound is not driven by any particular episode. Figure 5 plots the β coefficient estimates for equation (9) when the lower bounds $\mathbb{P}^L_{i,t}(\tau, 0.8)$ are used to forecast 20% crashes in nine non-overlapping three-year subsamples, at horizons of 1, 3, 6, and 12 months. The coefficient estimates are significantly greater than zero in all subsamples and across all forecasting horizons. The null $\beta = 1$ cannot be rejected in 31 out of the 36 cases at the 95% confidence level. For comparison, the coefficient estimates for risk-neutral probabilities exhibit more pronounced variation over time. Standard theory would suggest that risk-neutral probabilities are likely to overstate true crash probabilities, particularly at times of heightened market risk or risk aversion and for stocks that are highly exposed to systematic risk. Consistent with this view, we find that the coefficients on the risk-neutral probabilities are closer to zero (and insignificantly different from zero at horizons from 3 to 12 months) in subsamples covering the global financial crisis and the Covid pandemic.

The lower bound adjusts the risk-neutral probabilities to account for risk, and the adjustment is disciplined by theory. By contrast, ad hoc adjustments to the risk-neutral

probabilities do not seem to work as well. To illustrate, consider the following simple approach[11] to "debiasing" the risk-neutral probabilities:

1. At time t, calculate firm i's mean historical realized crash probability, $\hat{p}_i$, and mean historical risk-neutral crash probability, $\hat{p}_i^*$, by averaging, respectively, its historical crash event indicators and historical risk-neutral crash probabilities from the beginning of the sample until time t.

2. Generate an adjusted crash probability measure for firm i at time t by multiplying the time t risk-neutral probability of a crash in firm i by $\hat{p}_i/\hat{p}_i^*$.

3. Repeat for all time periods t and firms i.

We repeat the in-sample regression analysis for these adjusted risk-neutral probabilities and report the results in Table A9 in the appendix. For ease of comparability, the table also reports the corresponding results for the lower bounds and the raw risk-neutral probabilities (taken from Table 2). By construction, the adjusted risk-neutral probabilities are close to the true crash probability when averaged over the entire sample. But they are poor predictors of crashes, with low R^2s; and the estimated coefficients on the adjusted probabilities are significantly below one—and the associated intercepts significantly above zero—for all horizons and crash sizes.

Forecasting rallies. Result 7, in Appendix E, adapts the above analysis to derive bounds on the probability of a rally, rather than a crash, in a stock. In this case it is the *upper* bound that is attained if the return on the stock and on the market are comonotonic (or under the weaker condition that the return on the stock and on the market are comonotonic in the right tail), so a priori we expect the upper bound to perform better in forecasting rallies. Table A2 reports results of regressions of a realized rally indicator function onto the lower and upper bounds, and onto the risk-neutral probability of a rally. While the upper bound is a highly significant predictor of rallies, the coefficient is significantly smaller than 1 at horizons above 1 month, which suggests that the comonotonicity assumption is further from the truth in rally regimes than in crash regimes. This is broadly consistent with the finding that stock return correlations are countercyclical (Campbell et al., 2001).

[11] We consider other ways of adjusting the risk-neutral probabilities in the out-of-sample analysis discussed in Section 4.

Forecasting returns. It is natural to wonder whether the success in forecasting crashes documented in Table 2 translates directly into forecasting returns. It is not completely clear what to expect, however. One possibility is that crashes are simply the consequence of market mispricing. On this view, if our measure helps to detect stocks which are overpriced, and hence likely to crash, then it might also forecast low realized returns on average. At the opposite end of the spectrum, the market efficiency view would suggest that stocks with a high probability of crashing should have *high*, not low, expected returns in compensation for this risk.[12] In practice we find that the evidence does not come down decisively in favor of either side. Table A10 reports the results of regressions of returns onto the option-implied bounds and onto the risk-neutral crash probabilities. At the 12-month horizon there is some evidence of a statistically significant and positive relationship between crash probability and subsequent returns; at shorter horizons, the relationship remains positive but is not statistically significant.

3.2 Validity and tightness tests

We now carry out formal tests of the validity and tightness of the crash probability bounds based on conditional moment restrictions, following Back, Crotty, and Kazempour (2022) (henceforth, BCK).

Let $\boldsymbol{z}_t$ be a strictly positive vector of dimension d that incorporates conditioning variables known at time t. This vector includes a set of candidate variables that might help to determine crash probabilities, and it determines another vector, of the same length,

$$\lambda = \mathbb{E}\left[\left\{\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) - X_{i,t}(\tau, q)\right\} \boldsymbol{z}_t\right],$$

where X represents lower or upper bounds. As each element of $\boldsymbol{z}_t$ is strictly positive, we can assess the validity of the lower bound[13] by testing $\lambda \geq 0$ against the alternative that $\lambda \in \mathbb{R}^d$ (that is, λ is unrestricted). If the lower bound is valid, we can assess its *tightness* by testing $\lambda = 0$ against the alternative $\lambda \geq 0$. Similarly, we can assess the validity of the upper bound by testing $\lambda \leq 0$ against the alternative $\lambda \in \mathbb{R}^d$, and assess its tightness by

[12] As an example, a riskless bond is guaranteed not to crash, and earns no risk premium as a result.

[13] If the lower bound is valid then $\mathbb{E}_t\left[\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) - X_{i,t}(\tau, q)\right] \geq 0$, where $X_{i,t}(\tau, q) = \mathbb{P}^L_{i,t}(\tau, q)$. As $\boldsymbol{z}_t$ is known at time t and strictly positive, it follows that $\mathbb{E}_t\left[\left\{\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) - X_{i,t}(\tau, q)\right\} \boldsymbol{z}_t\right] \geq 0$, and hence that $\mathbb{E}\left[\left\{\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) - X_{i,t}(\tau, q)\right\} \boldsymbol{z}_t\right] \geq 0$ by the law of iterated expectations.

testing $\lambda = 0$ against the alternative $\lambda \leq 0$.

Following BCK, we include a constant in the vector $\boldsymbol{z}_t$, together with additional variables from Welch and Goyal (2008), transformed where necessary to guarantee positivity. We then construct the estimator

$$\hat{\lambda} = \frac{1}{T}\sum_{t=1}^{T}\left[\frac{1}{N_t}\sum_{i=1}^{N_t}\left\{\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) - X_{i,t}(\tau, q)\right\}\boldsymbol{z}_t\right],$$

where N_t is the number of firms at time t, and estimate the variance-covariance matrix of $\hat{\lambda}$ using the Driscoll and Kraay (1998) estimator to account for heteroskedasticity and serial correlation in the time series and cross sectional dependence across firms.

Table 3 reports the results of the BCK tests. The headline result is that we do not reject validity of either the upper or the lower bound at any horizon or crash size.

For all horizons and crash sizes (except the extreme one month/30% scenario, where we have lower power due to the relatively smaller number of crashes) we can, however, strongly reject the hypothesis that the upper bound is tight. This was to be expected: the upper bound is tight only if stock returns and the market return are countermonotonic—that is, if all individual stock returns are monotonically decreasing functions of the market return. This is implausible, even as an approximation, for a single stock; and it *cannot* hold for all stocks given that the market return is a weighted average of individual stock returns.

By contrast, as noted above, we expect a priori that the lower bound should be closer to the truth. Here the evidence of tightness is mixed. We do not reject tightness of the lower bound for any crash size at the 1-month horizon (with p-values on the null varying between 0.134 and 0.414) or 12-month horizon (p-values between 0.116 and 0.161), but at the 3- and 6-month horizons we can generally reject tightness with moderate confidence.

3.3 Comparison with other predictor variables

The previous section established that the theoretically motivated quantity $\mathbb{P}^L_{i,t}(\tau, q)$ is a strongly significant forecaster of crashes. We now investigate whether this empirical success survives the introduction of various stock characteristics, and compare the lower bound more directly with the forecasting performance of the risk-neutral probability of a

Table 3: Validity and tightness of the option-implied crash probability bounds: the Back, Crotty, and Kazempour (2022) tests

This table reports the p-values for testing the validity and tightness of our proposed bounds, using the methodology described in Back, Crotty, and Kazempour (2022). The data are monthly from January 1996 to December 2022. Firms under consideration are S&P 500 constituents. The return horizons, denoted by τ, are one month, three months, six months, and one year. For $q = 0.7, 0.8$ and 0.9, define

$$\lambda = \mathbb{E}\left[\{\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) - X_{it}(\tau, q)\}\, \boldsymbol{z}_t\right],$$

where X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^U$ (the upper bounds), or $\mathbb{P}^*$ (the risk-neutral probabilities); the elements of $\boldsymbol{z}_t$ are 1) a constant one, 2) the dividend yield of the market, 3) the earnings yield of the market, 4) the spread between five-year and three-month treasury yields, 5) the net equity issuance scaled by the market capitalization, 6) the month-to-month inflation rate, 7) the BAA-AAA credit spread, 8) the book-to-market ratio of the market, 9) the three-month treasury yield and 10) the VIX index. $H_0 : \lambda \geq 0$ vs. $H_1 : \lambda \in \mathbb{R}^d$ tests if a lower bound is valid; $H_0 : \lambda = 0$ vs. $H_1 : \lambda \geq 0$ tests if a lower bound is tight; $H_0 : \lambda \leq 0$ vs. $H_1 : \lambda \in \mathbb{R}^d$ tests if an upper bound is valid; $H_0 : \lambda = 0$ vs. $H_1 : \lambda \leq 0$ tests if an upper bound is tight.

	lower bound				upper bound			
horizon	1	3	6	12	1	3	6	12
Panel A: $q = 0.70$, 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: $q = 0.80$, 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: $q = 0.90$, 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

crash. From now on we focus on declines of at least 20% in the interest of brevity.

We consider several stock characteristics, including seven variables associated with the cross-section of expected stock returns: CAPM beta, relative size (the logarithms of a firm's market capitalization scaled by that of the S&P 500 index), book-to-market ratio, gross profitability (gross profits scaled by total assets), two momentum measures (stock returns from month -6 to month -1 and month -12 to month -1), and the most recent month's return (as a reversal signal).

We also consider various characteristics studied by a prior literature that has studied related (though not identical) topics: the volatility of market-adjusted returns and average monthly turnover (both of which are highlighted in Chen, Hong, and Stein (2001)), sales growth (Greenwood, Shleifer, and You, 2019), short interest scaled by institutional ownership (Asquith, Pathak, and Ritter, 2005; Daniel, Klos, and Rottke, 2023), and four variables motivated by the approach of Campbell, Hilscher, and Szilagyi (2008) to forecasting corporate bankruptcies and failures: the leverage (debt-to-asset) ratio, net income scaled by total assets, cash and short-term investment scaled by total assets, and log price per share. Appendix F gives further detail on the construction of all 15 characteristics, and Table A3 presents summary statistics.

Table 4 reports results for a crash of at least 20% over the next month. To make it easier to assess the economic significance of the forecasting variables, we rescale the lower bound, the risk-neutral probability, and all stock characteristics to have unit standard deviation, and we multiply coefficient estimates by 100. Each coefficient therefore measures the influence, in percentage points, of a one standard deviation move in the relevant variable. Asterisks indicate coefficients with t-statistics greater than 4 in absolute value.[14]

The first column of the table reports results for a multivariate regression of the crash indicator variable onto the stock characteristics described above. Together, the characteristics achieve an R^2 of 4.51%. Two of the characteristics are highly significant: the volatility measure of Chen, Hong, and Stein (2001) has a t-statistic around 7, and short interest scaled by institutional ownership has a t-statistic above 4. From now on, we refer to this volatility measure as *CHS volatility*.

The second column shows that the lower bound, on its own, performs better than the stock characteristics do collectively. It explains more of the variation in crashes, with R^2

[14]We choose a high threshold to avoid false positives, as recommended by Harvey, Liu, and Zhu (2016).

Table 4: Regression tests of the option-implied crash probability bounds: adjusted regressions for 20% crashes in one month

This table reports the results from the following regressions:

$$\boldsymbol{I}(R_{i,t\to t+1} \leq 0.8) = \beta \cdot X_{it}(1, 0.8) + \lambda \cdot \text{controls}_{it} + \varepsilon_{i,t+1},$$

in which X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^*$ (the risk-neutral probability), or both. The controls are 15 firm characteristics from the literature. All independent variables are transformed to have a unit standard deviation. Regression coefficients are reported in percentage points, and their two-way clustered standard errors are included in the parentheses. The first five columns are simple OLS estimates, and the sixth column reports estimates with time fixed effects, with a projected (within) R^2 replacing the standard ones. Asterisks indicate coefficients whose t-statistics exceed four in magnitude.

	$\boldsymbol{I}(R_{t\to t+1} \leq 0.8)$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{P}^L[R_{t\to t+1} \leq 0.8]$		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)	
beta	0.48		0.12	0.18	0.10	0.22
	(0.15)		(0.16)	(0.17)	(0.14)	(0.13)
relative size	0.06		−0.02	−0.04	−0.00	0.09
	(0.10)		(0.10)	(0.10)	(0.10)	(0.07)
book-to-market	−0.18		−0.20	−0.20	−0.20	−0.07
	(0.11)		(0.10)	(0.10)	(0.10)	(0.08)
gross profit	−0.13		−0.08	−0.10	−0.07	−0.04
	(0.09)		(0.09)	(0.09)	(0.08)	(0.07)
$r_{(t-1)\to t}$	−0.27		−0.08	−0.06	−0.09	−0.16
	(0.18)		(0.18)	(0.18)	(0.19)	(0.13)
$r_{(t-6)\to(t-1)}$	−0.45		−0.26	−0.26	−0.27	−0.34
	(0.20)		(0.19)	(0.20)	(0.20)	(0.16)
$r_{(t-12)\to(t-1)}$	−0.06		−0.06	−0.05	−0.07	−0.14
	(0.19)		(0.19)	(0.19)	(0.18)	(0.17)
CHS-volatility	2.27*		0.31	0.44	0.32	0.50
	(0.31)		(0.37)	(0.44)	(0.39)	(0.18)
turnover	0.18		−0.06	−0.07	−0.05	0.08
	(0.26)		(0.25)	(0.24)	(0.24)	(0.15)
sales growth	0.21		0.19	0.20	0.19	0.12
	(0.11)		(0.10)	(0.11)	(0.10)	(0.08)
short int.	0.39*		0.34*	0.37*	0.33*	0.27*
	(0.09)		(0.08)	(0.08)	(0.08)	(0.06)
leverage	−0.14		−0.10	−0.13	−0.08	−0.06
	(0.12)		(0.12)	(0.12)	(0.10)	(0.11)
net income-to-asset	−0.20		−0.13	−0.17	−0.12	−0.13
	(0.12)		(0.12)	(0.12)	(0.11)	(0.08)
cash-to-asset	−0.09		−0.09	−0.08	−0.09	−0.03
	(0.08)		(0.08)	(0.08)	(0.08)	(0.07)
log price	−0.34		0.13	0.05	0.15	0.06
	(0.16)		(0.15)	(0.15)	(0.17)	(0.13)
intercept	0.04	0.00	−0.02	−0.01	−0.03	
	(0.03)	(0.00)	(0.03)	(0.03)	(0.03)	
R^2/R^2-proj.	4.49%	5.65%	5.82%	5.69%	5.83%	4.72%

of 5.66%, and is highly statistically significant, with a t-statistic above 8. (This regression is identical, up to the rescaling, to the regression with an estimated coefficient of 0.92 reported in Panel A of Table 2.)

The third column reports results of a multivariate regression that uses both the lower bound and the stock characteristics to forecast crashes. The lower bound remains highly significant, with a t-statistic above 5. Of the stock characteristics, only short interest remains statistically significant, and the collective marginal contribution to explanatory power of the characteristics is small. The coefficient on the lower bound is roughly an order of magnitude greater than that on short interest: a one standard deviation move in the lower bound moves the implied crash probability by 3.02 percentage points, whereas a one standard deviation move in short interest moves the implied crash probability by 0.34 percentage points.

Columns (4) and (5) of the tables include the risk-neutral probability of a crash, either alone as an alternative to the lower bound, or together with it. At all three horizons, the risk-neutral probability enters strongly significantly when included on its own, but achieves a lower R^2 than the lower bound does. When both are included together, the coefficient on the lower bound is positive while that on the risk-neutral probability is negative; but the coefficients are imprecisely estimated, as the lower bound and risk-neutral probability are highly correlated.

Tables A11 and A12, in the Appendix, report similar results over horizons of one quarter and one year, respectively. As before, we rescale all right-hand side variables to have unit standard deviation so that coefficient estimates indicate the economic importance of the various potential predictors. The lower bound remains highly significant both in statistical (the t-statistic is large) and economic (the estimated coefficient is large) terms. In the univariate regression at the one-year horizon, for example, a one standard deviation increase in the lower bound represents a 6.90 percentage point increase in the probability of a crash, with a t-statistic above 12. When all stock characteristics are included, the coefficient estimate drops to 5.17, with a t-statistic above 7. The stock characteristics are more informative at this longer horizon: sales growth and short interest are highly statistically significant with estimated coefficients of 1.76 and 2.32, respectively.

3.3.1 A "kitchen sink" approach

We conclude the in-sample section by trying a "kitchen sink" approach that also includes 15 squared characteristics (for example, book-to-market$^2_{i,t}$) and $\binom{15}{2} = 105$ interactions between characteristics (for example, book-to-market$_{i,t}$ $\times$ sales growth$_{i,t}$).

Panel A of Table 5 reports the estimated coefficient on the lower bound across all horizons. As in Table 4, we standardize all variables so that the reported coefficient estimates can be interpreted as the effect of a 1 standard deviation move in the lower bound on the forecast crash probability, measured in percentage points.

The left and middle blocks report results with no characteristics included, and with the 15 characteristics included, at all horizons.[15] The rightmost block reports results when the characteristics, squared characteristics, and interaction terms are all included. The estimated coefficient on the lower bound remains highly significant, and is reasonably stable across specifications.

Panel B of the same table reports results when the risk-neutral probabilities are also included. At the 1-month horizon, the relatively high correlation of the lower bound and risk-neutral probabilities substantially increases the standard error on the lower bound estimate; at other horizons, the lower bound remains highly significant. Meanwhile, the coefficients on the risk-neutral probabilities are either insignificantly different from zero or significantly negative, and the inclusion of the risk-neutral probabilities has little impact on the regression R^2.

Summarizing, the risk-neutral probabilities do not contribute much incremental explanatory power, whereas the characteristics contain useful information in sample, especially in the "kitchen sink" implementation and for crashes over longer horizons. As we will now see, however, the characteristics and kitchen-sink approaches perform relatively poorly out of sample.

4 Out-of-sample forecasts

The literature has shown that there can be large gaps between in-sample fit and out-of-sample forecasting power, even if the true data generating process is stable over time

[15] At the 1-month horizon, the blocks with no controls and with the 15 characteristics repeat the coefficient estimates reported in columns (2) and (3) of Table 4.

Table 5: Regression tests of the option-implied crash probability bounds: additional adjusted regressions for 20% crashes

This table reports the results from the following regressions:

$$\boldsymbol{I}(R_{i,t\to t+\tau} \leq 0.8) = \beta \cdot X_{it}(\tau, 0.8) + \lambda \cdot \text{controls}_{it} + \varepsilon_{i,t+\tau}, \quad \tau = 1, 3, 6, 12$$

in which X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^*$ (the risk-neutral probability), or both. The first four columns do not include controls. Columns 5-8 include the 15 firm characteristics as controls. Columns 9-12 include the 15 firm characteristics, 15 squared characteristics, and 105 pairwise interaction terms. All independent variables are normalized to have unit standard deviation. Regression coefficients for the lower bound (OIB-LB) and risk-neutral probabilities (RN) are reported in percentage points, and two-way clustered standard errors are included in parentheses.

	control: none				control: characteristics				control: "kitchen sink"			
horizon	1	3	6	12	1	3	6	12	1	3	6	12
Panel A: the lower bound												
OIB-LB	3.4	5.7	6.8	6.9	3.0	4.1	5.5	5.2	2.8	3.9	4.9	4.4
	(0.4)	(0.5)	(0.5)	(0.5)	(0.6)	(0.7)	(0.8)	(0.7)	(0.7)	(0.8)	(0.8)	(0.7)
R^2	5.6%	5.1%	4.8%	3.7%	5.8%	5.6%	5.6%	5.0%	6.5%	6.8%	7.4%	7.5%
Panel B: the lower bound + risk-neutral probabilities												
OIB-LB	4.7	11.3	11.9	9.5	4.4	10.8	12.0	9.1	2.1	10.3	12.1	9.1
	(3.1)	(2.7)	(1.7)	(2.0)	(3.1)	(2.8)	(1.9)	(2.0)	(3.1)	(3.0)	(1.9)	(1.9)
RN	−1.3	−5.7	−5.3	−2.8	−1.4	−6.6	−7.0	−4.5	0.7	−6.4	−7.8	−5.6
	(3.2)	(2.8)	(1.8)	(2.2)	(3.4)	(2.9)	(1.9)	(2.1)	(3.4)	(3.1)	(2.0)	(2.1)
R^2	5.7%	5.4%	5.0%	3.8%	5.8%	5.9%	5.9%	5.2%	6.5%	7.0%	7.7%	7.7%

(Timmermann, 1993; Lewellen and Shanken, 2002; Martin and Nagel, 2022). In this section, we assess the lower bound's performance out of sample, using it to forecast crashes with coefficients α and β set equal to 0 and 1, respectively. This specification is well suited to out-of-sample testing because there are no free parameters to be estimated.

We assess forecasting performance using the out-of-sample R^2 measure

$$R^2_{\text{oos}} = 1 - \frac{\sum_t \sum_i \left\{\boldsymbol{I}(R_{i,t\to t+\tau} \leq 0.8) - \mathbb{P}^L_t\left(R_{i,t\to t+\tau} \leq 0.8\right)\right\}^2}{\sum_t \sum_i \left\{\boldsymbol{I}(R_{i,t\to t+\tau} \leq 0.8) - p_{i,t}\right\}^2}, \tag{10}$$

where τ denotes the forecasting horizon, $\mathbb{P}^L_t\left(R_{i,t\to t+\tau} \leq 0.8\right)$ is the lower bound on the probability of a 20% crash, and $p_{i,t}$ is the historical average crash probability for firm i estimated over the period from 1 to $(t-\tau)$. It follows from this definition that R^2_{oos} increases if and only if the forecaster performs better in the newly included sample than

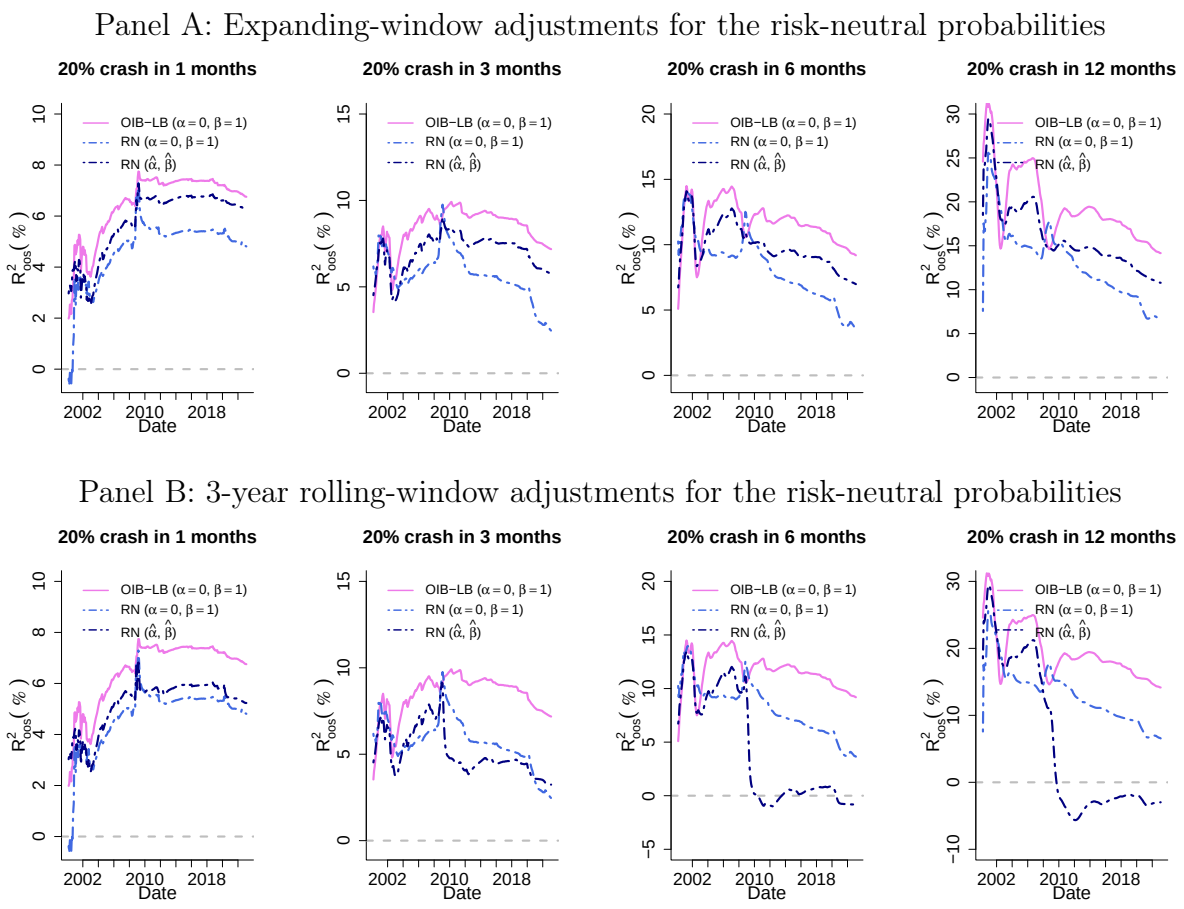


Figure 6: Out-of-sample R^2s: the lower bound, the risk-neutral, and the adjusted risk-neutral forecasts

This figure presents the out-of-sample R^2s (R^2_{oos}) for our option-implied lower bound (OIB-LB). At each time point t, we compare the sum of squared forecasting errors from OIB-LB to those from a firm-specific average probability of crashes, calculated over the period $1:(t-\tau)$ during which crash events are observable at time t; $\tau(=1,3,6,12)$ represents the forecasting horizon. The note $(\alpha=0,\beta=1)$ emphasizes that no adjustments have been made for the lower bound. For comparison, R^2_{oos} values are also presented for the risk-neutral probabilities (RN, $\alpha=0$, $\beta=1$) and their derived forecasts from expanding-window (Panel A) or 3-year rolling-window (Panel B) linear regressions (RN, $\hat{\alpha},\hat{\beta}$) in dark blue.

it does in the trailing sample.

For comparison, we calculate the corresponding out-of-sample R^2s for the risk-neutral probabilities, replacing $\mathbb{P}^L_t\left(R_{i,t\to t+\tau}\leq 0.8\right)$ in equation (10) with $\mathbb{P}^*_t\left(R_{i,t\to t+\tau}\leq 0.8\right)$. We also report results for an adjusted risk-neutral probability forecast $\hat{\alpha}_t+\hat{\beta}_t\mathbb{P}^*_t\left(R_{i,t\to t+\tau}\leq 0.8\right)$, where $(\hat{\alpha}_t,\hat{\beta}_t)$ are regression estimates, as in (9), to correct for the upward biases in the

Table 6: Diebold–Mariano tests of equal forecasting accuracy: the lower bound, the risk-neutral and the adjusted risk-neutral forecasts

This table reports p-values from Diebold–Mariano tests (Diebold and Mariano, 1995) of the predictive performance of the lower bound versus the risk-neutral probabilities and versus the adjusted risk-neutral forecasts. We consider 20% crashes at horizons of 1, 3, 6 or 12 months. The competing forecaster in Panel A is the risk-neutral probability without regression adjustments. Panel B and C adjust the risk-neutral probabilities through trailing regressions using expanding windows (Panel B) or 3-year rolling windows (Panel C). For a crash indicator y and a forecaster $\hat{y}$, the squared error loss and cross entropy loss are defined as $(y-\hat{y})^2$ and $-y\log\hat{y}-(1-y)\log(1-\hat{y})$ respectively. The spectral density of the loss differential is estimated using the Driscoll and Kraay (1998) estimator, accounting for general cross-firm crash correlations.

horizon	1	3	6	12
Panel A: risk-neutral ($\alpha=0,\beta=1$)				
Squared error	0.028	0.010	0.046	0.098
Cross entropy	0.059	0.007	0.053	0.189
Panel B: risk-neutral (expanding $\hat{\alpha},\hat{\beta}$)				
Squared error	0.124	0.056	0.056	0.048
Cross entropy	0.005	0.015	0.007	0.015
Panel C: risk-neutral (3-year rolling $\hat{\alpha},\hat{\beta}$)				
Squared error	0.045	0.120	0.046	0.004
Cross entropy	0.000	0.003	0.017	0.002

risk-neutral probabilities. We report in Panel A the results based on expanding-window regressions, and in Panel B based on 3-year rolling-window regressions.

Figure 6 shows that the lower bound outperforms the risk-neutral probabilities and the adjusted risk-neutral forecasts at all horizons. We note in particular that an attempt to correct the upward bias in the risk-neutral probabilities by re-estimating $\hat{\alpha}$ and $\hat{\beta}$ in rolling-regressions substantially worsens forecasting performance at the longer horizons. Simply chasing in-sample fittings on recent observations in search for more accurate crash forecasts can be counterproductive. The decline in the forecasting performance of the risk-neutral probabilities around 2008–2010 echoes our findings in Figure 5: consistent with standard theory, the extent to which risk-neutral probabilities overstate true crash risk is greatest at times of high market risk. Rolling regressions further exacerbate this issue in the aftermath of a realized crash.

To compare the predictive performance of the lower bound versus the risk-neutral probabilities and versus the adjusted risk-neutral forecasts formally, we report p-values

Table 7: Diebold–Mariano tests of equal forecasting accuracy: the lower bound and the regression forecasts aggregating all variables

This table reports the p-values from the Diebold–Mariano tests (Diebold and Mariano, 1995) of whether the lower bound (without any regression adjustments) is generating the same crash forecasting losses as the regression aggregates. We consider over 20% stock-level crashes in the next one, three, six or twelve months. The competing forecaster in the first block is the lower bound, the risk-neutral probabilities, and the 15 stock characteristics aggregated from trailing elastic net regressions (the tuning parameters are selected from cross-validations). The second block also include the 15 squared terms and the 105 interaction terms of the characteristics. Coefficients for regression aggregates are estimated from expanding windows (Panel A) or 3-year rolling windows (Panel B). For a crash indicator y and a forecaster $\hat{y}$, the squared error loss and cross entropy loss are defined as $(y-\hat{y})^2$ and $-y\log\hat{y}-(1-y)\log(1-\hat{y})$ respectively. The spectral density of the loss differential is estimated using the Driscoll and Kraay (1998) estimator, accounting for general cross-firm crash correlations.

	OIB-LB+RN+char.				OIB-LB+RN+char.+char.^{2}			
horizon	1	3	6	12	1	3	6	12
Panel A: expanding-window regressions								
Squared error	0.037	0.004	0.031	0.004	0.024	0.050	0.022	0.000
Cross entropy	0.000	0.000	0.031	0.001	0.000	0.001	0.000	0.000
Panel B: rolling 3-year window regressions								
Squared error	0.001	0.010	0.049	0.030	0.001	0.009	0.035	0.024
Cross entropy	0.000	0.000	0.108	0.099	0.000	0.000	0.053	0.050

of Diebold–Mariano tests (Diebold and Mariano, 1995) in Table 6, using both squared and cross-entropy loss functions (with the latter inducing heavier penalties on misaligned forecasts). The p-values on the null of equal predictive accuracy range from 0.000 to 0.189, allowing us to reject this null at conventional significant levels in most cases. The results are least favorable to our theory when we compare the lower bound to the raw risk-neutral probabilities at the 12-month horizon, perhaps because this is the horizon at which the comonotonic crash assumption underpinning the superior performance of the lower bound is most likely to be loose.

To consider a priori more competitive forecasters, we design a procedure to emulate an avid "data-snooper" that aggregates the 15 stock characteristics examined in Section 3.3—together with the risk-neutral probabilities and our lower bound—to construct crash predictors. We select models using the elastic net estimator (minimizing OLS losses subject to L_1 and L_2 penalties as in Zou and Hastie (2005)). The degree of shrinkage is determined from five-fold cross-validation using the training samples; we apply the same tuning parameter for both penalty terms to avoid the well-known indeterminacy issues in

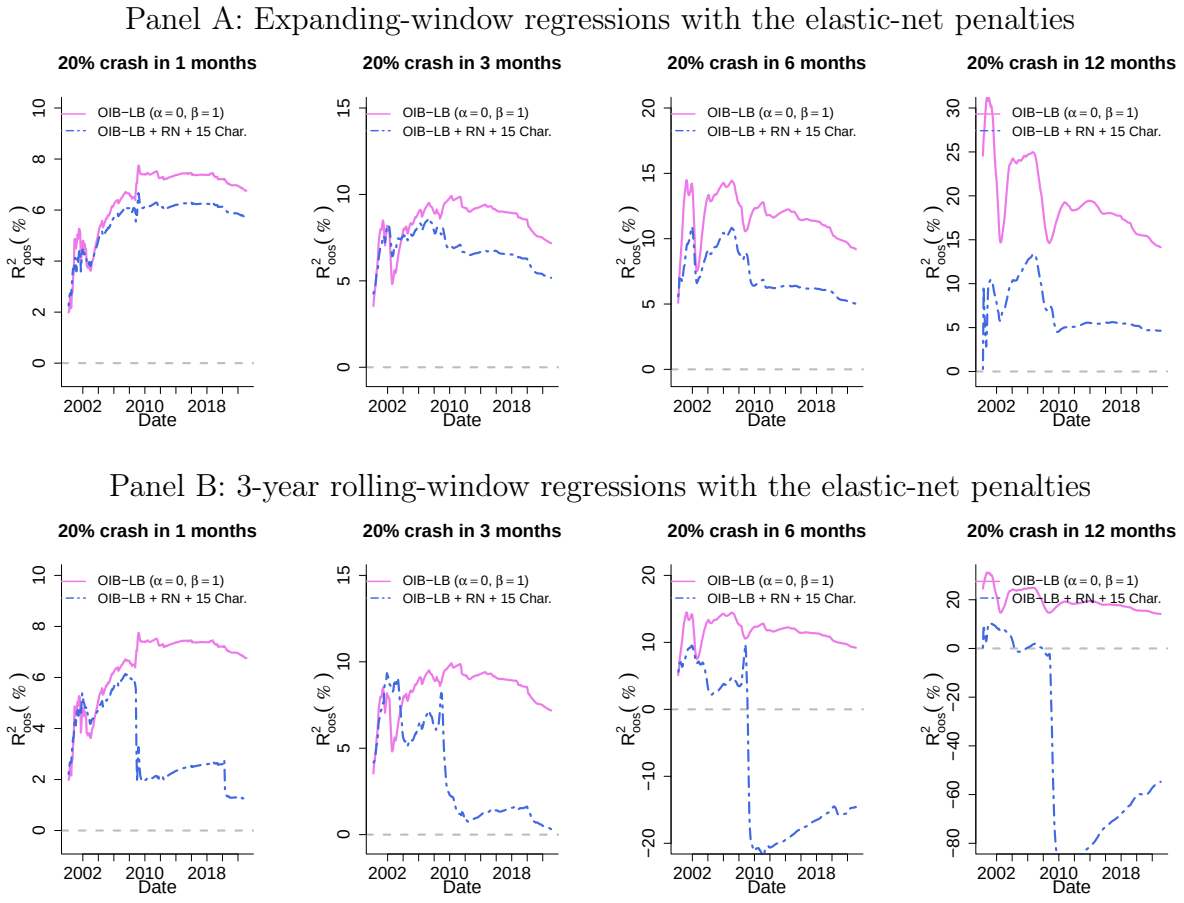


Figure 7: Out-of-sample R^2s: the lower bound and the regression forecasts aggregating all variables

This figure presents the out-of-sample R^2s (R^2_{oos}) for our option-implied lower bound (OIB-LB). At each time point t, we compare the sum of squared forecasting errors from OIB-LB to those from a firm-specific average probability of crashes, calculated over the period $1:(t-\tau)$ ($\tau = 1, 3, 6, 12$). For comparison, we also report R^2_{oos}s for a forecaster that aggregates the 15 stock characteristics considered in Session 3.3, the lower bound, and the risk-neutral probabilities. The variables are combined through expanding-window (Panel A) or 3-year rolling-window (Panel B) elastic-net regressions, the tuning parameters of which are chosen through 5-fold cross-validations.

cross-validating the two elastic net penalties simultaneously.

Figure 7 presents the results, using expanding windows in Panel A and 3-year rolling windows in Panel B. In both cases, the lower bound outperforms at all horizons. We find similar results using Lasso (Tibshirani, 1996) or ridge regularization: see Figures A3 and A4 in the appendix.

Figure 8 summarizes the coefficient estimates at each point in time. To make it possible

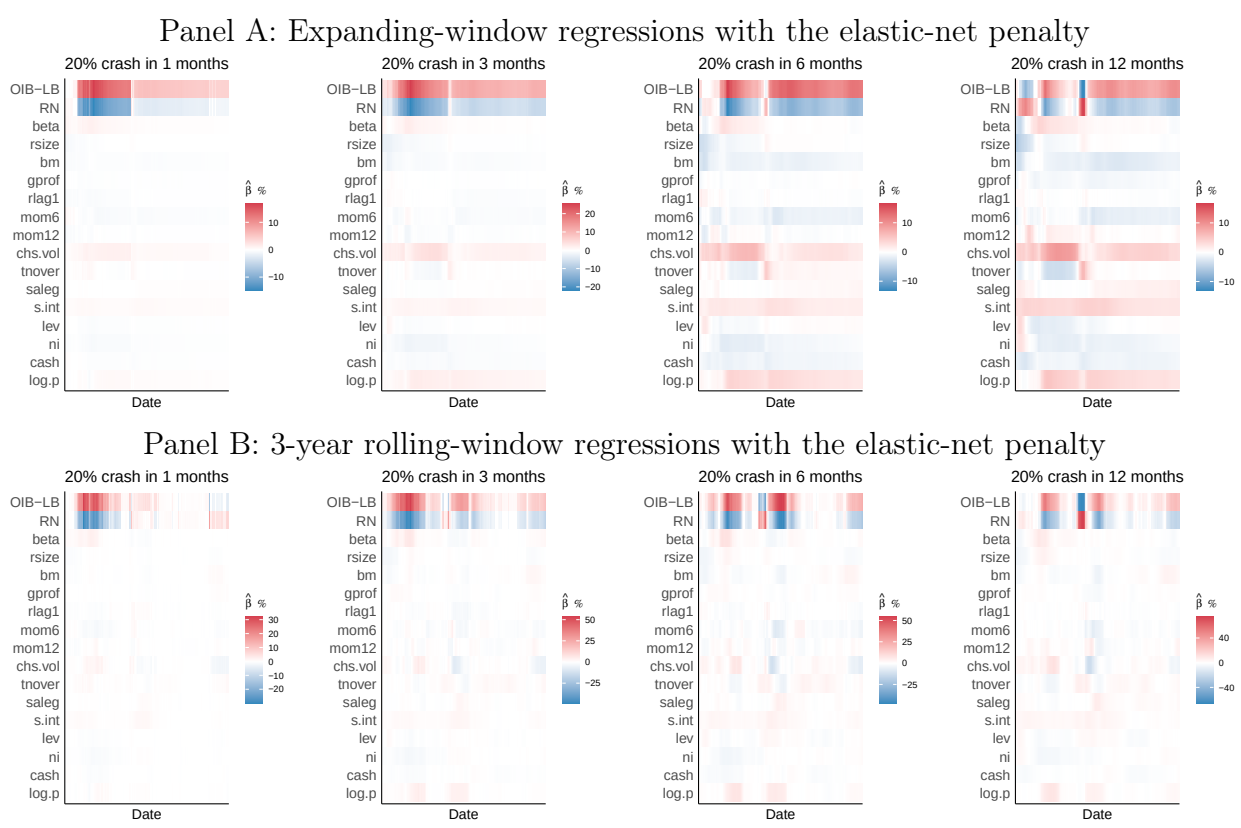


Figure 8: Regression coefficient estimates from the elastic-net regressions

This figure presents regression coefficient estimates using the elastic-net estimator for aggregating the 15 stock characteristics considered in Session 3.3 (their descriptions are presented in Appendix F), the lower bound, and the risk-neutral probabilities. The training samples come from expanding windows (Panel A) or 3-year rolling windows (Panel B). The tuning parameters are chosen through 5-fold cross-validations. To compare across the seventeen features, the estimates are divided by the standard deviations of their corresponding variables in the training samples. All values are in percentage points.

to compare the magnitudes of coefficients, all explanatory variables are standardized in each regression; the coefficients can therefore be interpreted as showing the change in crash probability (measured in percentage points) associated with a one standard deviation move in the associated variable. Panel A shows that in expanding-window regressions, the elastic net gradually "learns" that the lower bound is a useful predictor of crashes. At the 1-month horizon, in particular, the model ends up relying almost exclusively on the lower bound; at longer horizons, it loads positively on the lower bound and negatively on the risk-neutral probabilities, with some further role for other characteristics at the 6- and 12-month horizons.

During, and immediately following, the global financial crisis, the elastic nets switch to favor the risk-neutral probabilities, whose pessimistic predictions appear correct in hindsight. This pattern is particularly stark for the rolling window regressions shown in Panel B. But by "chasing fit" in this way, the elastic nets underperform dramatically in the aftermath of the crisis, as is visible in Figure 7, Panel B.

Figure A2, in the appendix, shows out-of-sample R^2 plots that document similarly poor performance at all horizons using the "kitchen sink" set of variables considered in Section 3.3.1.

Table 7 reports the p-values of Diebold–Mariano tests. With characteristics included, the null of equal forecasting accuracy is now always rejected. In sharp contrast to the in-sample evidence reported in Table 5, out-of-sample forecasting performance deteriorates when characteristics are included.

5 Crash probabilities across industries

In this section, we use our crash predictor variable to generate industry-specific crash probability measures. As we noted in the introduction, motivation for this exercise is provided by Baron, Verner, and Xiong (2021).

As a first step, however, we must check that the lower bound performs well within industries. Figure 9 reports coefficient estimates in regressions of crash indicators onto the lower bounds separately for stocks belonging to each of the 17 Fama–French industries. We focus on 20% crashes over horizons of 1, 3, 6, and 12 months. The estimated coefficients are significant for almost all industries and horizons, and are close to 1.

We therefore generate industry crash probability measures by averaging individual crash probabilities across stocks in a given industry. Figure 10 shows the average probabilities of 20% crashes in one year for various industries, defined using the Fama–French 49 industry classification. The top panel represents the financial sector (recall that unlike much of the cross-sectional asset pricing literature we do not remove financial firms because our theory applies quite generally). The middle panel represents the information technology sector, and the bottom panel the service sector. In each panel, we also include the average crash probabilities of stocks in the food industry as a relatively stable comparison industry.

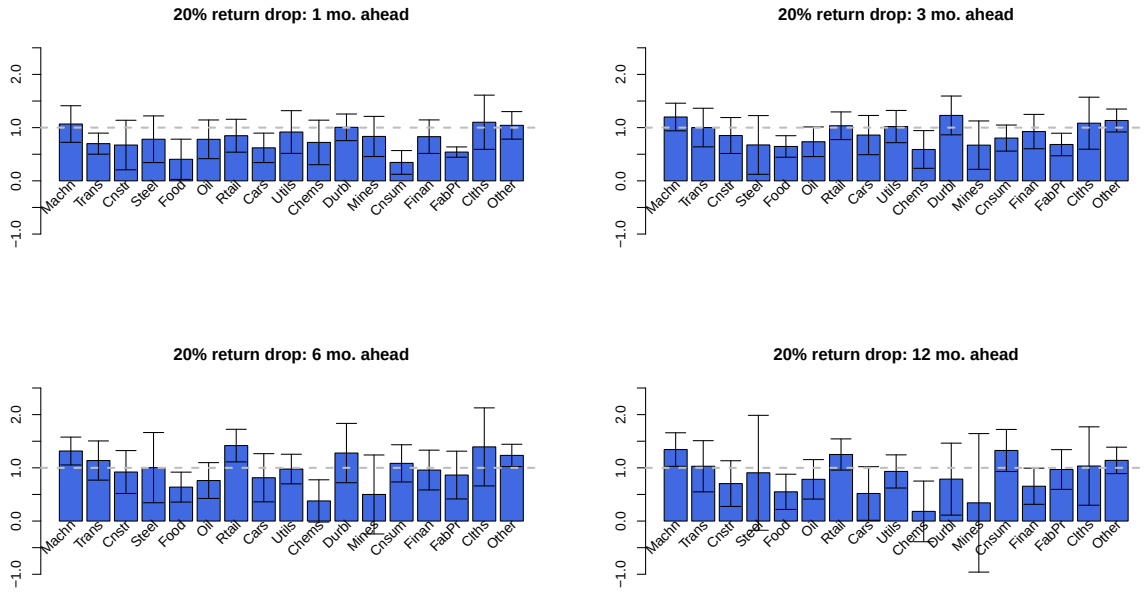


Figure 9: Regression coefficients β for the lower bounds by industry

This figure shows beta estimates for our baseline regression (9), using the lower bound as a predictor of 20% crashes over forecasting horizons of $\tau = 1, 3, 6$ and 12 months, within the 17 Fama–French industries. Error bars indicate 95% confidence intervals based on two-way clustered standard errors.

The three panels tell a coherent story. Related industries' crash probabilities comove, indicating that news about individual stocks' crash risk is not simply idiosyncratic, and there is substantial variation in crash probabilities over time and across industries.

6 Explaining the crash probabilities

As the empirical evidence suggests that we can treat the lower bound as a measure of the true forward-looking crash probability, we now use it to explore which characteristics help to explain the probability of a crash. To the extent that the lower bound is an accurate measure of the probability of a crash, it "de-noises" the realized crash event indicator, and this boosts the power of the exercise to detect variables that influence a stock's likelihood of crashing. We can therefore expect to find R^2 that are considerably higher than those documented in Table 4.

Defining a crash event as a net return of -20% or less over 1 month, we regress the lower bound on the 15 stock characteristics studied in Sections 3.3 and 4, normalized to

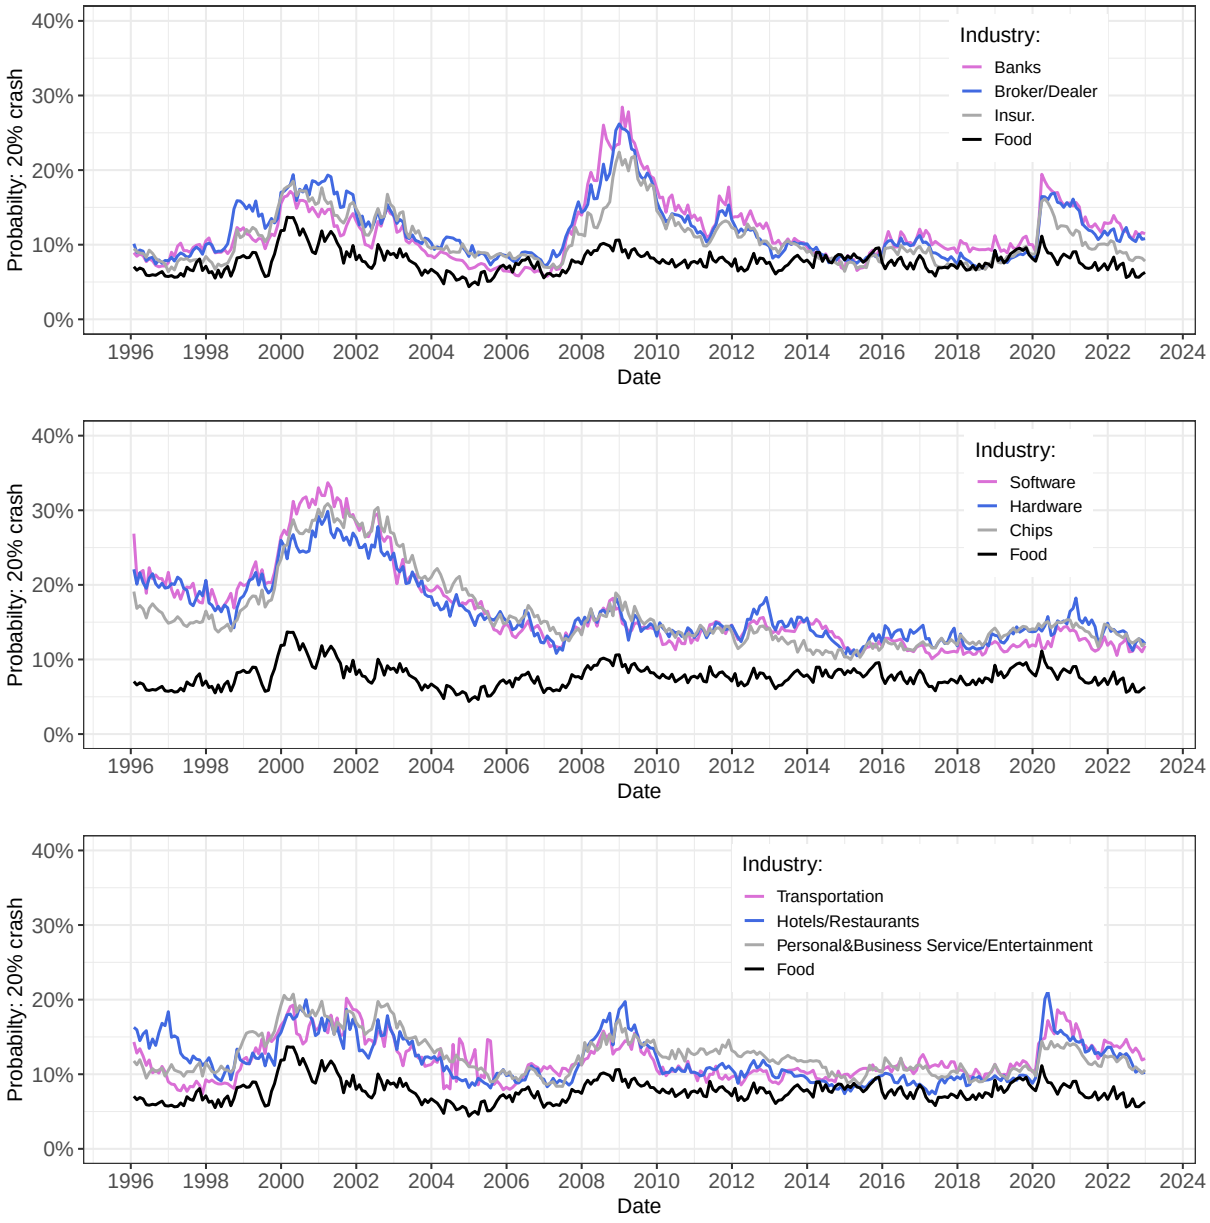


Figure 10: The average crash probabilities for stocks belonging to different industries

This figure presents the probabilities of 20% crashes in one year for three sets of industries. The crash probabilities are measured using the lower bounds, and then averaged for stocks being to the same industry. The industries are defined according to the 49 Fama–French industry classifications.

have zero mean and unit standard deviation. We estimate the regression using the full pooled sample, and also for each of the 17 Fama–French industry subsamples.

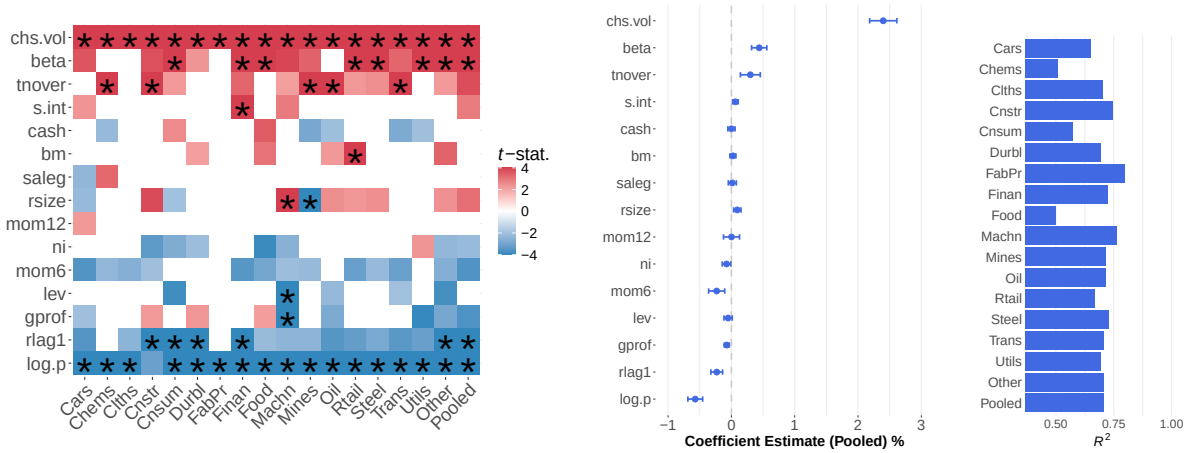
Panel A of Figure 11 shows a heatmap of t-statistics in industry-by-industry and pooled regressions (left), pooled coefficient estimates (middle), and R^2s (right). Several variables from the literature exhibit a consistent relationship with crash probabilities across industries. In the pooled sample, stocks with higher beta, CHS volatility (chs.vol), share turnover (tnover), and short interest (s.int) tend to have higher crash probabilities, consistent with the work of Chen, Hong, and Stein (2001) and Hong and Stein (2003). (Recall that these variables do not predict actual realized crashes when controlling for the lower bound, as demonstrated in columns (3) and (6) of Table 4.) The thresholded log share price variable, log.p, of Campbell, Hilscher, and Szilagyi (2008) enters with a negative, and highly statistically significant, coefficient: penny stocks tend to have higher crash risk. Similarly, crash risk is typically elevated for stocks with poor recent returns, either over the past month (rlag1) or from month -6 to -1 (mom6).

The strong statistical significance of these characteristics reflects the point noted above, that the lower bound filters out the noise in realized crash event indicators. (In contrast, when we use realized crash indicators as the dependent variable, we do not achieve the same level of significance due to reduced statistical efficiency, as shown in column (1) of Table 4.) As also anticipated above, this means that the regression R^2s reported in Figure 11 are an order of magnitude larger than those in reported in Table 4.

Some variables have significant predictive power for crashes in certain industries but not in others. High short interest is strongly associated with higher crash probabilities in financial firms. In the machinery and equipment industry, and in other capital-intensive sectors such as oil, transport, and consumer goods, high leverage is associated with lower crash risk.

We also run regressions that include all 15 squared characteristics and 105 interaction terms in addition to the 15 characteristics, as in the kitchen sink approach of Section 3.3.1. The left and middle sections of Panel B of Figure 11 report the t-statistics and full-sample coefficient estimates for the linear terms in these regressions; the main patterns are broadly consistent with Panel A. The right section reports the resulting R^2s. Including the squared and interaction terms increases R^2 fairly substantially in some industries (for example, chemicals and food) but has a more marginal impact—increasing R^2 from around 70% to

Panel A: Regressions of the lower bound onto 15 characteristics



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

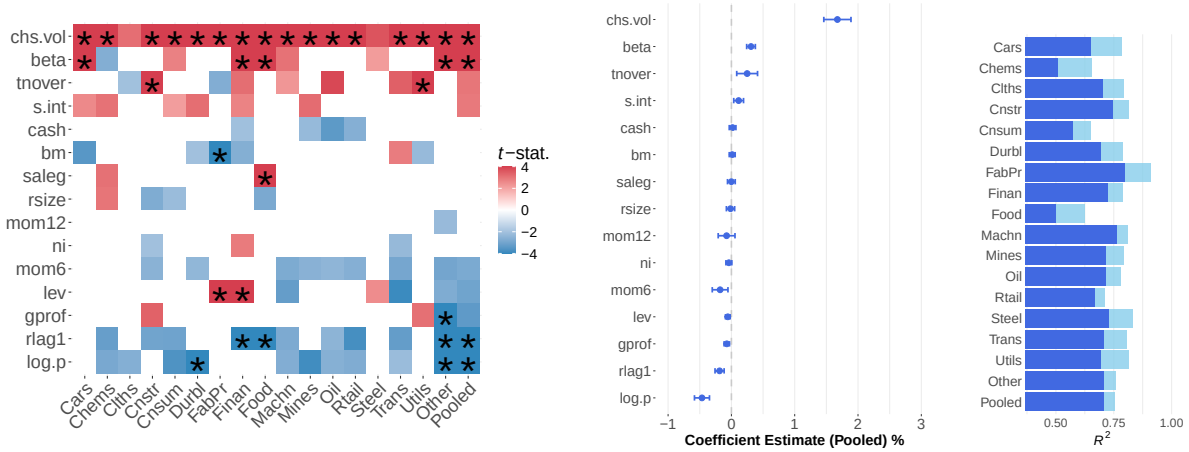


Figure 11: Explaining the crash probabilities

This figure presents the t-statistics, coefficient estimates with confidence intervals (for the entire sample), and R^2s from regressions of the lower bound—our proposed measure of expected crash probabilities—on stock characteristics. The standard errors used to calculate the t-statistics and confidence intervals are two-way clustered by firm and month. In Panel A, the independent variables are the 15 stock characteristics; in Panel B, 15 squared and 105 interaction terms are added (but only the results for individual variables are reported; the results for squared and interaction terms are reported in Figure A5 and Table A13 of the Appendix). In the heatmaps on the left, t-statistics between -2 and 2 are treated as zero, while those with absolute magnitudes above four are winsorized to plus/minus four—for better visualization—and marked with asterisks. Error bars in the middle plots indicate ± 2 standard errors. On the right, the barplots display R^2s; the increases in R^2s after including squared and interaction terms are shown in light blue in the lower panel.

around 75%—in the pooled regression.

Table A13, in the Appendix, reports coefficients and t-statistics for characteristics, squared characteristics, and interaction terms in the pooled regression. Here we briefly summarize some of the results, starting with the squared terms. Stocks with high CHS volatility, or with low lagged returns, profitability, or six-month momentum tend to have higher crash risk, as shown in Panel a; and the positive coefficients on the corresponding squared characteristics, reported in Panel b, indicate that these patterns become stronger as volatility rises or as lagged returns, profitability, or six-month momentum decline.

Turning to the interactions, reported in Panel c, we focus only on the most statistically significant coefficients with t-statistics above 4 in the interest of brevity.[16] The signs of the coefficients on cash×log.p, chs.vol×log.p, and r.size×log.p indicate that penny stocks that are small, or volatile, or that have low cash holdings, tend to be particularly exposed to crash risk. The positive coefficient on bm×gprof shows that an unprofitable stock has even higher crash risk if it also has a low book-to-market ratio. The negative coefficient on s.int×ni indicates high crash risk for stocks with both high short interest and low net income. Lastly, the positive sign on chs.vol×beta shows that the effects of CHS volatility and beta—each of which is individually linked to higher crash risk—are mutually reinforcing, so that volatile high-beta stocks have higher crash risk than would be suggested by a linear specification.

7 Conclusion

We introduce a new, theoretically motivated, forecasting variable that successfully predicts crashes in individual stocks by exploiting information in option prices. We do so as part of a more general framework that supplies bounds on the expectation of a general function of the market return and of an individual asset return.

We could, of course, have used option prices in a straightforward way to calculate *risk-neutral* probabilities of crashes. This approach is widely used by practitioners, and it has an appealing simplicity. Moreover, we find that the risk-neutral probability is a highly significant forecaster of crashes in full-sample tests.

[16] Figure A5, in the Appendix, plots the t-statistics for *all* squared and interaction terms in industry-specific and pooled regressions.

But standard theory suggests that risk-neutral probabilities should overstate the true probability of a crash, and should do so particularly at times of crisis when levels of risk, or of risk aversion, are particularly high. We find that this is the case in the data: the estimated coefficient on the risk-neutral probability is well below one, and at horizons above one month the coefficient is insignificant or even negative in the subperiods containing the subprime crisis and the Covid episode.

To move from risk-neutral probabilities to the true probabilities in which we are ultimately interested, we rely on an assumption on the form of the SDF that establishes a strong link between the true probabilities and option prices. Even after making this assumption, however, we face a further problem: to use option prices to measure the true probability that a given stock crashes, we need to understand the joint risk-neutral distribution of that stock and the market. But the prices we observe—of options on the market and of options on individual stocks—only reveal the *univariate* risk-neutral distributions of the market and of the individual stocks. We solve this problem with the final theoretical ingredient of the paper, the Fréchet–Hoeffding theorem, which places bounds on the relationship between the joint distribution and the marginal distributions that are tighter than those derived from the Cauchy–Schwarz inequality. We use the theorem to derive upper and lower bounds on the probability that an individual stock crashes, and show theoretically (and confirm empirically) that the upper and lower bounds are, respectively, higher and lower than the risk-neutral probability that the given stock crashes. We argue that the lower bound is likely to be closer to the truth, and find, empirically, that the lower bound is a highly statistically significant forecaster of crashes across crash sizes and forecasting horizons—even at horizons as short as one month—and across industries and subperiods. We find, moreover, that the estimated coefficient is close to one, as it should be if the lower bound is a good proxy for the true crash probability.

We compare the forecasting performance of the lower bound with a range of models based on 15 stock characteristics suggested by the prior literature. When the characteristics are included in multivariate regressions, the lower bound remains a highly statistically significant forecaster of crashes at all horizons. At the one month horizon, it drives out 14 of the 15 characteristics, and all 15 characteristics together contribute almost no incremental R^2 relative to the lower bound on its own. The one characteristic that is not driven out is a measure of short interest. But even short interest's economic importance

is limited by comparison with the lower bound: a one standard deviation increase in the lower bound raises the forecast crash probability by 3 percentage points, whereas a one standard deviation move in short interest raises the forecast probability by 0.3 percentage points.

The lower bound performs similarly well out of sample. We compare it to elastic net, ridge, and Lasso-regularized models that are allowed to exploit all 15 characteristics together with the risk-neutral probabilities and even the lower bound itself; and a "kitchen sink" approach that additionally includes squared characteristics and interactions of the characteristics. These models have considerable flexibility: we fit them to the data using both expanding-window and rolling regressions. In sharp contrast, the lower bound has no free parameters at all, as we use it without an intercept and with coefficient fixed at 1. And yet the lower bound outperforms at all horizons, for all three regularization methods, and for both the expanding-window and rolling regression approaches. The competitor models are led astray, in particular, by the financial crisis, which leads them to put too much weight on the risk-neutral probabilities, and hence to "cry wolf" too often in the aftermath of the crisis.

References

Aït-Sahalia, Y. and A. W. Lo (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance 53*(2), 499–547.

Ang, A., J. Chen, and Y. Xing (2006). Downside risk. *Review of Financial Studies 19*(4), 1191–1239.

Asquith, P., P. A. Pathak, and J. R. Ritter (2005). Short interest, institutional ownership, and stock returns. *Journal of Financial Economics 78*(2), 243–276.

Back, K., K. Crotty, and S. M. Kazempour (2022). Validity, tightness, and forecasting power of risk premium bounds. *Journal of Financial Economics 144*, 732–760.

Backus, D., M. Chernov, and I. Martin (2011). Disasters implied by equity index options. *Journal of Finance 66*(6), 1969–2012.

Baron, M., E. Verner, and W. Xiong (2021). Banking Crises Without Panics. *Quarterly Journal of Economics 136*(1), 51–113.

Barro, R. J. (2006). Rare disasters and asset markets in the 20th century. *Quarterly Journal of Economics 121*, 823–866.

Barro, R. J. and G. Y. Liao (2021). Rare disaster probability and options pricing. *Journal of Financial Economics 139*, 750–769.

Bates, D. S. (1991). The crash of '87: Was it expected? The evidence from options markets. *Journal of Finance 46*(3), 1009–1044.

Boyer, B., T. Mitton, and K. Vorkink (2009). Expected idiosyncratic skewness. *Review of Financial Studies 23*(1), 169–202.

Breeden, D. T. and R. H. Litzenberger (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business 51*(4), 621–651.

Campbell, J. Y., J. Hilscher, and J. Szilagyi (2008). In search of distress risk. *Journal of Finance 63*(6), 2899–2939.

Campbell, J. Y., M. Lettau, B. G. Malkiel, and Y. Xu (2001). Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. *Journal of Finance 56*(1), 1–43.

Carr, P. and D. Madan (2001). Towards a theory of volatility trading. In E. Jouini, J. Cvitanić, and M. Musiela (Eds.), *Handbooks in Mathematical Finance: Option Pricing, Interest Rates and Risk Management*, pp. 458–476. Cambridge University Press.

Carr, P. and L. Wu (2009). Variance risk premiums. *Review of Financial Studies 22*(3), 1311–1341.

Chabi-Yo, F., C. Dim, and G. Vilkov (2023). Generalized bounds on the conditional expected excess return on individual stocks. *Management Science 69*(2), 922–939.

Chabi-Yo, F. and J. Loudis (2020). The conditional expected market return. *Journal of Financial Economics 137*(3), 752–786.

Chen, J., H. Hong, and J. C. Stein (2001). Forecasting crashes: Trading volume, past returns, and conditional skewness in stock prices. *Journal of Financial Economics 61*(3), 345–381.

Cheung, K. C. (2009). Upper comonotonicity. *Insurance: Mathematics and Economics 45*(1), 35–40.

Christoffersen, P., M. Fournier, and K. Jacobs (2018). The factor structure in equity options. *Review of Financial Studies 31*(2), 595–637.

Christoffersen, P., K. Jacobs, and B. Y. Chang (2013). Forecasting with option-implied information. In G. Elliott and A. Timmermann (Eds.), *Handbook of Economic Forecasting*, Volume 2, pp. 581–656. Elsevier.

Daniel, K., A. Klos, and S. Rottke (2023). The dynamics of disagreement. *Review of Financial Studies 36*(6), 2431–2467.

Della Corte, P., C. Gao, and A. Jeanneret (2023). Expected currency returns and term structure of risk preferences. Working paper.

Diebold, F. X. and R. S. Mariano (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics 13*, 253–265.

Driscoll, J. C. and A. C. Kraay (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics 80*(4), 549–560.

Dybvig, P. H. (1988). Distributional analysis of portfolio choice. *Journal of Business*, 369–393.

Gandhi, M., N. J. Gormsen, and E. Lazarus (2022). Does the market understand time variation in the equity premium? Working paper.

Gao, C. and I. W. R. Martin (2021). Volatility, valuation ratios, and bubbles: An empirical measure of market sentiment. *Journal of Finance 76*(6), 3211–3254.

Gao, X. and J. R. Ritter (2010). The marketing of seasoned equity offerings. *Journal of Financial Economics 97*(1), 33–52.

Goetzmann, W. N., D. Kim, and R. J. Shiller (2022). Crash narratives. NBER Working paper 30195.

Greenwood, R., A. Shleifer, and Y. You (2019). Bubbles for Fama. *Journal of Financial Economics 131*(1), 20–43.

Harvey, C. R., Y. Liu, and H. Zhu (2016). ... and the cross-section of expected returns. *Review of Financial Studies 29*(1), 5–68.

Heckman, J. J., J. Smith, and N. Clements (1997). Making The Most Out Of Programme Evaluations and Social Experiments: Accounting For Heterogeneity in Programme Impacts. *Review of Economic Studies 64*(4), 487–535.

Hofer, M. and M. R. Iacò (2014). Optimal bounds for integrals with respect to copulas and applications. *Journal of Optimization Theory and Applications 3*, 999–1011.

Hong, H. and J. C. Stein (2003). Differences of opinion, short-sales constraints, and market crashes. *Review of Financial Studies 16*(2), 487–525.

Jackwerth, J. C. and M. Rubinstein (1996). Recovering probability distributions from option prices. *Journal of Finance 51*(5), 1611–1631.

Kadan, O. and X. Tang (2020). A bound on expected stock returns. *Review of Financial Studies 33*(4), 1565–1617.

Kelly, B. and H. Jiang (2014). Tail risk and asset prices. *Review of Financial Studies 27*(10), 2841–2871.

Kelly, B., H. Lustig, and S. Van Nieuwerburgh (2016). Too-systemic-to-fail: What option markets imply about sector-wide government guarantees. *American Economic Review 106*(6), 1278–1319.

Kremens, L. and I. W. R. Martin (2019). The quanto theory of exchange rates. *American Economic Review 109*(3), 810–843.

Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics 2*(1-2), 83–97.

Lewellen, J. and J. Shanken (2002). Learning, asset-pricing tests, and market efficiency. *Journal of Finance 57*(3), 1113–1145.

Manski, C. F. (1997). The mixing problem in programme evaluation. *Review of Economic Studies 64*, 537–553.

Martin, I. W. R. (2013). The Lucas orchard. *Econometrica 81*(1), 55–111.

Martin, I. W. R. (2017). What is the expected return on the market? *Quarterly Journal of Economics 132*(1), 367–433.

Martin, I. W. R. (2018). Options and the gamma knife. *Journal of Portfolio Management 44*(6), 47–55.

Martin, I. W. R. (2021). On the autocorrelation of the stock market. *Journal of Financial Econometrics 19*(1), 39–52.

Martin, I. W. R. and S. Nagel (2022). Market efficiency in the age of big data. *Journal of Financial Economics 145*(1), 154–177.

Martin, I. W. R. and C. Wagner (2019). What is the expected return on a stock? *Journal of Finance 74*(4), 1887–1929.

Pederzoli, P. (2021). Skewness swaps: Evidence from the cross-section of stocks. Working paper.

Ross, S. A. (1976). Options and efficiency. *Quarterly Journal of Economics 90*(1), 75–89.

Rubinstein, M. (1994). Implied binomial trees. *Journal of Finance 49*(3), 771–818.

Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris 8*, 229–231.

Tchen, A. H. (1980). Inequalities for distributions with given marginals. *Annals of Probability 8*(4), 814–827.

Thompson, S. B. (2011). Simple formulas for standard errors that cluster by both firm and time. *Journal of Financial Economics 99*(1), 1–10.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B (Statistical Methodology) 58*(1), 267–288.

Timmermann, A. G. (1993). How learning in financial markets generates excess volatility and predictability in stock prices. *Quarterly Journal of Economics 108*(4), 1135–1145.

Vilkov, G. and Y. Xiao (2013). Option-implied information and predictability of extreme returns. Working paper.

Welch, I. and A. Goyal (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies 21*(4), 1455–1508.

Zou, H. and T. Hastie (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B (Statistical Methodology) 67*(2), 301–320.

Appendix A Proofs

A.1 Proof of Result 1

Proof. When $g(x, y)$, defined on $[0, \infty) \times [0, \infty)$, is continuous and two-increasing, $k(u, v) = g\left(Q_m^{-1}(u),\, Q_i^{-1}(v)\right)$ is two-increasing in $[0, 1] \times [0, 1]$. We therefore have

$$\inf_{C \in \mathcal{C}} \int_{[0,1]^2} k\,(u, v)\ \mathrm{d}C(u, v) \leq \int g\,(x, y)\ \mathrm{d}Q_{mi}(x, y) \leq \sup_{C \in \mathcal{C}} \int_{[0,1]^2} k\,(u, v)\ \mathrm{d}C(u, v), \quad \text{(A.1)}$$

where we write $\mathcal{C}$ for the set of all two-dimensional copulas, and

$$k(u, v) = g\left(Q_m^{-1}(u),\, Q_i^{-1}(v)\right). \quad \text{(A.2)}$$

From Corollary 2.2 of Tchen (1980), we have

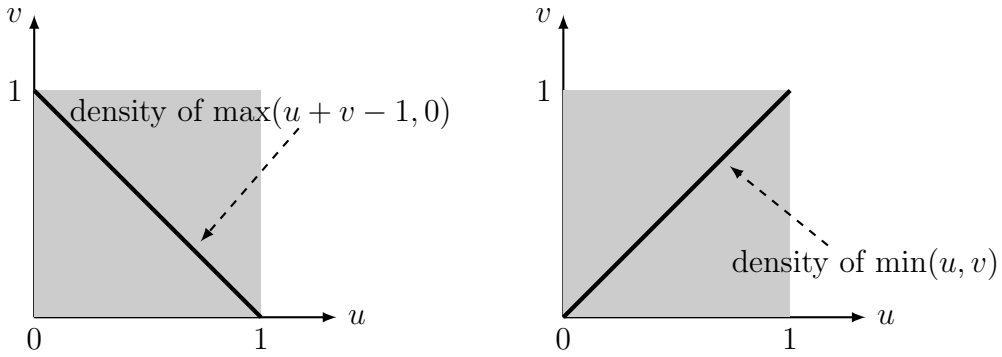
$$\inf_{C \in \mathcal{C}} \int_{[0,1]^2} k(u, v)\, \mathrm{d}C(u, v) = \int_{[0,1]^2} k(u, v)\, \mathrm{d}\left(\max(u + v - 1,\ 0)\right),$$

and

$$\sup_{C \in \mathcal{C}} \int_{[0,1]^2} k(u, v)\, \mathrm{d}C(u, v) = \int_{[0,1]^2} k(u, v)\, \mathrm{d}\left(\min(u, v)\right).$$

The probability densities of the Fréchet–Hoeffding lower bound, $\max(u + v - 1,\ 0)$, and the Fréchet–Hoeffding upper bound, $\min(u, v)$, are uniformly distributed along the two

diagonals of the square $[0,1]^2$ in $\mathbb{R}^2$, illustrated as follows:



Integrating the right-hand sides of the two equations above (with regard to these two densities), we have

$$\int_{[0,1]^2} k(u,v)\,\mathrm{d}\left(\max(u+v-1,\ 0)\right) = \int_0^1 k(u,1-u)\,\mathrm{d}u$$

and

$$\int_{[0,1]^2} k(u,v)\,\mathrm{d}\left(\min(u,v)\right) = \int_0^1 k(u,u)\,\mathrm{d}u.$$

Substituting these expressions back into (A.1) and using the definition (A.2) of $k(u,v)$, it follows that

$$\int_0^1 g\left(Q_m^{-1}(u),\, Q_i^{-1}(1-u)\right)\,\mathrm{d}u \le \int g\left(x,y\right)\,\mathrm{d}Q_{mi}(x,y) \le \int_0^1 g\left(Q_m^{-1}(u),\, Q_i^{-1}(u)\right)\,\mathrm{d}u\,. \tag{A.3}$$

The result follows on making the change of variable $R_m = Q_m^{-1}(u)$ in the left- and right-most integrals. □

A.2 Proof of Result 2

Proof. Under the stated assumptions, the function $g(x,y) = x^\gamma h(y)$ is continuous and two-increasing, as the product of two nondecreasing functions is two-increasing. From equation (A.3), we have

$$\int_0^1 \left[Q_m^{-1}(u)\right]^\gamma h\left(Q_i^{-1}(1-u)\right)\,\mathrm{d}u \le \mathbb{E}^*[R_m^\gamma h(R_i)] \le \int_0^1 \left[Q_m^{-1}(u)\right]^\gamma h\left(Q_i^{-1}(u)\right)\,\mathrm{d}u.$$

But, making the change of variables $R_m = Q_m^{-1}(u)$, we have

$$\int_0^1 \left[Q_m^{-1}(u)\right]^\gamma h\left(Q_i^{-1}(1-u)\right)\,\mathrm{d}u = \int_0^\infty R_m^\gamma h(Q_i^{-1}(1-Q_m(R_m)))\,\mathrm{d}Q_m(R_m)$$

and

$$\int_0^1 \left[Q_m^{-1}(u)\right]^\gamma h\left(Q_i^{-1}(u)\right)\,\mathrm{d}u = \int_0^\infty R_m^\gamma h(Q_i^{-1}(Q_m(R_m)))\,\mathrm{d}Q_m(R_m).$$

Combining these facts, we have the bounds stated in the result.

The lower bound is achieved when the copula linking Q_m and Q_i is $C(u,v) = \max(u+v-1,\,0)$. This implies that $Q_m(R_m)+Q_i(R_i) \equiv 1$. Similarly, the upper bound is achieved when the copula is $C(u,v) = \min(u,v)$, that is, when $Q_i(R_i) = Q_m(R_m)$. □

A.3 Proof of Result 3

Proof. Define $h_n(R_i) = -1/(1+e^{n(R_i-q)})$, which converges pointwise to $h(R_i) = -\boldsymbol{I}(R_i \leq q)$ almost everywhere. By equation (4) and the dominated convergence theorem ($|h_n| \leq 1$),

$$\mathbb{P}[R_i \leq q] = -\mathbb{E}[h(R_i)] = -\frac{\mathbb{E}^*[R_m^\gamma h(R_i)]}{\mathbb{E}^*[R_m^\gamma]} = -\lim_{n\to\infty}\frac{\mathbb{E}^*[R_m^\gamma h_n(R_i)]}{\mathbb{E}^*[R_m^\gamma]}.$$

Since h_n is continuous and increasing, by Result 2,

$$\mathbb{E}^*[R_m^\gamma h_n(Q_i^{-1}(1-Q_m(R_m))] \leq \mathbb{E}^*[R_m^\gamma h_n(R_i)] \leq \mathbb{E}^*[R_m^\gamma h_n(Q_i^{-1}(Q_m(R_m))]. \tag{A.4}$$

Taking limits and applying the dominated convergence theorem again,

$$-\frac{\mathbb{E}^*[R_m^\gamma h(Q_i^{-1}(Q_m(R_m))]}{\mathbb{E}^*[R_m^\gamma]} \leq \mathbb{P}[R_i \leq q] \leq -\frac{\mathbb{E}^*[R_m^\gamma h(Q_i^{-1}(1-Q_m(R_m))]}{\mathbb{E}^*[R_m^\gamma]}. \tag{A.5}$$

Simplifying terms, we have

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

The upper bound is attained when R_i and R_m are countermonotonic according to Result 2, as the upper bound in (A.5) is the negative of the lower bound in (A.4).

Similarly, the lower bound is attained when R_i and R_m are comonotonic, which means

that $Q_i(R_i) = Q_m(R_m) = U \sim \text{Uniform}(0,1)$ almost everywhere.

Now, if R_i and R_m are lower comonotonic with thresholds (q, q_l), Proposition 5 of Cheung (2009) implies that the joint distribution of $Q_i(R_i)$ and $Q_m(R_m)$ has the following distributional representation:

$$Q_i(R_i) = U \times \boldsymbol{I}(U \leq p^*) + [1 - (1-p^*)V_i] \times \boldsymbol{I}(U > p^*),$$

$$Q_m(R_m) = U \times \boldsymbol{I}(U \leq p^*) + [1 - (1-p^*)V_m] \times \boldsymbol{I}(U > p^*),$$

where $p^* = Q_i(q) = Q_m(q_l)$ by definition, and $U, V_i, V_m \overset{\text{iid}}{\sim} \text{Uniform}(0,1)$. (In the case in which $V_i = V_m = U$, R_i and R_m are comonotonic.) It is easy to verify that $Q_i(R_i) = Q_m(R_m) = U$ when $R_i \leq q$ and the lower bound can still hold as equality.

To see that the risk-neutral crash probability lies between the lower and upper bounds, note that by the continuous version of Chebyshev's sum inequality,[17] we have

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

$$\frac{\mathbb{E}^*\left[R_m^\gamma \boldsymbol{I}\left(R_m \geq q_u\right)\right]}{\mathbb{E}^*\left[R_m^\gamma\right]} \geq \frac{\mathbb{E}^*\left[R_m^\gamma\right]\mathbb{E}^*\left[\boldsymbol{I}\left(R_m \geq q_u\right)\right]}{\mathbb{E}^*\left[R_m^\gamma\right]} = 1 - Q_m(q_u) = Q_i(q) = \mathbb{P}^*[R_i \leq q]. \quad \square$$

A.4 Proof of Result 4

Proof. When $\gamma = 0$, the bounds become $\mathbb{P}^*[R_m \leq q_l] \leq \mathbb{P}[R_i \leq q] \leq \mathbb{P}^*[R_m \geq q_u]$. By definition, both the lower and upper bounds equal $\mathbb{P}^*[R_i \leq q]$.

To show that the lower bound is decreasing in γ, define the (decreasing) function $\psi(x) = \boldsymbol{I}(Q_m(x) \leq Q_i(q))$. The lower bound is then $\mathbb{E}^*[R_m^\gamma \psi(R_m)]/\mathbb{E}^*[R_m^\gamma]$ and

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\left\{\frac{\mathbb{E}^*[R_m^\gamma \psi(R_m)]}{\mathbb{E}^*[R_m^\gamma]}\right\} = \frac{\mathbb{E}^*[R_m^\gamma \log(R_m)\psi(R_m)]\mathbb{E}^*[R_m^\gamma] - \mathbb{E}^*[R_m^\gamma \psi(R_m)]\mathbb{E}^*[R_m^\gamma \log(R_m)]}{\{\mathbb{E}^*[R_m^\gamma]\}^2}.$$

[17] This inequality states that for functions f and g which are integrable over $[0,1]$, both non-increasing or both non-decreasing, then $\int_0^1 f(x)g(x)\,\mathrm{d}x \geq \int_0^1 f(x)\,\mathrm{d}x \int_0^1 g(x)\,\mathrm{d}x$. If one function is non-increasing and the other is non-decreasing, the inequality is reversed. Letting $f(x) = [Q_m^{-1}(x)]^\gamma$ (a non-decreasing function of x), we derive the first inequality by setting $g(x) = \boldsymbol{I}(x \leq Q_i(q))$ and the second by setting $g(x) = \boldsymbol{I}(x \geq Q_i(q))$.

The numerator of the right-hand side equals

$$\begin{aligned}
&\int_0^\infty x^\gamma \log(x)\psi(x)\,\mathrm{d}Q_m(x)\int_0^\infty y^\gamma\,\mathrm{d}Q_m(y) - \int_0^\infty x^\gamma\psi(x)\,\mathrm{d}Q_m(x)\int_0^\infty y^\gamma\log(y)\,\mathrm{d}Q_m(y)\\
&= \iint_{x\geq 0,\, y\geq 0} [x^\gamma\log(x)\psi(x)y^\gamma - x^\gamma\psi(x)y^\gamma\log(y)]\,\mathrm{d}Q_m(x)\mathrm{d}Q_m(y)\\
&= \iint_{x\geq y\geq 0} x^\gamma y^\gamma\psi(x)\log\left(\frac{x}{y}\right)\,\mathrm{d}Q_m(x)\,\mathrm{d}Q_m(y) + \iint_{0\leq x\leq y} x^\gamma y^\gamma\psi(x)\log\left(\frac{x}{y}\right)\,\mathrm{d}Q_m(x)\,\mathrm{d}Q_m(y).
\end{aligned}$$

Call the last term above N. Noticing that when $x \geq y$, $\psi(x) \leq \psi(y)$, we have

$$\begin{aligned}
N &\leq \iint_{x\geq y\geq 0} x^\gamma y^\gamma\psi(y)\log\left(\frac{x}{y}\right)\,\mathrm{d}Q_m(x)\,\mathrm{d}Q_m(y) + \iint_{0\leq x\leq y} x^\gamma y^\gamma\psi(x)\log\left(\frac{x}{y}\right)\,\mathrm{d}Q_m(x)\,\mathrm{d}Q_m(y)\\
&= \iint_{0\leq x\leq y} x^\gamma y^\gamma\psi(x)\log\left(\frac{y}{x}\right)\,\mathrm{d}Q_m(x)\,\mathrm{d}Q_m(y) + \iint_{0\leq x\leq y} x^\gamma y^\gamma\psi(x)\log\left(\frac{x}{y}\right)\,\mathrm{d}Q_m(x)\,\mathrm{d}Q_m(y)\\
&= 0.
\end{aligned}$$

As a result, the lower bound as a function of γ is decreasing. Applying the same logic to the increasing function $\psi(x) = \boldsymbol{I}(Q_m(x) \geq 1 - Q_i(q))$, we can conclude that the upper bound is increasing with regard to γ.

Next, note that the lower bound is such that

$$\frac{\mathbb{E}^*\left[R_m^\gamma \boldsymbol{I}\left(R_m \leq q_l\right)\right]}{\mathbb{E}^*\left[R_m^\gamma\right]} \leq \frac{q_l^\gamma}{\mathbb{E}^*[R_m^\gamma]}.$$

To show that the lower bound converges to zero as $\gamma \to \infty$, we must show that $\mathbb{E}^*[(R_m/q_l)^\gamma] \to \infty$ as $\gamma \to \infty$. This holds if $\mathbb{P}^*[R_m/q_l > 1] > 0$. If this condition does not hold, $R_m \leq q_l = Q_m^{-1}(Q_i(q))$ with probability one, which violates the assumption that $Q_i(q) < 1$. Thus, $\mathbb{E}^*[(R_m/q_l)^\gamma] \to \infty$ and the lower bound converges to zero as $\gamma \to \infty$.

To show that the upper bound goes to one as $\gamma \to \infty$, note that

$$1 = \frac{\mathbb{E}^*[R_m^\gamma \boldsymbol{I}(R_m < q_u)] + \mathbb{E}^*[R_m^\gamma \boldsymbol{I}(R_m \geq q_u)]}{\mathbb{E}^*[R_m^\gamma]}.$$

The result follows because $\mathbb{E}^*[R_m^\gamma \boldsymbol{I}(R_m < q_u)]/\mathbb{E}^*[R_m^\gamma] \to 0$ as $\gamma \to \infty$. This is satisfied because $\mathbb{P}^*[R_m < q_u] < 1$. (If not, we would have $R_m \leq q_u = Q_m^{-1}(1 - Q_i(q))$, and hence $1 - Q_i(q) = 1$; but this violates the assumption that $Q_i(q) > 0$.) □

A.5 Proof of Result 5

Proof. By the Carr–Madan formula (Carr and Madan, 2001), for any smooth function $g(\cdot)$ we have

$$g(S) = g(F)+g'(F)(S-F)+\int_0^F g''(K)\max\{K-S,\,0\}\,\mathrm{d}K+\int_F^\infty g''(K)\max\{S-K,\,0\}\,\mathrm{d}K.$$

Let S_0 and $F = S_0R_f$ be the spot and forward level of the market index, the function $g(S)$ be S^γ. Treating S, a random variable, as the level of market index next period, taking the risk-neutral expectations on both sides of the equation above (changing orders of integrals when needed), we have

$$\begin{aligned}\mathbb{E}^*\left[S^\gamma\right] =& S_0^\gamma R_f^\gamma + \gamma S_0^{\gamma-1}R_f^{\gamma-1}(\mathbb{E}^*[S]-F)\\ &+\int_0^F \gamma(\gamma-1)K^{\gamma-2}R_f\mathrm{put}(K)\,\mathrm{d}K+\int_F^\infty \gamma(\gamma-1)K^{\gamma-2}R_f\mathrm{call}(K)\,\mathrm{d}K.\end{aligned}$$

Dividing both sides by S_0^γ and noticing that $R_m = S/S_0$ and that $\mathbb{E}^*[S] = F$, we have the first equation.

Next, noticing that

$$\begin{aligned}\mathbb{E}^*\left[R_m^\gamma \boldsymbol{I}(R_m\le q_l)\right] &= \frac{\mathbb{E}^*\left[S^\gamma \boldsymbol{I}(S\le K_l)\right]}{S_0^\gamma} = \frac{R_f}{S_0^\gamma}\int_0^{K_l} K^\gamma \mathrm{put}''(K)\,\mathrm{d}K,\\ \mathbb{E}^*\left[R_m^\gamma \boldsymbol{I}(R_m\ge q_u)\right] &= \frac{\mathbb{E}^*\left[S^\gamma \boldsymbol{I}(S\ge K_u)\right]}{S_0^\gamma} = \frac{R_f}{S_0^\gamma}\int_{K_u}^\infty K^\gamma \mathrm{call}''(K)\,\mathrm{d}K,\end{aligned}$$

where the second equations follow from the static replication logic of Breeden and Litzenberger (1978). Integrating the last integrals above by parts and using the facts that $\mathrm{put}(0) = \mathrm{put}'(0) = 0$ and that $\mathrm{call}(\infty) = \mathrm{call}'(\infty) = 0$, we have

$$\begin{aligned}\int_0^{K_l} K^\gamma \mathrm{put}''(K)\,\mathrm{d}K &= K^\gamma \mathrm{put}'(K)\Big|_0^{K_l} - \int_0^{K_l}\gamma K^{\gamma-1}\mathrm{put}'(K)\,\mathrm{d}K\\ &= K_l^\gamma \mathrm{put}'(K_l) - \left(\gamma K^{\gamma-1}\mathrm{put}(K)\Big|_0^{K_l} - \int_0^{K_l}\gamma(\gamma-1)K^{\gamma-2}\mathrm{put}(K)\,\mathrm{d}K\right)\\ &= K_l^\gamma \mathrm{put}'(K_l) - \left(\gamma K_l^{\gamma-1}\mathrm{put}(K_l) - \int_0^{K_l}\gamma(\gamma-1)K^{\gamma-2}\mathrm{put}(K)\,\mathrm{d}K\right)\end{aligned}$$

and

$$\begin{aligned}\int_{K_u}^{\infty} K^{\gamma}\text{call}''(K)\,\mathrm{d}K &= K^{\gamma}\text{call}'(K)\Big|_{K_u}^{\infty} - \int_{K_u}^{\infty} \gamma K^{\gamma-1}\text{call}'(K)\,\mathrm{d}K \\ &= -K_u^{\gamma}\text{call}'(K_u) - \left(\gamma K^{\gamma-1}\text{call}(K)\Big|_{K_u}^{\infty} - \int_{K_u}^{\infty} \gamma(\gamma-1)K^{\gamma-2}\text{call}(K)\,\mathrm{d}K\right) \\ &= -K_u^{\gamma}\text{call}'(K_u) + \left(\gamma K_u^{\gamma-1}\text{call}(K_u) + \int_{K_u}^{\infty} \gamma(\gamma-1)K^{\gamma-2}\text{call}(K)\,\mathrm{d}K\right).\end{aligned}$$

Multiplying these formulas by R_f/S_0^{γ} yields the second and the third equations. □

Appendix B Bounds for general contingent payoffs

Result 1, which underpins our empirical work, requires that the function $g(x,y) = x^{\gamma}h(y)$ is two-increasing. When this is not the case, we can modify our approach by exploiting a result of Hofer and Iacò (2014). Specifically, for any well-behaved function k, we can approximate the optimization over the functional space of all copula functions $\mathcal{C}$ as follows:

$$\max_{C\in\mathcal{C}} \int_{[0,1]^2} k\,(u,v)\ \mathrm{d}C(u,v) \approx \max_{\pi\in\mathcal{P}_n}\ \frac{1}{n}\sum_{i=1}^{n} k\left(\frac{i}{n}, \frac{\pi(i)}{n}\right), \tag{A.6}$$

where $\mathcal{P}_n$ is the set of permutations of $\{1,\ldots,n\}$, and the approximate equality can be made to hold up to arbitrarily small error by choosing n sufficiently large. The two conditions for k to be well-behaved are that (i) it must be such that the integral is finite and (ii) it must be Lipschitz continuous almost everywhere.

The right-hand side of (A.6) is the canonical linear assignment problem in combinatorial optimization. The so-called Hungarian algorithm (Kuhn, 1955) reduces the complexity of solving this problem from $O(n!)$ (based on brute-force search) to $O(n^3)$, and offers the upper bound. Similarly, to obtain lower bounds, we can apply the Hungarian algorithm to the integral involving $-k(u,v)$. Using this approach, together with Sklar's theorem, we have the following result.

Result 6. *Let h be Lipschitz continuous almost everywhere. If $\pi_{\min}$ is a permutation of $\{1,\ldots,n\}$ that minimizes $\sum_{k=1}^{n} k\left(\frac{i}{n}, \frac{\pi(i)}{n}\right)$ and $\pi_{\max}$ is a permutation that maximizes $\sum_{k=1}^{n} k\left(\frac{i}{n}, \frac{\pi(i)}{n}\right)$, then (up to errors that can be made arbitrarily small by choosing n*

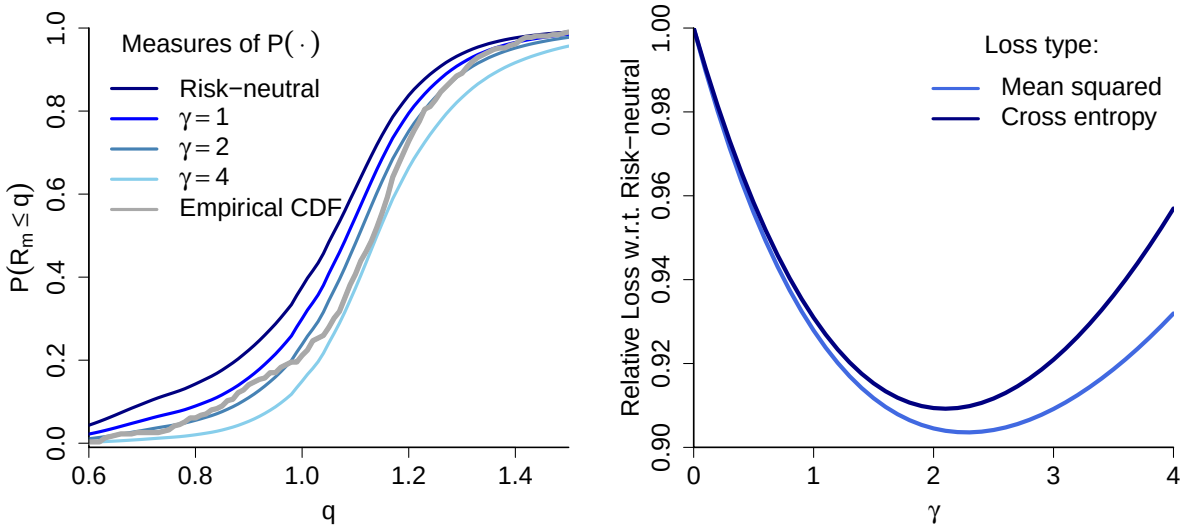


Figure A1: The probability distribution of market returns: the left panel compares the empirical CDFs of one-year-ahead market returns and our theory-based distribution forecasts, recovered from option prices; the right panel shows two loss functions measuring forecasting errors for when the crash probabilities are calculated using different γs.

sufficiently large)

$$\frac{1}{nC}\sum_{i=1}^{n} k\left(\frac{i}{n}, \frac{\pi_{\min(i)}}{n}\right) \leq \mathbb{E}[h(R_i)] \leq \frac{1}{nC}\sum_{i=1}^{n} k\left(\frac{i}{n}, \frac{\pi_{\max(i)}}{n}\right),$$

where the constant C *equals* $\int_0^1 [Q_m^{-1}(u)]^\gamma \, \mathrm{d}u$ *and* $k(u,v) = g\left(Q_m^{-1}(u), Q_i^{-1}(v)\right)$.

As a simple example, if we are interested in evaluating the probability that a stock's return lies in some interval, Result 6 can be applied with $h(R_i) = \boldsymbol{I}(q_1 \leq R_i \leq q_2)$.

Appendix C Calibrating risk aversion

We use equation (7) (in which the risk-neutral expectations can be evaluated using option prices according to Result 5) to compute the one-year-ahead crash probabilities $\mathbb{P}[R_m \leq q]$ for a range of values of q, and for risk-aversion parameters γ ranging from zero to four. We focus on the one-year horizon because of the limited number of market crash events over shorter horizons in our sample (Jan. 1996–Dec. 2022).

The left panel of Figure A1 shows the unconditional distributions of the gross S&P 500 return R_m recovered from index option prices assuming risk-neutrality ($\gamma = 0$) and

for $\gamma \in \{1, 2, 4\}$. The empirical CDF $(1/T)\sum_{t=0}^{T-1} \boldsymbol{I}(R_{m,t+1} \leq q)$ is also included. The plot confirms the intuition that the risk-neutral probabilities overstate the likelihood of market downturns.

The right panel of Figure A1 presents the loss functions for forecasting market "crashes" $y_{t+1} = \boldsymbol{I}(R_{m,t+1} \leq q)$ using their conditional expectations $p_t = \mathbb{E}_t[y_{t+1}] = \mathbb{P}_t[R_{m,t+1} \leq q]$ calculated based on equation (7). The mean-squared, $\sum_t (y_{t+1} - p_t)^2$, and cross-entropy, $\sum_t [-y_{t+1} \log p_t - (1 - y_{t+1}) \log(1 - p_t)]$, losses are both reported. All forecasting losses are normalized by dividing by the realized loss when the risk-neutral distribution is used to forecast. The best forecasting performance is achieved when γ is around two.

Appendix D Calculating the Option-Implied Bounds

Here we provide further implementation details on calculating the option-implied bounds.

Filtering. We applied four criteria to filter the implied volatility data in our sample: 1) spot prices must be available from the CRSP database; 2) strike prices must be positive; 3) the OptionMetrics dispersion variable, a goodness-of-fit measure for OptionMetrics' proprietary multinomial tree algorithm of constructing the volatility surface, must be smaller than 0.05 and greater than zero; 4) for any firm-month-maturity combination, the implied volatilities must be available at more than 10 different strike prices.

Interpolation and extrapolation. At time t, we denote by $\{\sigma_{it}(K_1, \tau), \ldots, \sigma_{it}(K_n, \tau)\}$ the Black-Scholes implied volatilities of firm i's options at strike prices $K_1 \leq \cdots \leq K_n$, maturing at time $(t+\tau)$. These are observable volatility surface data from OptionMetrics. We linearly interpolate implied volatility observations for any K such that $K_1 < K < K_n$. For strikes outside the observable range $[K_1, K_n]$, we extrapolate a flat volatility surface. That is, for $K \leq K_1$, we set $\sigma_{it}(K, \tau) = \sigma_{it}(K_1, \tau)$, and, for $K \geq K_n$, we set $\sigma_{it}(K, \tau) = \sigma_{it}(K_n, \tau)$.

Risk-free rates. All risk-free rates are sources from the OptionMetrics yield curve data. At time t, for maturities at which the risk-free rates are not directly observable, we use values linearly interpolated from the OptionMetrics yield curves.

The "clean" option prices. We construct option prices by applying the Black–Scholes formula for a given strike $K > 0$, maturity τ, implied volatility $\sigma_{it}(K, \tau)$, risk-free rate $r_{f,t}(\tau)$, and spot price S_{it}. These are European option prices assuming zero dividend yield.

We compute these prices on a grid of 2000 steps within the interval $K/S_{it} \in [1/L,\, L]$, where $L = 3$ for one-, three-, and six-month horizons and $L = 5$ for the one-year horizon. We only consider out-of-the-money options. That is, when $K \leq S_{it}R_{f,t}$, we compute put prices, where $R_{f,t} = \exp(r_{f,t}(\tau)\tau)$; when $K > S_{it}R_{f,t}$, we compute call prices.

The risk-neutral marginals. Given put and call option prices on the grid of strikes, we numerically compute the following gradients to recover the risk-neutral marginals:

$$Q_{it}\left(\frac{K}{S_0}\right) = \begin{cases} R_{f,t}\text{put}'_{it}(K), & K \leq R_{f,t}S_{it} \\ R_{f,t}\text{call}'_{it}(K) + 1, & K > R_{f,t}S_{it} \end{cases}.$$

Only out-of-the-money option prices (derived from the corresponding implied volatilities) are used to compute the risk-neutral marginals. We fit isotonic regressions to the raw option-implied risk-neutral CDFs to guarantee monotonicity. We then winsorize the fitted curve to ensure the CDFs are within [0, 1].

We use Result 3 to compute the quantiles q_l and q_u, which involves both the risk-neutral marginals for individual stocks and those for the market index. When applying Result 5 to compute the numerators and denominators of our bounds, we use the S&P 500 index option prices on the fine grid to numerically evaluate the integrals according to the midpoint rule.

Appendix E Forecasting Rallies

Here we state our result on bounding rally probabilities in parallel to Result 3 on crashes, prove this result, and tabulate our in-sample test results for forecasting rallies.

Result 7. *The probability of a rally in stock i, $\mathbb{P}[R_i \geq q]$, satisfies the bounds*

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

where $q_u = Q_m^{-1}(Q_i(q))$ and $q_l = Q_m^{-1}(1 - Q_i(q))$.

The upper bound is attained if the return on the market R_m and return on the stock R_i are upper comonotonic with thresholds (q_u, q); in particular this holds if R_m and R_i are comonotonic. The lower bound is attained if the two returns are countermonotonic.

Proof. Define $h_n(R_i) = 1/(1+e^{-n(R_i-q)})$, which converges pointwise to $h(R_i) = \boldsymbol{I}(R_i \geq q)$ almost everywhere. By equation (4) and the dominated convergence theorem ($|h_n| \leq 1$),

$$\mathbb{P}[R_i \geq q] = \mathbb{E}[h(R_i)] = \frac{\mathbb{E}^*[R_m^\gamma h(R_i)]}{\mathbb{E}^*[R_m^\gamma]} = \lim_{n\to\infty} \frac{\mathbb{E}^*[R_m^\gamma h_n(R_i)]}{\mathbb{E}^*[R_m^\gamma]}.$$

Since h_n is continuous and increasing, by Result 2,

$$\mathbb{E}^*[R_m^\gamma h_n(Q_i^{-1}(1-Q_m(R_m))] \leq \mathbb{E}^*[R_m^\gamma h_n(R_i)] \leq \mathbb{E}^*[R_m^\gamma h_n(Q_i^{-1}(Q_m(R_m))]. \tag{A.7}$$

Taking limits and applying the dominated convergence theorem again,

$$\frac{\mathbb{E}^*[R_m^\gamma h(Q_i^{-1}(1-Q_m(R_m))]}{\mathbb{E}^*[R_m^\gamma]} \leq \mathbb{P}[R_i \geq q] \leq \frac{\mathbb{E}^*[R_m^\gamma h(Q_i^{-1}(Q_m(R_m))]}{\mathbb{E}^*[R_m^\gamma]}. \tag{A.8}$$

Simplifying terms, we have

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

By Result 2, the lower and upper bounds are attained when R_i and R_m are, respectively, countermonotonic and comonotonic. In the latter case, $Q_i(R_i) = Q_m(R_m) = U \sim$ Uniform$(0,1)$ almost everywhere. Noticing that if R_i and R_m are upper comonotonic with thresholds (q, q_u), then their joint distribution can be characterized as

$$Q_i(R_i) = U \times \boldsymbol{I}(U \geq p^*) + p^* V_i \times \boldsymbol{I}(U < p^*),$$

$$Q_m(R_m) = U \times \boldsymbol{I}(U \geq p^*) + p^* V_m \times \boldsymbol{I}(U < p^*),$$

where $U, V_i, V_m \overset{\text{iid}}{\sim}$ Uniform$(0,1)$ and $p^* = Q_i(q)$. It follows that $Q_i(R_i) = Q_m(R_m) = U$ when $R_i \geq q$, so that equality still holds for the upper bound under the weaker assumption of upper comonotonicity. □

We expect the upper bound to be tighter than the lower bound, as stock returns are less likely to move in the opposite direction of the market. The summary statistics and OLS test results in Tables A1 and A2 confirm this point.

Table A1: Summary statistics (rally probabilities)

This table presents summary statistics of realized rally events, our rally probability bounds defined in Result 7, and risk-neutral crash probabilities. The sample data are monthly from January 1996 to December 2022. The rally events (realized rallies) under consideration are $\boldsymbol{I}(R_{i,t\to t+\tau} \geq q)$ for q = 1.3, 1.2, 1.1 and τ = 1, 3, 6, 12 months. The bounds and risk-neutral probabilities are measures of the conditional probabilities of these rally events.

		averaged across firms (number of obs. $T = 324$)				averaged across time (number of obs. $N = 1044$)			
	horizon	1	3	6	12	1	3	6	12
Panel A: $q = 1.3$, up by over 30%									
realized	mean	0.008	0.046	0.115	0.238	0.011	0.055	0.122	0.230
	s.d.	0.018	0.059	0.107	0.165	0.022	0.059	0.090	0.131
lower bound	mean	0.006	0.025	0.046	0.072	0.008	0.030	0.053	0.079
	s.d.	0.007	0.016	0.019	0.020	0.013	0.025	0.030	0.029
risk-neutral	mean	0.009	0.044	0.092	0.163	0.012	0.052	0.102	0.173
	s.d.	0.011	0.031	0.042	0.043	0.018	0.038	0.046	0.045
upper bound	mean	0.012	0.061	0.132	0.249	0.015	0.070	0.144	0.261
	s.d.	0.015	0.045	0.066	0.079	0.022	0.048	0.061	0.064
Panel B: $q = 1.2$, up by over 20%									
realized	mean	0.027	0.112	0.221	0.355	0.032	0.123	0.223	0.347
	s.d.	0.044	0.106	0.154	0.196	0.037	0.078	0.107	0.157
lower bound	mean	0.021	0.063	0.098	0.131	0.027	0.071	0.107	0.137
	s.d.	0.017	0.026	0.025	0.024	0.026	0.038	0.038	0.030
risk-neutral	mean	0.030	0.100	0.169	0.248	0.037	0.110	0.179	0.255
	s.d.	0.025	0.045	0.048	0.040	0.033	0.049	0.050	0.039
upper bound	mean	0.037	0.129	0.229	0.357	0.045	0.141	0.241	0.365
	s.d.	0.033	0.064	0.076	0.074	0.039	0.060	0.064	0.053
Panel C: $q = 1.1$, up by over 10%									
realized	mean	0.127	0.290	0.401	0.500	0.136	0.293	0.393	0.490
	s.d.	0.124	0.176	0.201	0.219	0.067	0.094	0.128	0.176
lower bound	mean	0.097	0.177	0.218	0.231	0.108	0.186	0.222	0.232
	s.d.	0.037	0.030	0.021	0.034	0.046	0.044	0.031	0.022
risk-neutral	mean	0.123	0.239	0.312	0.367	0.134	0.247	0.317	0.367
	s.d.	0.050	0.050	0.036	0.026	0.052	0.050	0.035	0.023
upper bound	mean	0.142	0.290	0.397	0.497	0.154	0.299	0.402	0.497
	s.d.	0.063	0.073	0.063	0.051	0.058	0.058	0.043	0.031

Table A2: Regression tests of the option-implied rally probability measures

This table reports the results of running linear regressions

$$\boldsymbol{I}(R_{i,t\to t+\tau} \geq q) = \alpha + \beta X_{it}(\tau, q) + \varepsilon_{i,t+\tau},$$

in which $q = 1.3, 1.2$ and 1.1, and X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^U$ (the upper bounds), or $\mathbb{P}^*$ (the risk-neutral probabilities). The data are monthly from January 1996 to December 2022. The stocks under consideration are S&P 500 constituents. The return horizon τ is one month, three months, six months, or one year. Values in parentheses are firm-month two-way clustered standard errors following Thompson (2011). Values in square brackets are standard errors following the panel bootstrap procedures of Martin and Wagner (2019) using 2500 bootstrap samples.

	lower bound				risk-neutral				upper bound			
horizon	1	3	6	12	1	3	6	12	1	3	6	12
Panel A: $q = 1.3$, up by over 30%												
α	0.00	0.00	0.03	0.15	0.00	−0.01	0.00	0.08	0.00	−0.01	0.00	0.07
	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)	(0.02)
	[0.00]	[0.00]	[0.01]	[0.02]	[0.00]	[0.00]	[0.01]	[0.02]	[0.00]	[0.00]	[0.01]	[0.03]
β	1.44	1.80	1.88	1.15	1.02	1.17	1.24	0.95	0.82	0.88	0.90	0.67
	(0.19)	(0.15)	(0.14)	(0.14)	(0.14)	(0.11)	(0.09)	(0.11)	(0.12)	(0.09)	(0.07)	(0.08)
	[0.17]	[0.19]	[0.22]	[0.20]	[0.17]	[0.15]	[0.15]	[0.16]	[0.12]	[0.12]	[0.09]	[0.11]
R^2	5.91%	6.76%	4.68%	1.03%	6.17%	7.70%	6.34%	2.28%	6.21%	7.81%	6.73%	2.70%
Panel B: $q = 1.2$, up by over 20%												
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%
Panel C: $q = 1.1$, up by over 10%												
α	0.00	0.12	0.38	0.66	−0.01	0.08	0.25	0.64	−0.01	0.08	0.21	0.49
	(0.01)	(0.01)	(0.03)	(0.04)	(0.01)	(0.02)	(0.03)	(0.05)	(0.01)	(0.02)	(0.04)	(0.06)
	[0.01]	[0.02]	[0.03]	[0.06]	[0.01]	[0.02]	[0.05]	[0.07]	[0.01]	[0.03]	[0.06]	[0.10]
β	1.30	0.97	0.11	−0.70	1.11	0.89	0.47	−0.38	0.96	0.74	0.47	0.01
	(0.09)	(0.08)	(0.11)	(0.17)	(0.08)	(0.09)	(0.12)	(0.13)	(0.08)	(0.08)	(0.10)	(0.13)
	[0.09]	[0.12]	[0.14]	[0.24]	[0.09]	[0.12]	[0.17]	[0.22]	[0.08]	[0.11]	[0.16]	[0.20]
R^2	6.16%	1.81%	0.01%	0.47%	6.58%	2.44%	0.34%	0.14%	6.71%	2.66%	0.67%	0.00%

Appendix F Constructing Firm Characteristics

Firm characteristics used in the multiple regressions for crash probabilities are listed and described below. All variables are constructed using a merged CRSP-Compustat firm-month panel, unless otherwise noted.

Beta (beta). The stock betas are estimated using daily return data within the windows of the last 12 months.

Relative size (rsize). The relative size is the difference between the market capitalization of a firm and the total market capitalization of the S&P 500 index in logarithmic terms.

Book to market ratio (bm). The ratio is firms' book value of equities divided by their market capitalizations, calculated and updated at each fiscal quarter end.

Gross profitability (gprof). The numerator of this measure is the net revenue or the gross profit of a firm at the end of each fiscal quarter. If both of these quantities are missing, we use the sum of operating income and operating expenditures. The denominator is the market value of assets, calculated as a firm's market capitalization plus its book value of debt. Dividing the market value of total assets creates measures that are more sensitive to new firm-specific information (Campbell, Hilscher, and Szilagyi, 2008). Similarly, firm characteristics such as *leverage*, *net income to asset*, and *cash to asset* ratios will also be scaled by the market value of assets throughout our analysis, as will be discussed later.

Momentum and reversals ($r_{(t-6)\to(t-1)}$, $r_{(t-12)\to(t-1)}$ and $r_{(t-1)\to t}$, also named mom6, mom12, and rlag1 in figures). These variables are lagged (net) equity returns of firms from month -6 to -1 and -12 to -1 (two momentum signals), as well as lagged one-month returns (reversals).

CHS volatility (chs.vol). This measure is proposed in Chen, Hong, and Stein (2001) (thus the acronym CHS) for crash forecasting. The volatility is the rolling-window standard deviation of the excess of market returns ($R_i - R_m$), calculated based on daily return samples spanning the last six month.

Turnover (tnover). The stock turnover variable is defined as the monthly trading volume scaled by the number of shares outstanding. Following Chen, Hong, and Stein (2001), we use the average turnover over the lagged six-month data samples as our turnover characteristic. The trading volume on Nasdaq is adjusted according to the procedure

detailed in Appendix B of Gao and Ritter (2010). Specifically, we divide Nasdaq volume by (1) 2.0 from January 1996 to January 2001; (2) 1.8 from February 2001 to December 2001; (3) 1.6 from January 2002 to December 2003; (4) 1.0 after January 2004 to the end of our sample.

Sales growth (saleg). This variable is proposed for predicting industry-level stock crashes by Greenwood, Shleifer, and You (2019). To be considered, firms must have at least two consecutive years of revenue data. We calculate one-year sales growth based on the most recent observations of the changes in revenue.

Short interest (s.int). The fraction of shares held by institutional investors that have been sold short, as considered in Asquith, Pathak, and Ritter (2005); Daniel, Klos, and Rottke (2023). We divide the number of shares held short (available from Compustat) by the number of shares held by institutional investors (aggregated using Thomson-Reuters Institutional 13-F filings).

Leverage. (lev) We compute the leverage of a firm as its total liability (book value of debt) divided by the market value of its total assets (calculated as a firm's market capitalization plus its book value of debt).

Net income over the market value of total assets (ni). This variable is an earnings-based measure of profitability based on net income, proposed by Campbell, Hilscher, and Szilagyi (2008) in the study of bankruptcy forecast.

Cash over the market value of total asset (cash). This characteristic is a liquidity measure with the numerator being cash and short-term investments. It is also incorporated in the econometric model of Campbell, Hilscher, and Szilagyi (2008) to forecast bankruptcy.

Log price per share (log.p). We also include the log share prices, winsorized from above at \$15 per share (before taking logs) following Campbell, Hilscher, and Szilagyi (2008). This variable is mainly used to isolate the tendency of firms traded at low prices to go bankrupt or experience share price crashes.

To eliminate outliers, each of the 15 characteristics described above are winsorized within a 2.5 to 97.5 percentile interval. Table A3 tabulates the summary statistics for all the characteristics used in our sample.

Table A3: Summary statistics of firm characteristics in our sample

char.	mean	sd	median	q25	q75	min	max
beta	1.029	0.498	0.984	0.691	1.293	0.129	2.509
relative size	−6.905	1.105	−7.007	−7.644	−6.272	−10.779	−2.618
book-to-market	0.467	0.330	0.377	0.233	0.614	0.054	1.546
gross profit.	0.158	0.096	0.144	0.092	0.203	0.014	0.485
$r_{(t-1)\to t}$	0.011	0.087	0.011	−0.039	0.060	−0.210	0.247
$r_{(t-6)\to(t-1)}$	0.064	0.219	0.060	−0.066	0.183	−0.451	0.723
$r_{(t-12)\to(t-1)}$	0.136	0.334	0.117	−0.068	0.306	−0.564	1.232
CHS-volatility	0.018	0.009	0.015	0.011	0.021	0.007	0.053
turnover	0.184	0.135	0.144	0.095	0.225	0.037	0.717
sales growth	0.085	0.201	0.060	−0.009	0.141	−0.334	0.912
short int.	0.046	0.065	0.026	0.015	0.047	0.003	0.410
leverage	0.442	0.221	0.401	0.270	0.598	0.108	0.913
net income-to-asset	0.025	0.024	0.026	0.015	0.038	−0.075	0.080
cash-to-asset	0.070	0.074	0.046	0.020	0.089	0.002	0.340
log price	2.679	0.136	2.708	2.708	2.708	1.738	2.708

Appendix G Additional Tables and Figures

Table A4: The risk-neutral correlation between R_m^γ and the crash indicator $\boldsymbol{I}(R_i \leq q)$: bounds based on the Fréchet–Hoeffding inequalities

This table reports summary statistics of the bounds for risk-neutral correlation $\rho^*[R_m^\gamma, \boldsymbol{I}(R_i \leq q)]$, taking the marginal distributions of the market and stock returns as given (from the observed option prices). The correlation bounds are calculated based on the Fréchet–Hoeffding inequalities, where the lower bound is achieved when R_m and R_i are comonotonic and the upper bound is achieved when the two are countermonotonic. We consider four different forecasting horizons and three crash sizes q. The parameter γ is fixed to two according to our calibration exercise for the market returns, which is also the value we use throughout our empirical analysis. The sample covers all S&P 500 firms for each month (end) from January 1996 to December 2022.

		lower bound				upper bound			
q	horizon	mean	median	min	max	mean	median	min	max
30%	1	−0.11	−0.04	−0.82	−0.00	0.09	0.03	0.00	0.79
30%	3	−0.39	−0.38	−0.86	−0.00	0.32	0.31	0.00	0.83
30%	6	−0.58	−0.62	−0.88	−0.00	0.48	0.49	0.00	0.85
30%	12	−0.67	−0.70	−0.89	−0.00	0.61	0.62	0.00	0.85
20%	1	−0.31	−0.28	−0.84	−0.00	0.23	0.21	0.00	0.81
20%	3	−0.64	−0.68	−0.87	−0.00	0.48	0.49	0.00	0.84
20%	6	−0.73	−0.74	−0.88	−0.00	0.59	0.59	0.00	0.87
20%	12	−0.74	−0.75	−0.89	−0.00	0.67	0.68	0.00	0.88
10%	1	−0.63	−0.65	−0.87	−0.00	0.46	0.46	0.00	0.82
10%	3	−0.76	−0.77	−0.87	−0.00	0.60	0.60	0.00	0.84
10%	6	−0.78	−0.79	−0.88	−0.00	0.67	0.68	0.00	0.87
10%	12	−0.78	−0.78	−0.89	−0.00	0.72	0.73	0.00	0.88

Table A5: Relative bound widths: Fréchet–Hoeffding vs. Cauchy–Schwarz

This table reports summary statistics of the relative widths of bounds calculated according to Result 3 (based on the Fréchet–Hoeffding theorem) and Equation 8 (based on the Cauchy–Schwarz inequality). The ratios reported here are the ranges of the Fréchet–Hoeffding bounds divided by the ranges of the Cauchy–Schwarz bounds. We consider four different forecasting horizons and three crash sizes. For every month from January 1996 to December 2022, we compute the two types of bounds for each S&P 500 firm and report the mean, standard deviation (sd), median, 25% and 75% sample quantile (q25 and q75), the minimum and the maximum of the ratio in our full firm-month panel.

crash size (q)	horizon	mean	sd	median	q25	q75	min	max
30%	1	0.099	0.138	0.036	0.005	0.138	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

Table A6: Regression tests of the option-implied crash probability bounds: OLS with time fixed effects

This table reports the results of running linear regressions with time fixed effects

$$\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) = \alpha_t + \beta X_{it}(\tau, q) + \varepsilon_{i,t+\tau},$$

in which $q = 0.7, 0.8$ and 0.9, and X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^U$ (the upper bounds), or $\mathbb{P}^*$ (the risk-neutral probabilities). The data are monthly from January 1996 to December 2022. The stocks under consideration are S&P 500 constituents. The return horizon τ is one month, three months, six months, or one year. Values in parentheses are firm-month two-way clustered standard errors following Thompson (2011). Values in square brackets are standard errors following the panel bootstrap procedures of Martin and Wagner (2019) using 2500 bootstrap samples. Projected R^2s are also reported.

	lower bound				risk-neutral				upper bound			
horizon	1	3	6	12	1	3	6	12	1	3	6	12
Panel A: $q = 0.7$, down by over 30%												
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%
Panel B: $q = 0.8$, down by over 20%												
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%
Panel C: $q = 0.9$, down by over 10%												
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%

Table A7: Regression tests of the option-implied crash probability bounds: OLS with firm fixed effects

This table reports the results of running linear regressions with firm fixed effects

$$\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) = \alpha_i + \beta X_{it}(\tau, q) + \varepsilon_{i,t+\tau},$$

in which $q = 0.7, 0.8$ and 0.9, and X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^U$ (the upper bounds), or $\mathbb{P}^*$ (the risk-neutral probabilities). The data are monthly from January 1996 to December 2022. The stocks under consideration are S&P 500 constituents. The return horizon τ is one month, three months, six months, or one year. Values in parentheses are firm-month two-way clustered standard errors following Thompson (2011). Values in square brackets are standard errors following the panel bootstrap procedures of Martin and Wagner (2019) using 2500 bootstrap samples. Projected R^2s are also reported.

	lower bound				risk-neutral				upper bound			
horizon	1	3	6	12	1	3	6	12	1	3	6	12
Panel A: $q = 0.7$, down by over 30%												
β	0.80	0.77	0.69	0.25	0.55	0.43	0.33	0.13	0.43	0.30	0.20	0.08
	(0.15)	(0.13)	(0.13)	(0.11)	(0.11)	(0.08)	(0.08)	(0.07)	(0.09)	(0.06)	(0.05)	(0.05)
	[0.15]	[0.21]	[0.23]	[0.15]	[0.13]	[0.12]	[0.11]	[0.13]	[0.08]	[0.10]	[0.08]	[0.08]
R^2-proj	2.26%	2.03%	1.25%	0.13%	2.21%	1.63%	0.85%	0.12%	2.14%	1.48%	0.68%	0.10%
Panel B: $q = 0.8$, down by over 20%												
β	0.78	0.73	0.64	0.25	0.57	0.46	0.36	0.15	0.46	0.33	0.23	0.09
	(0.12)	(0.12)	(0.12)	(0.11)	(0.10)	(0.09)	(0.09)	(0.09)	(0.08)	(0.07)	(0.07)	(0.06)
	[0.12]	[0.17]	[0.18]	[0.16]	[0.09]	[0.12]	[0.13]	[0.14]	[0.08]	[0.10]	[0.11]	[0.11]
R^2-proj	2.90%	1.58%	0.85%	0.11%	2.87%	1.37%	0.66%	0.11%	2.82%	1.26%	0.51%	0.08%
Panel C: $q = 0.9$, down by over 10%												
β	0.88	0.67	0.52	0.08	0.72	0.50	0.33	0.04	0.60	0.36	0.20	0.01
	(0.10)	(0.11)	(0.11)	(0.11)	(0.10)	(0.11)	(0.10)	(0.11)	(0.09)	(0.09)	(0.08)	(0.08)
	[0.10]	[0.14]	[0.15]	[0.16]	[0.10]	[0.15]	[0.16]	[0.17]	[0.10]	[0.13]	[0.12]	[0.14]
R^2-proj	2.46%	0.83%	0.39%	0.01%	2.58%	0.81%	0.31%	0.00%	2.58%	0.74%	0.21%	0.00%

Table A8: Regression tests of the option-implied crash probability bounds: OLS with both time and firm fixed effects

This table reports the results of running linear regressions with time and firm fixed effects

$$\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) = \alpha_i + \lambda_t + \beta X_{it}(\tau, q) + \varepsilon_{i,t+\tau},$$

in which $q = 0.7, 0.8$ and 0.9, and X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^U$ (the upper bounds), or $\mathbb{P}^*$ (the risk-neutral probabilities). The data are monthly from January 1996 to December 2022. The stocks under consideration are S&P 500 constituents. The return horizon τ is one month, three months, six months, or one year. Values in parentheses are firm-month two-way clustered standard errors following Thompson (2011). Values in square brackets are standard errors following the panel bootstrap procedures of Martin and Wagner (2019) using 2500 bootstrap samples. Projected R^2s are also reported.

	lower bound				risk-neutral				upper bound			
horizon	1	3	6	12	1	3	6	12	1	3	6	12
Panel A: $q = 0.7$, down by over 30%												
β	0.77	0.71	0.55	0.16	0.56	0.47	0.38	0.13	0.45	0.37	0.29	0.10
	(0.13)	(0.09)	(0.08)	(0.07)	(0.09)	(0.06)	(0.05)	(0.05)	(0.08)	(0.05)	(0.04)	(0.04)
	[0.12]	[0.13]	[0.10]	[0.09]	[0.10]	[0.07]	[0.07]	[0.05]	[0.08]	[0.09]	[0.05]	[0.03]
R^2-proj	1.74%	1.30%	0.64%	0.05%	1.71%	1.16%	0.60%	0.06%	1.68%	1.11%	0.58%	0.07%
Panel B: $q = 0.8$, down by over 20%												
β	0.73	0.59	0.38	0.05	0.57	0.46	0.30	0.06	0.49	0.38	0.24	0.06
	(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]	[0.09]	[0.07]	[0.07]	[0.07]	[0.07]	[0.05]	[0.05]	[0.06]	[0.05]	[0.05]	[0.04]
R^2-proj	1.75%	0.73%	0.25%	0.00%	1.70%	0.70%	0.24%	0.01%	1.67%	0.68%	0.23%	0.01%
Panel C: $q = 0.9$, down by over 10%												
β	0.65	0.35	0.19	−0.03	0.58	0.31	0.18	−0.01	0.53	0.27	0.16	0.01
	(0.05)	(0.05)	(0.05)	(0.06)	(0.04)	(0.04)	(0.05)	(0.05)	(0.04)	(0.04)	(0.04)	(0.04)
	[0.05]	[0.05]	[0.05]	[0.06]	[0.04]	[0.05]	[0.06]	[0.05]	[0.05]	[0.05]	[0.04]	[0.05]
R^2-proj	0.89%	0.19%	0.05%	0.00%	0.88%	0.18%	0.05%	0.00%	0.87%	0.18%	0.05%	0.00%

Table A9: Regression tests of the option-implied crash probability bounds: comparing with the bias-adjusted risk-neutral probabilities

This table reports the results of linear regressions

$$\boldsymbol{I}(R_{i,t\to t+\tau} \leq q) = \alpha + \beta X_{it}(\tau, q) + \varepsilon_{i,t+\tau},$$

in which $q = 0.7, 0.8$ and 0.9, and X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^U$ (the upper bounds), or $\mathbb{P}^*$ (the risk-neutral probabilities). Results for the lower bound and the raw risk-neutral probabilities are identical to those shown in Table 2: they are shown here for comparison. The adjusted risk-neutral probabilities are calculated by correcting for their biases, as examined in the new row of the average differences between the outcome variable and the predictors ("mean diff."). This is done by multiplying the trailing ratios of the average realized crash events over the average risk-neutral probabilities. The data are monthly from January 1996 to December 2022. Firms under consideration are S&P 500 constituents. The return horizon τ is one month, three months, six months, or one year. The values in parentheses are firm-month two-way clustered standard errors following Thompson (2011).

	lower bound				risk-neutral (raw)				risk-neutral (adjusted)			
horizon	1	3	6	12	1	3	6	12	1	3	6	12
Panel A: $q = 0.70$, down by over 30%												
mean diff.	−0.13%	−0.35%	−0.69%	−1.72%	0.10%	1.52%	4.04%	7.38%	0.43%	0.77%	1.01%	1.44%
α	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.01	0.02	0.04	0.08
	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	(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	0.03	0.17	0.19	0.11
	(0.15)	(0.12)	(0.11)	(0.10)	(0.11)	(0.08)	(0.07)	(0.07)	(0.01)	(0.03)	(0.03)	(0.02)
R^2	3.90%	5.37%	5.17%	3.91%	3.77%	4.56%	4.01%	3.06%	0.28%	1.46%	1.45%	0.59%
Panel B: $q = 0.80$, down by over 20%												
mean diff.	0.14%	0.41%	−0.83%	−2.90%	1.03%	4.38%	6.30%	8.35%	0.36%	0.69%	0.53%	0.61%
α	0.00	−0.01	−0.01	0.02	0.00	−0.01	−0.02	0.00	0.02	0.04	0.08	0.13
	(0.00)	(0.01)	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)
β	0.92	1.03	1.15	1.07	0.68	0.69	0.73	0.66	0.21	0.36	0.27	0.13
	(0.11)	(0.09)	(0.09)	(0.08)	(0.09)	(0.07)	(0.07)	(0.07)	(0.08)	(0.04)	(0.03)	(0.02)
R^2	5.65%	5.15%	4.76%	3.69%	5.48%	4.50%	3.89%	2.96%	1.66%	2.53%	1.64%	0.51%
Panel C: $q = 0.90$, down by over 10%												
mean diff.	1.31%	−0.42%	−1.56%	−2.92%	3.98%	5.54%	7.52%	10.31%	0.15%	0.03%	−0.28%	−0.15%
α	−0.02	−0.01	−0.01	0.03	−0.02	−0.02	−0.02	0.00	0.04	0.11	0.16	0.22
	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.02)	(0.02)	(0.03)	(0.01)	(0.01)	(0.01)	(0.01)
β	1.05	1.07	1.12	1.01	0.88	0.83	0.80	0.68	0.52	0.36	0.24	0.10
	(0.08)	(0.07)	(0.07)	(0.08)	(0.08)	(0.08)	(0.08)	(0.09)	(0.05)	(0.04)	(0.03)	(0.03)
R^2	5.46%	3.71%	3.38%	2.41%	5.46%	3.39%	2.80%	1.83%	3.38%	1.72%	1.02%	0.23%

Table A10: Crash probability and returns: OLS tests

This table reports the results of running linear regressions

$$R_{i,t\to t+\tau} = \alpha + \beta X_{it}(\tau, q) + \varepsilon_{i,t+\tau},$$

in which $q = 0.7, 0.8$ and 0.9, and X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^U$ (the upper bounds), or $\mathbb{P}^*$ (the risk-neutral probabilities). The data are monthly from January 1996 to December 2022. The stocks under consideration are S&P 500 constituents. The return horizon τ is one month, three months, six months, or one year. Values in parentheses are firm-month two-way clustered standard errors following Thompson (2011). Values in square brackets are standard errors following the panel bootstrap procedures of Martin and Wagner (2019) using 2500 bootstrap samples.

	lower bound				risk-neutral				upper bound			
horizon	1	3	6	12	1	3	6	12	1	3	6	12
Panel A: $q = 0.7$, down by over 30%												
α	1.01	1.03	1.05	1.08	1.01	1.02	1.03	1.05	1.01	1.02	1.03	1.04
	(0.00)	(0.00)	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)	(0.02)	(0.00)	(0.00)	(0.01)	(0.02)
	[0.00]	[0.01]	[0.01]	[0.02]	[0.00]	[0.01]	[0.01]	[0.03]	[0.00]	[0.01]	[0.01]	[0.03]
β	0.19	0.17	0.23	0.43	0.16	0.16	0.24	0.39	0.14	0.14	0.21	0.31
	(0.25)	(0.16)	(0.16)	(0.20)	(0.20)	(0.11)	(0.11)	(0.13)	(0.16)	(0.09)	(0.08)	(0.09)
	[0.24]	[0.23]	[0.24]	[0.31]	[0.19]	[0.16]	[0.16]	[0.22]	[0.18]	[0.11]	[0.14]	[0.15]
R^2	0.10%	0.13%	0.19%	0.42%	0.13%	0.32%	0.61%	0.92%	0.16%	0.40%	0.87%	1.20%
Panel B: $q = 0.8$, down by over 20%												
α	1.01	1.02	1.04	1.08	1.01	1.02	1.02	1.04	1.01	1.01	1.02	1.03
	(0.00)	(0.01)	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.02)	(0.00)	(0.01)	(0.01)	(0.03)
	[0.00]	[0.01]	[0.01]	[0.02]	[0.00]	[0.01]	[0.02]	[0.04]	[0.00]	[0.01]	[0.02]	[0.04]
β	0.07	0.10	0.13	0.28	0.06	0.11	0.20	0.32	0.06	0.10	0.19	0.27
	(0.10)	(0.09)	(0.12)	(0.15)	(0.08)	(0.07)	(0.09)	(0.12)	(0.07)	(0.06)	(0.07)	(0.09)
	[0.11]	[0.13]	[0.18]	[0.21]	[0.08]	[0.11]	[0.14]	[0.20]	[0.08]	[0.09]	[0.11]	[0.15]
R^2	0.07%	0.11%	0.10%	0.24%	0.10%	0.27%	0.45%	0.66%	0.12%	0.37%	0.72%	0.94%
Panel C: $q = 0.9$, down by over 10%												
α	1.01	1.02	1.04	1.07	1.01	1.01	1.01	1.00	1.01	1.00	0.99	0.97
	(0.00)	(0.01)	(0.01)	(0.02)	(0.00)	(0.01)	(0.02)	(0.04)	(0.00)	(0.01)	(0.02)	(0.04)
	[0.00]	[0.01]	[0.02]	[0.03]	[0.00]	[0.01]	[0.03]	[0.06]	[0.00]	[0.02]	[0.03]	[0.07]
β	0.03	0.07	0.08	0.21	0.03	0.10	0.18	0.34	0.03	0.10	0.18	0.31
	(0.04)	(0.06)	(0.08)	(0.12)	(0.04)	(0.06)	(0.08)	(0.12)	(0.04)	(0.05)	(0.07)	(0.10)
	[0.05]	[0.09]	[0.12]	[0.19]	[0.04]	[0.07]	[0.12]	[0.21]	[0.04]	[0.08]	[0.11]	[0.17]
R^2	0.04%	0.08%	0.05%	0.13%	0.06%	0.24%	0.36%	0.59%	0.08%	0.36%	0.66%	0.93%

Table A11: Regression tests of the option-implied crash probability bounds: adjusted regressions for 20% crashes in one quarter

This table reports the results from the following regressions:

$$\boldsymbol{I}(R_{i,t\to t+3} \leq 0.8) = \beta \cdot X_{it}(3, 0.8) + \lambda \cdot \text{controls}_{it} + \varepsilon_{i,t+3},$$

in which X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^*$ (the risk-neutral probability), or both. The controls are 15 firm characteristics from the literature. All independent variables are transformed to have a unit standard deviation. Regression coefficients are reported in percentage points, and their two-way clustered standard errors are included in the parentheses. The first five columns are simple OLS estimates, and the sixth column reports estimates with time fixed effects, with a projected (within) R^2 replacing the standard ones. Asterisks indicate coefficients whose t-statistics exceed four in magnitude.

	$\boldsymbol{I}(R_{t\to t+3} \leq 0.8)$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{P}^L[R_{t\to t+3} \leq 0.8]$		5.72*	4.10*		10.76	3.56*
		(0.51)	(0.74)		(2.83)	(0.32)
$\mathbb{P}^*[R_{t\to t+3} \leq 0.8]$				2.54	−6.63	
				(0.80)	(2.91)	
beta	1.15		0.36	0.75	0.12	0.90
	(0.30)		(0.34)	(0.35)	(0.31)	(0.26)
relative size	−0.30		−0.23	−0.31	−0.09	0.00
	(0.26)		(0.26)	(0.26)	(0.25)	(0.17)
book-to-market	−0.26		−0.30	−0.28	−0.33	0.14
	(0.21)		(0.21)	(0.21)	(0.21)	(0.16)
gross profit	−0.04		−0.04	−0.05	−0.01	0.09
	(0.19)		(0.18)	(0.19)	(0.18)	(0.15)
$r_{(t-1)\to t}$	−0.51		−0.31	−0.35	−0.40	−0.11
	(0.37)		(0.37)	(0.38)	(0.36)	(0.20)
$r_{(t-6)\to(t-1)}$	−0.74		−0.58	−0.61	−0.66	−0.49
	(0.47)		(0.47)	(0.48)	(0.45)	(0.27)
$r_{(t-12)\to(t-1)}$	−0.02		−0.10	−0.03	−0.19	−0.65
	(0.47)		(0.47)	(0.47)	(0.47)	(0.31)
CHS-volatility	4.12*		1.47	2.42	1.59	1.12
	(0.52)		(0.67)	(0.71)	(0.65)	(0.37)
turnover	0.55		0.15	0.30	0.15	0.69
	(0.54)		(0.51)	(0.52)	(0.50)	(0.25)
sales growth	0.51		0.45	0.50	0.37	0.23
	(0.20)		(0.20)	(0.20)	(0.19)	(0.12)
short int.	1.06*		0.97*	1.04*	0.86*	0.70*
	(0.19)		(0.18)	(0.19)	(0.17)	(0.14)
leverage	−0.19		−0.06	−0.16	0.06	−0.06
	(0.23)		(0.23)	(0.23)	(0.21)	(0.20)
net income-to-asset	−0.49		−0.38	−0.48	−0.24	−0.34
	(0.23)		(0.23)	(0.23)	(0.22)	(0.15)
cash-to-asset	−0.25		−0.41	−0.33	−0.45	−0.15
	(0.17)		(0.17)	(0.17)	(0.16)	(0.13)
log price	0.36		0.77	0.55	0.94	0.48
	(0.22)		(0.22)	(0.21)	(0.24)	(0.17)
intercept	−0.12	−0.01	−0.18	−0.15	−0.20*	
	(0.05)	(0.01)	(0.05)	(0.05)	(0.05)	
R^2/R^2-proj.	5.03%	5.15%	5.63%	5.30%	5.90%	5.20%

Table A12: Regression tests of the option-implied crash probability bounds: adjusted regressions for 20% crashes in one year

This table reports the results from the following regressions:

$$\boldsymbol{I}(R_{i,t\to t+12} \leq 0.8) = \beta \cdot X_{it}(12, 0.8) + \lambda \cdot \text{controls}_{it} + \varepsilon_{i,t+12},$$

in which X stands for $\mathbb{P}^L$ (the lower bounds), $\mathbb{P}^*$ (the risk-neutral probability), or both. The controls are 15 firm characteristics from the literature. All independent variables are transformed to have a unit standard deviation. Regression coefficients are reported in percentage points, and their two-way clustered standard errors are included in the parentheses. The first five columns are simple OLS estimates, and the sixth column reports estimates with time fixed effects, with a projected (within) R^2 replacing the standard ones. Asterisks indicate coefficients whose t-statistics exceed four in magnitude.

	$\boldsymbol{I}(R_{t\to t+12} \leq 0.8)$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{P}^L[R_{t\to t+12} \leq 0.8]$		6.90*	5.17*		9.06*	4.24*
		(0.53)	(0.71)		(2.01)	(0.41)
$\mathbb{P}^*[R_{t\to t+12} \leq 0.8]$				2.46	−4.54	
				(0.92)	(2.14)	
beta	0.91		−0.29	0.52	−0.46	1.12
	(0.41)		(0.42)	(0.46)	(0.41)	(0.40)
relative size	−0.98		−0.41	−0.82	−0.28	0.21
	(0.43)		(0.44)	(0.45)	(0.43)	(0.33)
book-to-market	−1.31		−1.36	−1.33	−1.36	−0.02
	(0.40)		(0.39)	(0.40)	(0.39)	(0.32)
gross profit	−0.60		−0.62	−0.62	−0.61	−0.01
	(0.36)		(0.36)	(0.36)	(0.35)	(0.31)
$r_{(t-1)\to t}$	−0.43		−0.24	−0.30	−0.34	−0.47
	(0.53)		(0.52)	(0.54)	(0.50)	(0.24)
$r_{(t-6)\to(t-1)}$	−1.46		−1.38	−1.38	−1.47	−0.72
	(0.65)		(0.63)	(0.65)	(0.60)	(0.33)
$r_{(t-12)\to(t-1)}$	1.21		0.95	1.15	0.86	0.33
	(0.62)		(0.62)	(0.63)	(0.60)	(0.48)
CHS-volatility	5.30*		2.26	3.68	2.95	2.31
	(0.76)		(0.88)	(0.91)	(0.85)	(0.55)
turnover	0.75		0.41	0.56	0.51	0.57
	(0.80)		(0.78)	(0.77)	(0.76)	(0.41)
sales growth	1.90*		1.76*	1.88*	1.69*	0.85
	(0.34)		(0.33)	(0.34)	(0.33)	(0.23)
short int.	2.45*		2.32*	2.44*	2.24*	2.06*
	(0.40)		(0.40)	(0.40)	(0.39)	(0.31)
leverage	−0.30		0.03	−0.22	0.12	−0.08
	(0.42)		(0.41)	(0.42)	(0.40)	(0.36)
net income-to-asset	−0.87		−0.68	−0.87	−0.54	−0.36
	(0.38)		(0.37)	(0.38)	(0.36)	(0.30)
cash-to-asset	−0.54		−0.81	−0.64	−0.82	−0.34
	(0.31)		(0.30)	(0.32)	(0.30)	(0.25)
log price	1.43		1.97*	1.59*	2.08*	1.20
	(0.35)		(0.36)	(0.35)	(0.37)	(0.30)
intercept	−0.30	0.02	−0.39*	−0.34*	−0.37*	
	(0.08)	(0.01)	(0.08)	(0.08)	(0.08)	
R^2/R^2-proj.	4.49%	3.69%	5.03%	4.63%	5.19%	4.86%

Table A13: Explaining the crash probabilities: the kitchen sink regression

This table reports the results from pooled regressions of the lower bound—our measure of expected crash probabilities—onto the 15 stock characteristics, their squares, and all 105 interaction terms. Only regression coefficients with t-statistics greater than two in magnitude are reported in Panel (c).

(a) Characteristics

	est.	s.e.	t-stat.
chs.vol	1.674	0.109	15.29
beta	0.310	0.035	8.88
log.p	−0.465	0.060	−7.69
rlag1	−0.187	0.037	−5.10
gprof	−0.073	0.022	−3.35
lev	−0.057	0.019	−3.03
tnover	0.248	0.084	2.96
s.int	0.114	0.039	2.90
mom6	−0.178	0.063	−2.82
ni	−0.040	0.024	−1.68
mom12	−0.076	0.068	−1.12
cash	0.019	0.027	0.70
bm	0.013	0.024	0.54
rsize	−0.014	0.033	−0.42
saleg	−0.001	0.031	−0.02

(b) Squared characteristics

	est.	s.e.	t-stat.
rlag1^2	0.143	0.022	6.34
mom12^2	0.122	0.019	6.32
gprof^2	0.043	0.009	4.99
mom6^2	0.145	0.034	4.19
beta^2	−0.105	0.027	−3.91
rsize^2	0.049	0.013	3.66
chs.vol^2	0.158	0.055	2.84
log.p^2	−0.032	0.012	−2.71
s.int^2	−0.022	0.009	−2.38
lev^2	0.019	0.019	1.01
bm^2	−0.013	0.013	−0.99
ni^2	−0.008	0.008	−0.98
tnover^2	−0.010	0.021	−0.48
cash^2	0.001	0.009	0.13
saleg^2	0.001	0.009	0.09

(c) Interaction terms

	est.	s.e.	t-stat.
cash×log.p	0.113	0.014	8.10
bm×gprof	0.108	0.015	7.32
chs.vol×log.p	−0.154	0.028	−5.52
s.int×ni	−0.056	0.011	−5.26
chs.vol×beta	0.173	0.036	4.76
rsize×log.p	0.088	0.022	4.07
lev×rlag1	−0.056	0.015	−3.81
beta×s.int	0.053	0.015	3.60
chs.vol×rsize	−0.103	0.029	−3.49
cash×ni	−0.049	0.015	−3.34
chs.vol×bm	−0.076	0.023	−3.34
cash×rlag1	0.035	0.011	3.28
chs.vol×gprof	−0.078	0.024	−3.18
tnover×ni	0.048	0.015	3.16
chs.vol×lev	−0.082	0.026	−3.11
chs.vol×saleg	0.070	0.023	3.09
s.int×bm	0.034	0.011	3.09
mom6×log.p	0.052	0.017	2.98
ni×log.p	−0.048	0.016	−2.94
chs.vol×cash	0.056	0.019	2.88
chs.vol×mom6	−0.110	0.040	−2.75
s.int×log.p	0.030	0.011	2.63
mom6×lev	−0.068	0.027	−2.52
beta×rsize	0.050	0.020	2.47
chs.vol×ni	−0.050	0.021	−2.36
tnover×cash	−0.040	0.017	−2.31
s.int×gprof	−0.023	0.010	−2.26
beta×gprof	−0.042	0.019	−2.21
beta×rlag1	−0.045	0.021	−2.16

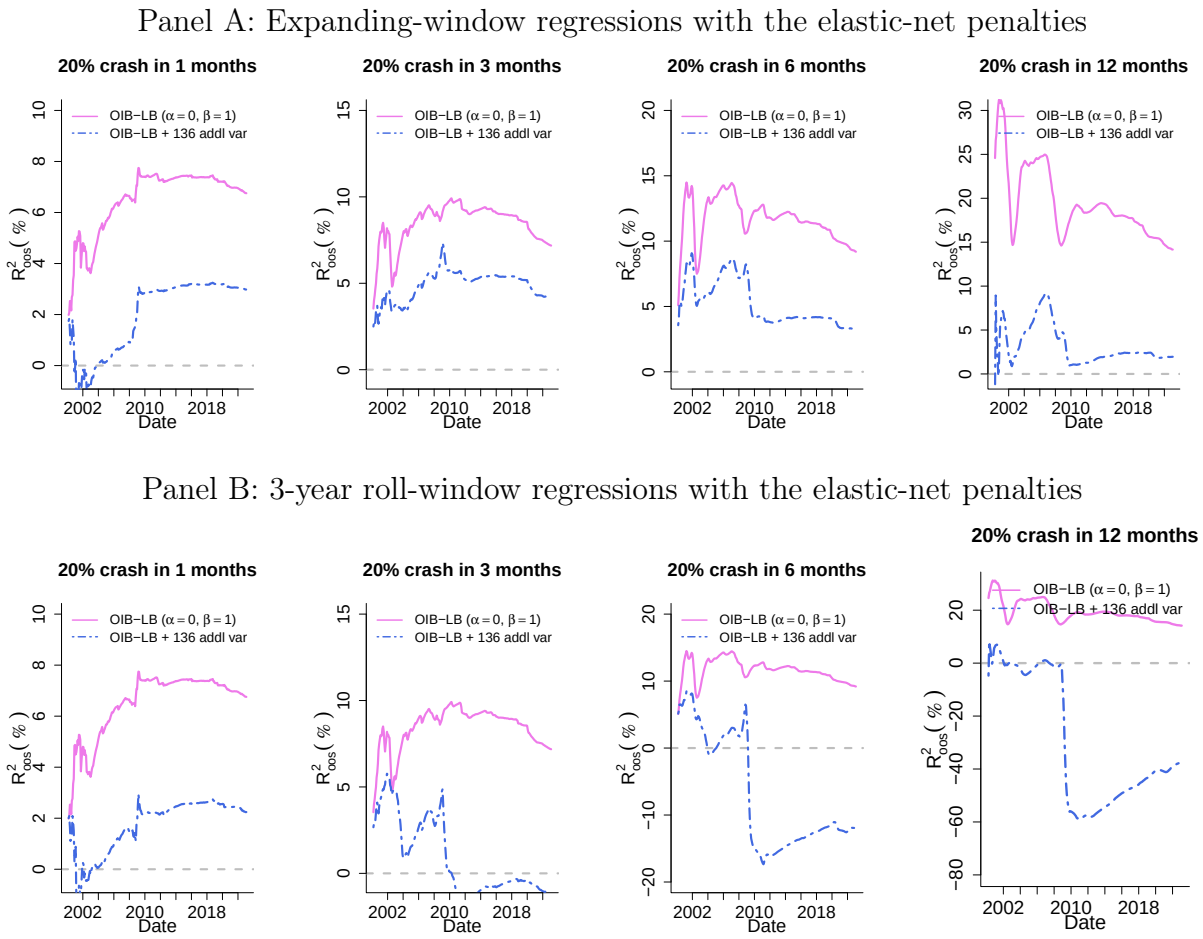


Figure A2: Out-of-sample R^2s: the lower bound and rolling-window regression forecasts aggregating all variables (with squared and interaction terms aggregated through elastic net regressions)

This figure presents the out-of-sample R^2s (R^2_{oos}) for our option-implied lower bound (OIB-LB). At each time point t, we compare the sum of squared forecasting errors from OIB-LB to those from a firm-specific average probability of crashes, calculated over the period $1:(t-\tau)$ ($\tau = 1, 3, 6, 12$).

For comparison, we also report R^2_{oos}s for a forecaster that aggregates 137 variables that include [1] the lower bound and the risk-neutral probabilities (two variables); [2] the 15 stock characteristics considered in Session 3.3, their squared terms, and all their 105 interaction terms (15+15+105=135 variables in total). The variables are combined through expanding-window (Panel A) or 3-year rolling-window (Panel B) elastic net regressions, the tuning parameters of which are chosen through 5-fold cross-validations.

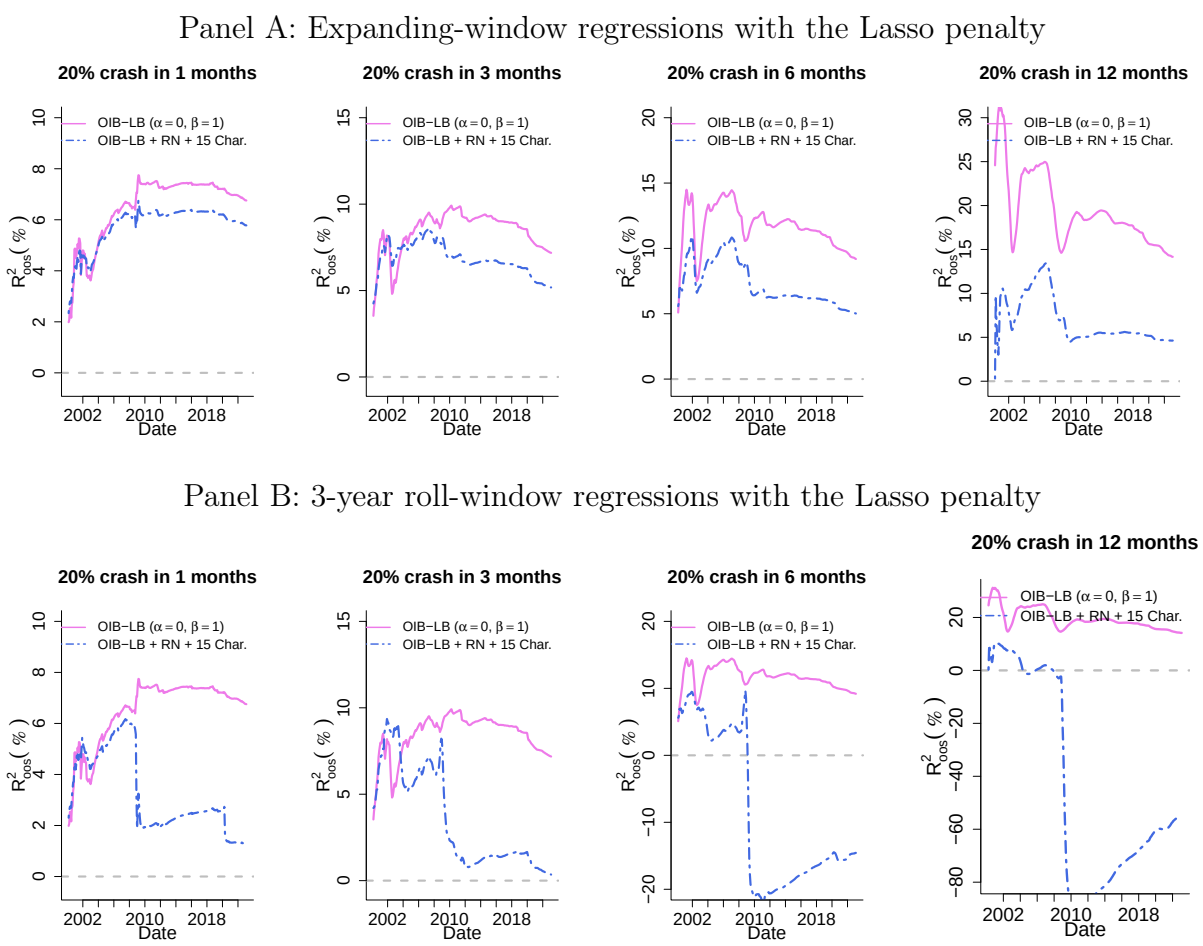


Figure A3: Out-of-sample R^2s: the lower bound and rolling-window regression forecasts aggregating all variables (Lasso version)

This figure presents the out-of-sample R^2s (R^2_{oos}) for our option-implied lower bound (OIB-LB). At each time point t, we compare the sum of squared forecasting errors from OIB-LB to those from a firm-specific average probability of crashes, calculated over the period $1:(t-\tau)$ ($\tau = 1, 3, 6, 12$). For comparison, we also report R^2_{oos}s for a forecaster that aggregates the 15 stock characteristics considered in Session 3.3, the lower bound, and the risk-neutral probabilities. The variables are combined through expanding-window (Panel A) or 3-year rolling-window (Panel B) Lasso regressions, the tuning parameters of which are chosen through 5-fold cross-validations.

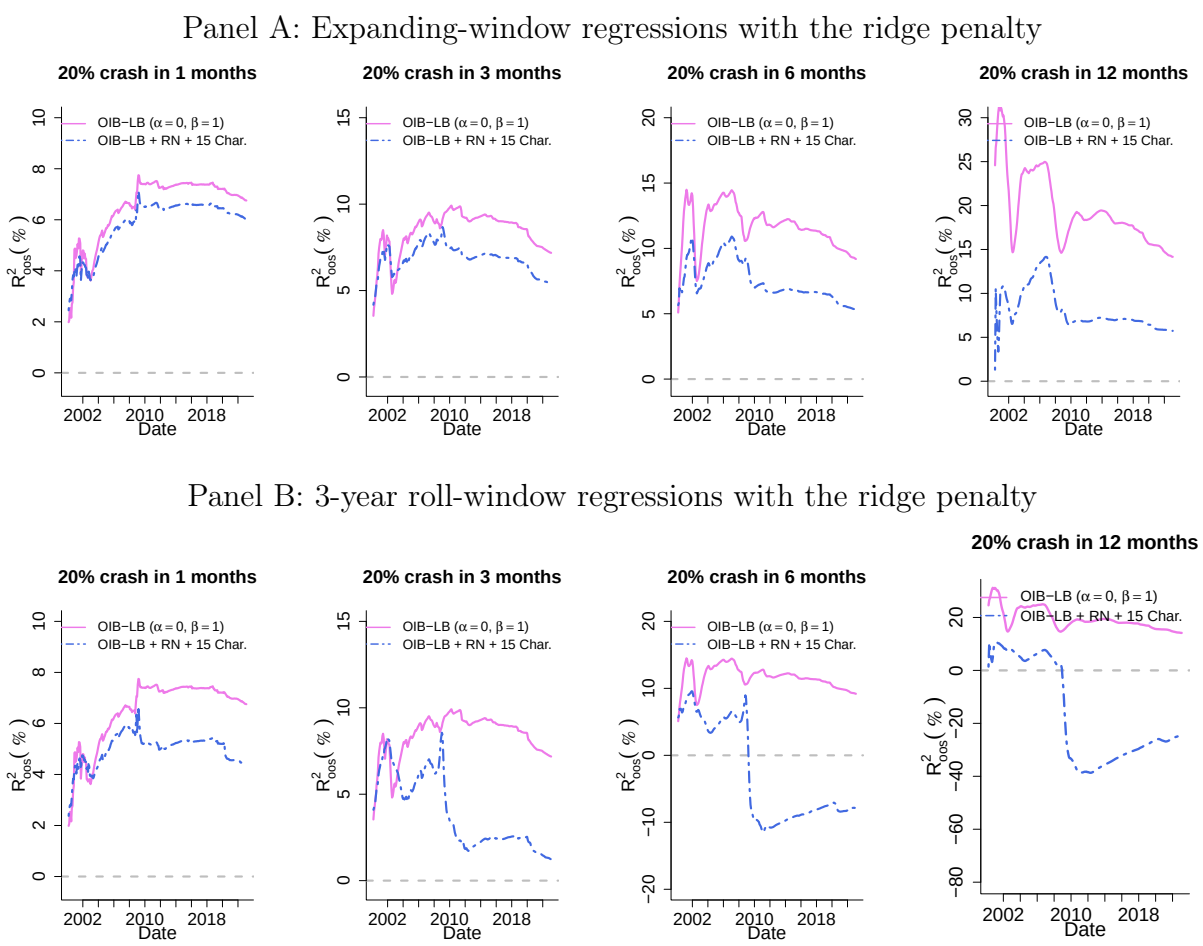


Figure A4: Out-of-sample R^2s: the lower bound and rolling-window regression forecasts aggregating all variables (ridge version)

This figure presents the out-of-sample R^2s (R^2_{oos}) for our option-implied lower bound (OIB-LB). At each time point t, we compare the sum of squared forecasting errors from OIB-LB to those from a firm-specific average probability of crashes, calculated over the period $1:(t-\tau)$ ($\tau = 1, 3, 6, 12$). For comparison, we also report R^2_{oos}s for a forecaster that aggregates the 15 stock characteristics considered in Session 3.3, the lower bound, and the risk-neutral probabilities. The variables are combined through expanding-window (Panel A) or 3-year rolling-window (Panel B) ridge regressions, the tuning parameters of which are chosen through 5-fold cross-validations.

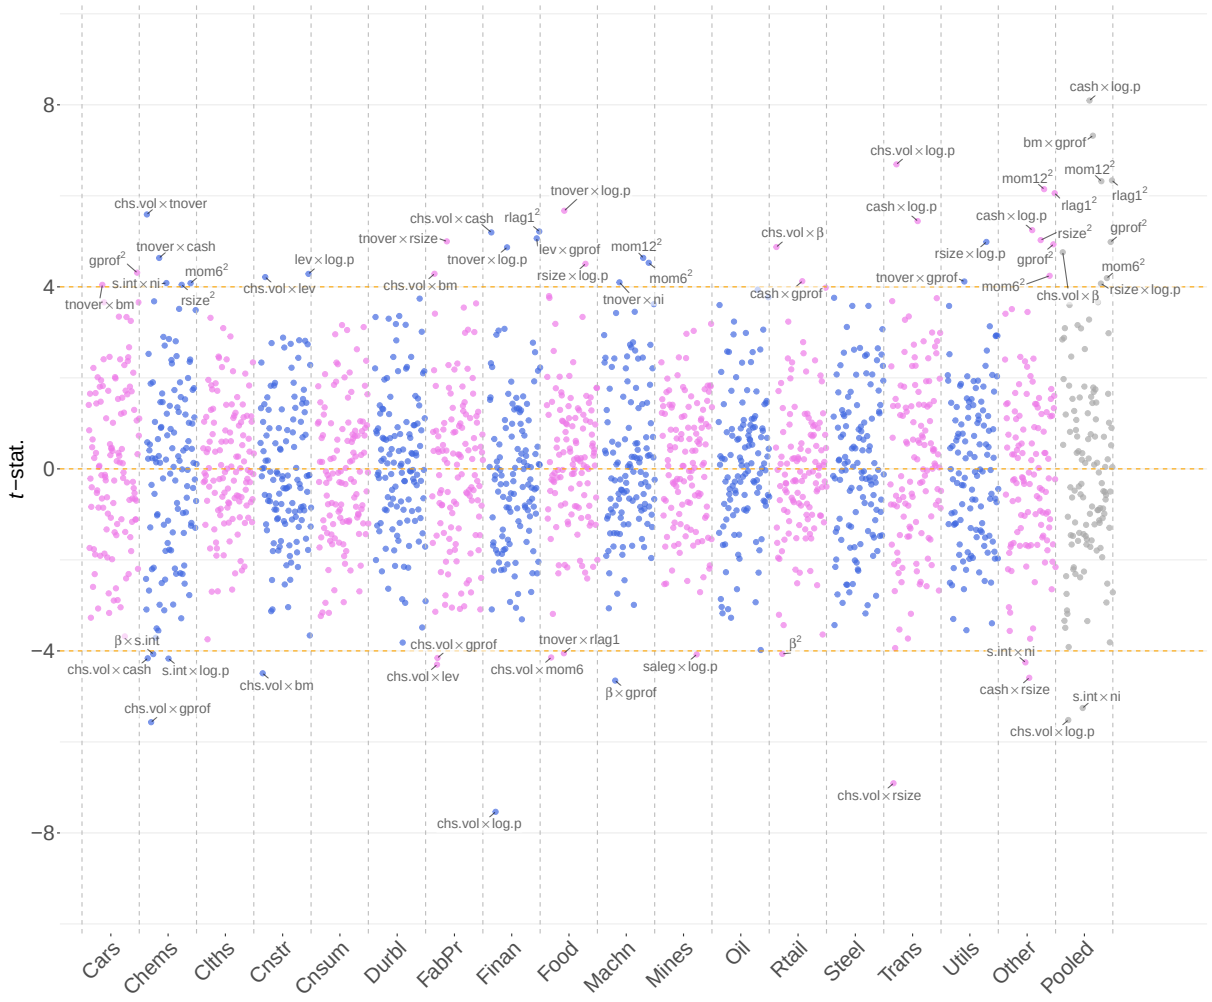


Figure A5: Explaining the crash probabilities: adding squared and interaction terms to the industry-specific and pooled regressions

This figure shows the t-statistics for interaction terms in regressions of the lower bound—our measure of expected crash probabilities—on the 15 stock characteristics and their 105 interactions. The data is divided into sub-samples based on the 17 Fama–French industries. The gray dots in the rightmost column represent results from the entire pooled sample. We label all dots with t-statistics above 4 in absolute value.

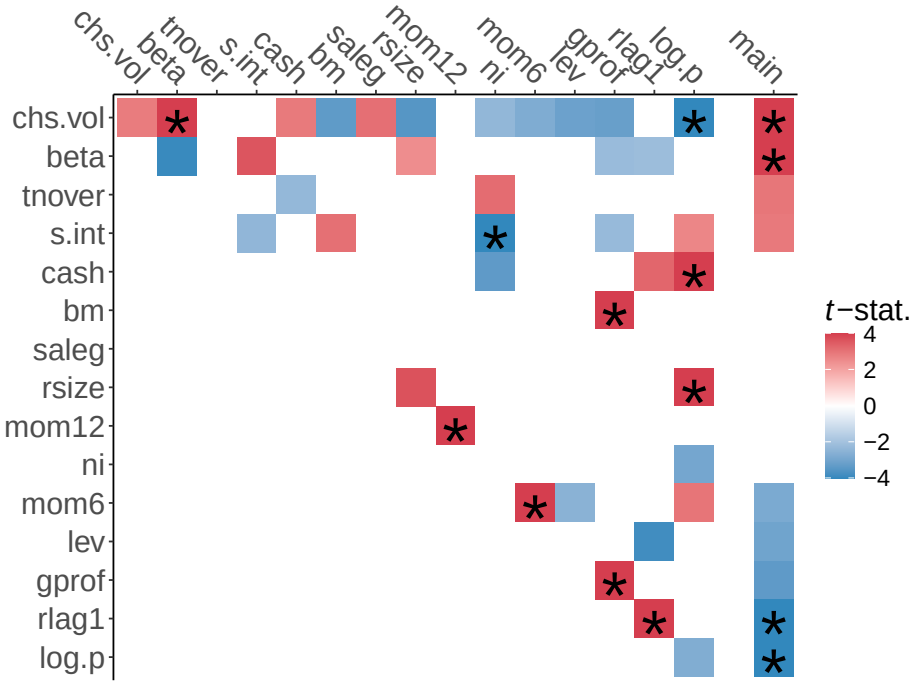


Figure A6: Explaining the crash probabilities: squared and interaction terms

This figure shows the t-statistics for interaction terms in the full-sample regression of the lower bound—our measure of expected crash probabilities—on the 15 stock characteristics, their squaring transforms, and their 105 interactions. The last column, labelled "main", reports the t-statistic for the coefficient on the linear term for each variable, as previously reported in Table A13, Panel a.