

7498462-the-physical-crash-frontier-what-finite-option-quotes-can-and-cannot-reveal

SUMMARY: This paper examines how the physical crash frontier is revealed by observable bid and ask quotes in option markets. The authors characterize the pairs of physical crash probability and expected loss below the crash threshold that these quotes admit. They find that the quotes admit a set of pairs that is convex and can be computed exactly using second-order cone programs at a calibrated risk aversion of two. The authors also note that the quotes beyond the two nearest puts to a 10 percent decline shrink the range of admissible crash probabilities by about 80 percent, yet the upper end remains two to three times its lower end. The paper concludes that a positive floor is a joint statement about prices and a tail restriction, and anything tighter than the frontier is an assumption.

METHOD: The authors collect all risk-neutral distributions that price each quote inside its spread, have the forward as their mean, and are supported on a bounded interval. They characterize the sets of crash probability and crash depth that these distributions generate under the risk-neutral measure and call their boundary the risk-neutral crash frontier. They then pass each admissible distribution through the pricing kernel of a power utility investor with relative risk aversion γ , who holds the market. The resulting physical distribution is then used to define the physical crash frontier of the cross section, which separates the crash scenarios the quotes admit from those they rule out. The set is compact and convex, and its support function in any direction is the value of a finite second-order cone program. The frontier is sharp, meaning that every point in it is generated by an admissible distribution, and every point outside it lies at a positive distance from all of them.

KEY FINDINGS:

- The quotes admit a convex set of pairs (crash probability, expected loss) below the crash threshold.
- The quotes beyond the two nearest puts to a 10 percent decline shrink the range of admissible crash probabilities by about 80 percent.
- The quotes admit a physical crash frontier that is convex and can be computed exactly using second-order cone programs.
- The quotes admit a positive floor, which is a joint statement about prices and a tail restriction.

IMPLICATION FOR TRADING: The physical crash frontier provides a clear boundary of physical crash risk that can be inferred from observable bid and ask quotes. Practitioners can use this information to better understand and manage physical crash risk, especially in times of stress. The frontier can help identify the most extreme scenarios that are plausible given the available quotes, and it can be used to evaluate the robustness of different risk management strategies. However, practitioners should be aware that the frontier is not the only possible boundary, and assumptions may be necessary to further refine the understanding of physical crash risk.

The Physical Crash Frontier: What Finite Option Quotes Can and Cannot Reveal

Jirong Zhuang

Department of Mathematics, University of Macau
Avenida da Universidade Taipa, Macau 999078, China
yc27478@um.edu.mo

September 2026

Abstract

Physical crash probabilities recovered from option prices depend on a pricing kernel and on a risk-neutral distribution that finitely many bid and ask quotes do not identify. For a power utility investor, we characterize the pairs of physical crash probability and expected loss below the crash threshold that the quotes admit; the boundary of this set is the physical crash frontier. Both coordinates are ratios of moments, yet when the index is bounded above the set is convex, and second-order cone programs compute it exactly at the calibrated risk aversion of two. In a decade of weekly S&P 500 cross sections, the quotes beyond the two puts nearest a 10 percent decline shrink the range of admissible crash probabilities by about 80 percent, yet its upper end remains two to three times its lower end. That lower end exists only because the index is bounded. Otherwise, for any investor more risk averse than the log investor, a vanishing probability far in the right tail inflates the denominator and drives the crash probability to zero while every quote stays inside its spread. A positive floor is therefore a joint statement about prices and a tail restriction; anything tighter than the frontier is an assumption.

Keywords: physical crash frontier, physical crash probability, partial identification, bid and ask quotes, pricing kernel

JEL Classification: G12, G13, C14, C61

1 Introduction

Severe market declines are hard to measure for the same reason they matter: they are rare. Over a horizon of a few weeks, realized returns say almost nothing about the probability of a crash or about how deep the losses would be. Option markets price this risk every day, and puts struck below any given threshold trade at observable bid and ask quotes, so it is natural to read crash probabilities from them. The difficulty is that option prices are state prices. They measure what investors pay for protection, and the same fear that makes the measurement urgent also raises that price, so the probabilities implied by option prices overstate physical crash risk, most severely in times of stress ([Martin and Shi, 2026](#)). Getting from prices to physical probabilities requires a restriction on some investor's preferences. This paper asks how much the contracts that actually trade reveal about physical crash risk once that restriction is in place.

The usual approach fits a smooth curve through the observed option prices and differentiates it twice to obtain a single risk-neutral density. The fitted curve blends what the market says with assumptions about strikes that never traded, and it gives no indication of how much of the answer comes from each. We take a different route and keep the ambiguity in view. For a given cross section, we collect every risk-neutral distribution that prices each quote inside its spread, has the forward as its mean, and is supported on a bounded interval. [Zhuang \(2026\)](#) characterizes the sets of crash probability and crash depth that these distributions generate under the risk-neutral measure and calls their boundary the risk-neutral crash

frontier. Those sets describe the price of crash insurance, and a price of insurance mixes the probability of the loss with the compensation for bearing it.

This paper takes the remaining step. Each admissible distribution is passed through the pricing kernel of a power utility investor with relative risk aversion γ who holds the market. [Martin \(2017\)](#) introduced this restriction for the log investor, and [Martin and Shi \(2026\)](#) generalized it to arbitrary risk aversion and showed that the resulting crash probability bounds forecast realized crashes. What comes out is the physical distribution as that investor sees it. Every statement in the paper is conditional on that kernel: the results describe what the quotes imply for her.

For each admissible distribution, two numbers summarize the physical crash risk it implies: the probability that the index ends below the threshold, and the expected loss below the threshold, which we call the shortfall. Collecting these pairs across all admissible distributions gives an identified set in the sense of [Manski \(2003\)](#), and we call its boundary the physical crash frontier of the cross section. It separates the crash scenarios the quotes admit from those they rule out. The set is compact and convex, and its support function in any direction is the value of a finite second-order cone program. Under a regularity condition that can be verified cross section by cross section, the set is sharp: every point in it is generated, or approached as closely as one likes, by an admissible distribution, and every point outside it lies at a positive distance from all of them. Given the kernel and the support condition, the frontier is therefore exactly what the quotes pin down; anything tighter is an assumption. Setting risk aversion to zero recovers, up to the support bound, the risk-neutral sets of [Zhuang \(2026\)](#), so the construction places them in a family indexed by risk aversion.

The joint set matters because risk scenarios come in pairs. A stress test that specifies a crash probability together with an expected loss needs to know whether one distribution can deliver both, and the two marginal intervals cannot say, because their endpoints may come from different distributions. The frontier removes the combinations that no distribution can deliver, cutting most deeply where the shortfall is stressed, and checking any candidate pair amounts to one conic feasibility problem (Corollary 3.5).

Moving from prices to physical probabilities changes the mathematics, and the change turns out to matter economically. Under the risk-neutral measure, crash probability and expected loss are linear in the distribution, and their identified sets come from linear programs. Under the physical measure they are ratios. Our benchmark sets risk aversion equal to two, so the pricing kernel weights each state by the square of the normalized market payoff and the physical crash probability becomes the second moment of the crash states divided by the second moment of all states. The denominator is the problem. It involves every state, including payoffs far above the highest quoted strike, where nothing trades, so two distributions that price every quote identically can disagree about it, and the disagreement has no limit: a vanishing amount of probability placed at an ever larger payoff pushes the denominator toward infinity and the physical crash probability toward zero while every quote stays inside its spread. We call this the dilution channel.

The channel is why no finite set of vanilla quotes can by itself keep the crash probability of an investor more risk averse than the log investor away from zero, and why a positive lower bound requires a restriction on the far right tail. The bound on the index that we maintain is such a restriction. It is a substantive assumption rather than a technicality, and Section 3.5 shows exactly what it does and does not deliver. Conditional depth, the expected further decline given that a crash occurs, is the ratio of shortfall to probability. Both vanish along the same sequence, but the troublesome denominator cancels from their ratio, and depth is untouched.

The empirical part of the paper asks how much the quotes actually narrow these sets. They narrow them a great deal, and most of the narrowing comes from a handful of strikes. We use more than a thousand S&P 500 cross sections, sampled weekly from 2013 through 2023 at three fixed horizons close to one month. Expanding the quote set from the two puts nearest a 10 percent decline to all retained out-of-the-money quotes shrinks the interval of admissible crash probabilities by a median of about 80 percent. Most of that shrinkage arrives with the six additional puts closest to the threshold, and even with every quote in place the upper end of the interval remains two to three times the lower end. The intervals rise in every major stress episode of the sample, from the August 2015 selloff through early 2018 and the

COVID-19 crisis to the 2022 selloff. In March 2020 both ends jump at once: stress raises the level of crash risk and at the same time widens the range of crash probabilities the quotes can accommodate.

Section 2 places the analysis in the literature. Section 3 develops the identification framework and the exact computation, Section 4 describes the data and design, Sections 5 and 6 report the marginal and joint results, and Section 7 concludes. Proofs, the log-investor benchmark, and supplementary evidence are in the appendices.

2 Related Literature

With a continuum of option prices, the risk-neutral distribution is identified under regularity conditions (Breeden and Litzenberger, 1978). Empirical work builds on that identification. Pan (2002) estimates jump risk premia from the joint dynamics of spot and option prices in a parametric model, Backus et al. (2011) compare disaster distributions implied by index options with those calibrated to macroeconomic data, and Bollerslev and Todorov (2011) separate compensation for tail events from the events themselves, reading risk-neutral jump tails from short-maturity options and physical jump tails from high-frequency returns. The premium component of these tail measures forecasts returns (Bollerslev et al., 2015). Each approach reaches its point estimates by adding structure, whether a parametric model or a tail extrapolation. We start one step earlier and ask what is identified before that structure is imposed.

Physical crash probabilities can be approached from three directions. The most direct is to ask: survey respondents report subjective crash probabilities far above historical crash frequencies (Goetzmann et al., 2016). Recovery arguments attempt to extract physical probabilities from state prices while replacing preference restrictions with structural ones: Hansen and Scheinkman (2009) develop the operator decomposition that underlies these arguments, Ross (2015) establishes recovery for a Markov state with a transition-independent kernel, and Borovička et al. (2016) show that the recovered measure differs from the physical one whenever the kernel has a nondegenerate martingale component, so an explicit restriction is unavoidable. A third route supplies the restriction directly. Ait-Sahalia and Lo (2000) and Jackwerth (2000) infer risk aversion by comparing option-implied state prices with return histories, Rosenberg and Engle (2002) estimate time-varying pricing kernels projected on market returns, and Bliss and Panigirtzoglou (2004) calibrate utility so that the transformed densities best forecast realized returns.

The restriction maintained here comes from Martin (2017), whose benchmark investor has log utility and holds the market. Martin and Shi (2026) generalize his expectation identity to arbitrary risk aversion and convert risk-neutral crash probabilities into physical ones. They document that the risk-neutral versions overstate crash risk most severely in crises, and they calibrate risk aversion to two on market returns. Their stock-level bounds rely on Fréchet–Hoeffding arguments because marginal option prices do not identify the joint distribution of a stock with the market. For the market index itself the dependence problem disappears, and identification uncertainty comes instead from finite strikes, bid–ask intervals, and the unobserved right tail.

A finite cross section turns identification into a generalized moment problem, and the idea of answering with a set is old. Lo (1987) derives upper bounds for option prices and expected payoffs from moment restrictions, and Bassett (1997) bounds the probabilities of future prices by the set of distributions that could have generated the observed option values. Bertsimas and Popescu (2002) derive best-possible bounds on option prices from moment information by semidefinite programming, and King et al. (2005) compute arbitrage-free price bounds for a claim from the prices of traded options through convex duality. Davis and Hobson (2007) characterize when a finite set of option prices is consistent with an arbitrage-free model. Cox and Obłój (2011) derive model-free bounds for barrier claims from call and digital call quotes, where mass exactly at a barrier changes the answer. In a single-maturity bid–ask martingale transport problem, Liang et al. (2026) derive a square-root convergence rate for upper-semicontinuous payoffs of bounded variation as spreads vanish.

We read the resulting solution sets as identified sets in the sense of Manski (2003) and describe them through their support functions, following Beresteanu and Molinari (2008). The closest computational relatives are Barratt et al. (2024), who bound the cumulative distribution and distributional risk measures

over risk-neutral probabilities on a discretized price grid, and [van Eekelen \(2023\)](#), who extends generalized moment duality to conditional expectations, the fractional form that conditional depth takes here. The direct predecessor is [Zhuang \(2026\)](#), which introduces the risk-neutral crash frontier, the boundary of the sharp joint set of risk-neutral crash probability and conditional depth implied by finite bid–ask quotes, and shows that separate marginal extrema need not be jointly attainable. We differ in the object that is identified: the bounds here concern physical crash risk under the benchmark kernel. Positive risk aversion makes the transform to physical probabilities normalize by a moment of the unknown distribution. The normalization turns linear objectives into ratios, replaces the linear programs of the risk-neutral analysis with second-order cone programs, and demands a support condition that the risk-neutral problem does not need. That condition is the source of the tail identification problem studied below.

3 From Option Quotes to Sets of Physical Crash Probabilities

We work one cross section at a time. This section defines the admissible risk-neutral distributions, the physical crash quantities they imply, and the identified sets, and then shows how to compute the sets exactly.

3.1 Normalized prices and admissible measures

Fix a valuation date t and an option maturity T . Let D denote the discount factor and F the forward price of the cross section. The random variable S_T is the terminal index level, and we work with the forward-normalized terminal level

$$X = \frac{S_T}{F}. \quad (1)$$

The normalization makes X dimensionless, and under the normalized risk-neutral measure $\mathbb{Q}$ its first moment equals one.

Quote i has strike κ_i in index points and normalized strike $k_i = \kappa_i/F$. We write $\underline{c}_i$ and $\bar{c}_i$ for the normalized bid and ask of a call payoff at that strike, so a retained call quote imposes

$$\underline{c}_i \leq \mathbb{E}_{\mathbb{Q}}[(X - k_i)^+] \leq \bar{c}_i. \quad (2)$$

Put quotes enter the same form through put–call parity: their normalized endpoints are the put bid and ask divided by DF , with $1 - k_i$ added to each endpoint.

Let $\mathcal{I}$ denote a finite set of retained quote indices, and let $B > 1$ be an upper support bound. The admissible set is

$$\mathcal{M}(\mathcal{I}) = \left\{ \mathbb{Q} : \mathbb{Q} \text{ is a probability measure on } [0, B], \right. \\ \left. \mathbb{E}_{\mathbb{Q}}[X] = 1, \quad \underline{c}_i \leq \mathbb{E}_{\mathbb{Q}}[(X - k_i)^+] \leq \bar{c}_i \quad \text{for every } i \in \mathcal{I} \right\}. \quad (3)$$

In general $\mathbb{Q}$ is the normalized state-price measure for maturity T , and when discounting is deterministic it is the usual risk-neutral measure. Support on $[0, B]$ is a maintained economic restriction whose role [Section 3.5](#) isolates. The set $\mathcal{M}(\mathcal{I})$ also depends on B and on (D, F) , which are held fixed throughout, so we suppress them and keep only the quote set $\mathcal{I}$, the object the empirical design varies.

3.2 The power utility investor and the crash quantities

Assumption 3.1 (Power utility pricing kernel). *The investor has constant relative risk aversion $\gamma > 0$, is fully invested in market wealth, and is marginal in the market. Her gross return on market wealth is $R_m = cX$ for a constant $c > 0$, gross returns are strictly positive, and her stochastic discount factor is proportional to $R_m^{-\gamma}$.*

For each $\gamma > 0$ and admissible $\mathbb{Q}$, let $M_\gamma(\mathbb{Q}) = \mathbb{E}_\mathbb{Q}[X^\gamma]$. Under Assumption 3.1, the physical measure $\mathbb{P}_\gamma^\mathbb{Q}$ perceived by the investor satisfies

$$\frac{d\mathbb{P}_\gamma^\mathbb{Q}}{d\mathbb{Q}} = \frac{X^\gamma}{M_\gamma(\mathbb{Q})}. \quad (4)$$

Normalized state pricing makes $d\mathbb{Q}/d\mathbb{P}$ proportional to the stochastic discount factor, hence to $X^{-\gamma}$, and inverting and normalizing under $\mathbb{Q}$ gives the displayed density. The transform restates, for the normalized terminal level, the expectation identity that [Martin \(2017\)](#) derives for a fully invested log investor and [Martin and Shi \(2026\)](#) extend to arbitrary risk aversion. The support and forward restrictions imply $0 < M_\gamma(\mathbb{Q}) \leq B^\gamma$, so the transform is well defined for every admissible measure.¹

Fix a threshold $K \in (0, 1)$, held constant throughout, and define the crash event $\{X \leq K\}$. For each exponent $r \geq 0$, define the truncated moment

$$A_r(\mathbb{Q}) = \mathbb{E}_\mathbb{Q}[X^r \mathbf{1}\{X \leq K\}].$$

The crash probability and the unconditional shortfall, the expected loss below the threshold, are

$$q_{\gamma,K}(\mathbb{Q}) = \mathbb{P}_\gamma^\mathbb{Q}(X \leq K) = \frac{A_\gamma(\mathbb{Q})}{M_\gamma(\mathbb{Q})}, \quad (5)$$

$$\ell_{\gamma,K}(\mathbb{Q}) = \mathbb{E}_{\mathbb{P}_\gamma^\mathbb{Q}}[(K - X)^+] = \frac{KA_\gamma(\mathbb{Q}) - A_{\gamma+1}(\mathbb{Q})}{M_\gamma(\mathbb{Q})}. \quad (6)$$

When $A_\gamma(\mathbb{Q}) > 0$, conditional depth is the ratio

$$s_{\gamma,K}(\mathbb{Q}) = \mathbb{E}_{\mathbb{P}_\gamma^\mathbb{Q}}[K - X \mid X \leq K] = K - \frac{A_{\gamma+1}(\mathbb{Q})}{A_\gamma(\mathbb{Q})} = \frac{\ell_{\gamma,K}(\mathbb{Q})}{q_{\gamma,K}(\mathbb{Q})}. \quad (7)$$

Thus $\ell_{\gamma,K}$ measures the unconditional distance below K , while $1 - K + s_{\gamma,K}$ is the conditional decline from the normalized forward.² Setting $\gamma = 0$ in (5)–(6) makes the transform the identity and recovers the risk-neutral crash probability and normalized put value, the p_K and L_K whose joint sets [Zhuang \(2026\)](#) characterizes. That paper places no upper bound on the state, so the $\gamma = 0$ sets here are the support-capped counterparts of its joint sets. Positive γ reweights each admissible measure by X^γ , undoing the premium that state prices attach to low-payoff states. Probability mass shifts toward high-payoff states, so physical crash probabilities fall below their risk-neutral counterparts. The price of this correction is that the reweighting depends on $M_\gamma(\mathbb{Q})$, a global moment of the unknown distribution.

Mass exactly at the threshold needs its own coordinate. Every retained payoff $(X - k_i)^+$ is continuous, so a finite quote cross section cannot separate an atom at K from mass just above it, yet the two differ for the closed event $\{X \leq K\}$. Let $a_K = \mathbb{Q}\{K\}$ and let $A_r^<(\mathbb{Q}) = \mathbb{E}_\mathbb{Q}[X^r \mathbf{1}\{X < K\}]$ denote the truncated moment over the open event. Then

$$A_r = A_r^< + K^r a_K, \quad KA_\gamma - A_{\gamma+1} = KA_\gamma^< - A_{\gamma+1}^<, \quad (8)$$

so an atom at the threshold adds $K^\gamma a_K / M_\gamma(\mathbb{Q})$ to the crash probability and nothing to the shortfall.

3.3 The identified sets and their geometry

Each admissible measure delivers one probability–shortfall pair. The identified set collects every pair that admissible measures attain or approach:

$$\Theta_{q_\ell}^\gamma(I) = \text{cl} \left\{ (q_{\gamma,K}(\mathbb{Q}), \ell_{\gamma,K}(\mathbb{Q})) : \mathbb{Q} \in \mathcal{M}(I) \right\}, \quad (9)$$

¹In Assumption 3.1, the relation $R_m = cX$ treats X as market wealth after deterministic dividend and numeraire adjustments. Admissible measures with mass at zero enter as a closure relaxation of the positive-return model and leave the closed identified sets unchanged (Lemma A.1).

²Throughout, shortfall refers to this expected loss below the fixed threshold, not to the quantile-based risk measure that shares the name.

where cl denotes Euclidean closure. The closure matters because the event indicator jumps at K : a sequence of admissible measures can slide mass toward the threshold and change the event probability only in the limit, so the raw image need not be closed. Conversely, any point outside $\Theta_{q\ell}^\gamma(I)$ is at a positive distance from every pair an admissible measure can generate.

Conditional depth is a ratio and needs its own set,

$$\Theta_s^\gamma(I) = \text{cl} \left\{ K - \frac{A_{\gamma+1}(\mathbb{Q})}{A_\gamma(\mathbb{Q})} : \mathbb{Q} \in \mathcal{M}(I), A_\gamma(\mathbb{Q}) > 0 \right\}, \quad (10)$$

which is empty if no admissible measure gives the event positive weighted mass. Dividing the shortfall interval by the probability interval can overstate the range, because the four endpoints need not be driven by the same measure. The set in (10) keeps numerator and denominator tied to a common measure.

Theorem 3.2 (Identified sets under a power kernel). *Fix $\gamma > 0$ and $K \in (0, 1)$. Suppose $B < \infty$ and $\mathcal{M}(I)$ is nonempty. Then $\Theta_{q\ell}^\gamma(I)$ is a nonempty compact convex subset of*

$$\{(q, \ell) : 0 \leq q \leq 1, 0 \leq \ell \leq Kq\}.$$

If some admissible $\mathbb{Q}$ has $A_\gamma(\mathbb{Q}) > 0$, then $\Theta_s^\gamma(I)$ is a nonempty compact interval contained in $[0, K]$.

Convexity is not automatic, because mixing two measures does not mix their probability–shortfall pairs linearly: each pair carries its own normalizing denominator $M_\gamma(\mathbb{Q})$. The proof gets around this by mixing the measures with weights inversely proportional to their normalizing moments, which delivers any prescribed convex combination of the two pairs; admissibility survives because the quote and forward restrictions are linear in the measure. Appendix A records the algebra together with the compactness argument.

The boundary of the joint set records where the maintained restrictions bind, and we call it the cross section’s physical crash frontier, the physical counterpart of the risk-neutral crash frontier of [Zhuang \(2026\)](#). By Theorem 3.2 the set is compact and convex, so it is recovered from its support function ([Rockafellar, 1970](#)). For a compact convex $\Theta \subset \mathbb{R}^2$ and a direction $u = (u_q, u_\ell) \in \mathbb{R}^2$, write

$$H_\Theta(u) = \sup_{(q, \ell) \in \Theta} (u_q q + u_\ell \ell). \quad (11)$$

The value $H_\Theta(u)$ locates the supporting line with outward normal u , and the set is the intersection of the half planes those lines bound:

$$\Theta = \bigcap_{\|u\|_2=1} \{(q, \ell) : u_q q + u_\ell \ell \leq H_\Theta(u)\}.$$

The directions $(\pm 1, 0)$ and $(0, \pm 1)$ return the marginal interval endpoints, and intermediate directions carry the joint information that the marginals miss.

Adding quotes can only shrink the sets. If $\mathcal{I}_1 \subseteq \mathcal{I}_2$, then $\mathcal{M}(\mathcal{I}_2) \subseteq \mathcal{M}(\mathcal{I}_1)$, so both identified sets under $\mathcal{I}_2$ are subsets of their counterparts under $\mathcal{I}_1$. This monotonicity organizes the empirical design in Section 4. The sets are monotone in the support bound as well: enlarging B enlarges the admissible set, so both identified sets expand weakly with B , and Section 5.4 traces this monotone family.

3.4 Exact computation under $\gamma = 2$

Theorem 3.2 holds for every $\gamma > 0$, and integer values of γ admit exact finite representations. The empirical benchmark is $\gamma = 2$, the calibrated choice in [Martin and Shi \(2026\)](#), and Appendix B develops the $\gamma = 1$ log investor.

From here on we fix $\gamma = 2$, drop the superscript on the sets, and write

$$q_{2,K}, \quad \ell_{2,K}, \quad s_{2,K}, \quad M(\mathbb{Q}) = \mathbb{E}_{\mathbb{Q}}[X^2], \quad \Theta_{q\ell}(I), \quad \Theta_s(I).$$

The exact statements below maintain one regularity condition.

Assumption 3.3 (Interior regularity (a Slater-type condition)). *Some $\mathbb{Q}^\circ \in \mathcal{M}(\mathcal{I})$ prices every retained quote strictly inside its interval, and there are bounded open intervals $U_- \subset (0, 1)$ and $U_+ \subset (1, B)$, disjoint from all strike and threshold boundaries, with $\mathbb{Q}^\circ(U_-) > 0$ and $\mathbb{Q}^\circ(U_+) > 0$.*

The condition gives room to perturb an admissible measure without leaving $\mathcal{M}(\mathcal{I})$: the strict quote slack absorbs small changes in option values, and the interior mass on both sides of the forward restores the mean. Relative to the strict quote interior condition of [Zhuang \(2026\)](#), the new ingredient is the interior mass, which the mean restoration needs. The limiting arguments in [Appendix A](#) use exactly this room, and [Section 4.3](#) verifies the condition cross section by cross section.

Computing a support value means maximizing a linear functional of $(q_{2,K}, \ell_{2,K})$ over all admissible measures, an infinite-dimensional problem whose objectives are ratios of moments. Two transformations turn it into a finite convex program: a change of scale removes the ratio, and a partition of the support at the strikes compresses the measure into finitely many moments whose feasible values are characterized exactly.

The change of scale is the linear-fractional normalization of [Charnes and Cooper \(1962\)](#). Define $\rho = M(\mathbb{Q})^{-1}$ and the scaled measure $d\tilde{\mathbb{Q}} = \rho d\mathbb{Q}$. The mass, forward, and second-moment restrictions become

$$\int d\tilde{\mathbb{Q}} = \rho, \quad \int X d\tilde{\mathbb{Q}} = \rho, \quad \int X^2 d\tilde{\mathbb{Q}} = 1, \quad (12)$$

and the two crash coordinates become linear in the scaled measure:

$$q_{2,K} = \int_{[0,K]} X^2 d\tilde{\mathbb{Q}}, \quad \ell_{2,K} = \int_{[0,K]} (K - X) X^2 d\tilde{\mathbb{Q}}. \quad (13)$$

Each quote band is multiplied by ρ , which keeps the quote restrictions linear in the pair $(\tilde{\mathbb{Q}}, \rho)$. Nothing is lost in the change of variables: since $0 \leq X \leq B$ and $\mathbb{E}_{\mathbb{Q}}[X] = 1$, the global moment satisfies $1 \leq M(\mathbb{Q}) \leq B$, so $\rho \in [1/B, 1]$ and $\mathbb{Q}$ is recovered from $(\tilde{\mathbb{Q}}, \rho)$ by division.

The second transformation partitions the support at the strikes. Let $0 = b_0 < b_1 < \dots < b_J = B$ be the ordered distinct elements of $\{0, K, B\} \cup \{k_i : i \in \mathcal{I}, 0 < k_i < B\}$, and call the intervals $I_j = [b_j, b_{j+1}]$ cells. Let j_K be the index with $b_{j_K} = K$, so that I_{j_K} is the cell beginning at the threshold, and let $\mathcal{J}_- = \{j : j < j_K\}$ collect the cells below it. Decompose the scaled measure into local components,

$$\tilde{\mathbb{Q}} = \sum_j \mu_j,$$

where μ_j is a nonnegative measure supported on I_j ; such a decomposition always exists. An atom at a shared strike endpoint may be assigned to either adjacent cell, and no statement below depends on that choice, because every retained payoff is continuous there. The threshold is the one exception: any atom at K is assigned to the cell beginning at K and tracked by the explicit coordinate introduced below. For each cell j and integer order r , define the scaled cell moment

$$y_{r,j} = \int x^r d\mu_j(x),$$

and let the vector $\mathbf{y}$ collect these moments, with orders $r = 0, 1, 2, 3$ on the cells in $\mathcal{J}_-$ and orders $r = 0, 1, 2$ elsewhere. These are all the moments the problem uses. The normalization (12) reads

$$\sum_j y_{0,j} = \rho, \quad \sum_j y_{1,j} = \rho, \quad \sum_j y_{2,j} = 1. \quad (14)$$

Every retained call payoff is affine on every cell, because its only kink sits at a strike and every strike is a cell boundary. The scaled value of quote i is therefore linear in the cell masses and first moments,

$$C_i(\mathbf{y}) = \sum_{j: b_j \geq k_i} (y_{1,j} - k_i y_{0,j}), \quad \rho \underline{c}_i \leq C_i(\mathbf{y}) \leq \rho \bar{c}_i, \quad (15)$$

and the targets (13) are linear in the second and third moments of the cells below the threshold.

The remaining question is which vectors $\mathbf{y}$ come from actual local measures. This is the truncated Hausdorff moment problem, and on bounded cells it has a complete answer. On a cell $[a, b]$ with orders up to three, a vector (y_0, y_1, y_2, y_3) is the moment sequence of some nonnegative measure on $[a, b]$ if and only if two localizing matrices are positive semidefinite:

$$\begin{pmatrix} y_1 - ay_0 & y_2 - ay_1 \\ y_2 - ay_1 & y_3 - ay_2 \end{pmatrix} \succeq 0, \quad \begin{pmatrix} by_0 - y_1 & by_1 - y_2 \\ by_1 - y_2 & by_2 - y_3 \end{pmatrix} \succeq 0. \quad (16)$$

On a cell with orders up to two, the conditions are the moment matrix

$$\begin{pmatrix} y_0 & y_1 \\ y_1 & y_2 \end{pmatrix} \succeq 0, \quad (17)$$

combined with the support inequalities

$$ay_0 \leq y_1 \leq by_0, \quad y_2 \leq (a + b)y_1 - aby_0. \quad (18)$$

These conditions are necessary and sufficient on a bounded interval (see Krein and Nudelman, 1977; Bertsimas and Popescu, 2002). A two-by-two positive semidefinite constraint is in turn one second-order cone constraint (Alizadeh and Goldfarb, 2003), since

$$\begin{pmatrix} u & v \\ v & w \end{pmatrix} \succeq 0 \iff u \geq 0, \quad w \geq 0, \quad \|(2v, u - w)\|_2 \leq u + w.$$

Cell by cell, the feasible moment vectors form an intersection of linear restrictions and second-order cone constraints.

One coordinate remains: the threshold atom of Section 3.2. Let $\tilde{a}_K \geq 0$ denote the scaled event mass at $X = K$. The atom is assigned to the threshold cell I_{j_K} , whose moment conditions then apply net of the atom; Appendix A records this step. The targets are the linear functions

$$q_{2,K} = \sum_{j \in \mathcal{J}_-} y_{2j} + K^2 \tilde{a}_K, \quad \ell_{2,K} = \sum_{j \in \mathcal{J}_-} (Ky_{2j} - y_{3j}), \quad (19)$$

so the atom adds probability and no shortfall, exactly as in (8).

The program is now complete. Its decision variables are the cell moments $\mathbf{y}$ together with the two scalars ρ and $\tilde{a}_K$. Its constraints are of three kinds: the global equalities (14); the pair of scaled quote inequalities (15) for each retained quote; and one moment cone per cell, (16) on the cells in $\mathcal{J}_-$ and (17)–(18) elsewhere, with the threshold cell taken net of the atom. Maximizing the linear objective $u_q q_{2,K} + u_\ell \ell_{2,K}$ from (19) is therefore a finite second-order cone program (SOCP) for each direction u .³

Corollary 3.4 (Exact computation under $\gamma = 2$). *Suppose Assumption 3.3 holds. Then $\Theta_{q\ell}(\mathcal{I})$ equals the image of the SOCP feasible region in $(\mathbf{y}, \rho, \tilde{a}_K)$ under the linear map (19), and every support value $H_{\Theta_{q\ell}}(u)$ is both the SOCP optimum and the supremum over admissible measures. If some admissible measure has $A_2(\mathbb{Q}) > 0$, analogous programs give the endpoints of the closed interval $\Theta_s(\mathcal{I})$.*

Appendix A proves the corollary in two directions: every admissible measure produces a feasible vector, and every feasible vector is produced by an admissible measure or by a limit of admissible measures, with Assumption 3.3 supplying the approximating sequences. Conditional depth needs a second normalization, because its denominator is the truncated rather than the global second moment; scaling by $A_2(\mathbb{Q})^{-1}$ makes the depth objective linear, and the same cells and moment cones apply. Unlike ρ , the depth scaling has no upper bound, so the depth programs need not attain their optimal values, which nevertheless equal the endpoints of the closed depth interval.

³For integer γ , the same construction applies with cell moments up to order $\gamma + 1$ below the threshold and order γ elsewhere. The truncated moment conditions become positive semidefinite constraints on larger moment and localizing matrices (Krein and Nudelman, 1977; Bertsimas and Popescu, 2002), so the support values are computed by semidefinite programming. Second-order cones suffice exactly when $\gamma \leq 2$, where every matrix is two by two.

Corollary 3.5 (Scenario feasibility). *Under Assumption 3.3, a pair (q°, ℓ°) lies in $\Theta_{q\ell}(\mathcal{I})$ if and only if the SOCP feasible region is nonempty once the two linear targets (19) are fixed at (q°, ℓ°) . Screening a candidate crash scenario is one conic feasibility problem.*

The corollary follows from Corollary 3.4: the identified set equals the image of the feasible region under a linear map, so membership of a candidate pair is equivalent to feasibility with the image coordinates pinned.

Assembling the identified sets from the programs is then direct. Each direction u is one SOCP whose solution returns the support value $H_{\Theta_{q\ell}}(u)$ together with an optimal pair (q, ℓ) on the frontier. The four coordinate directions give the marginal probability and shortfall intervals, and the two depth programs give the depth interval. For the frontier itself, a finite grid of unit directions produces an inner and an outer polygon that sandwich the set. The convex hull of the optimal pairs is contained in $\Theta_{q\ell}(\mathcal{I})$, while the intersection over the grid of the supporting half planes $\{(q, \ell) : u_q q + u_\ell \ell \leq H_{\Theta_{q\ell}}(u)\}$ contains it. The gap between the polygons measures what the grid can still miss, so the grid is doubled until the gap falls below a preset tolerance, and the figures plot the outer polygon. Half planes computed for a coarser quote set remain valid under every finer set, so intersecting the accumulated half planes yields frontiers that are nested by construction.

3.5 What the right tail identifies

The support cap matters because it controls the denominator of the $\gamma = 2$ transform. Quotes can be dense near the crash threshold and still leave the far right tail weakly constrained, and the finite cap is what prevents a small probability at a very large value of X from carrying an arbitrarily large second moment. Let $\mathcal{M}_\infty(\mathcal{I})$ denote the admissible measures of (3) with $[0, B]$ replaced by $[0, \infty)$ and with finite second moment.

Proposition 3.6 (Zero lower bound under a perturbable right tail). *Suppose $\mathcal{M}_\infty(\mathcal{I})$ contains a measure $\mathbb{Q}^\circ$ whose option values lie strictly inside every retained quote interval and which has positive mass on bounded open intervals $U_- \subset (0, 1)$ and $U_+ \subset (1, \infty)$. Then*

$$\sup_{\mathbb{Q} \in \mathcal{M}_\infty(\mathcal{I})} M(\mathbb{Q}) = \infty, \quad \inf_{\mathbb{Q} \in \mathcal{M}_\infty(\mathcal{I})} q_{2,K}(\mathbb{Q}) = \inf_{\mathbb{Q} \in \mathcal{M}_\infty(\mathcal{I})} \ell_{2,K}(\mathbb{Q}) = 0, \quad (20)$$

and the closure of the joint image $\{(q_{2,K}(\mathbb{Q}), \ell_{2,K}(\mathbb{Q})) : \mathbb{Q} \in \mathcal{M}_\infty(\mathcal{I})\}$ contains the origin.

The mechanism is the dilution channel of the introduction. Mass of order $H^{-3/2}$ placed at a state H far in the right tail changes fixed-strike option values by a vanishing amount but raises the second moment by order $H^{1/2}$. A small transfer between U_- and U_+ restores total mass and the forward restriction, and strict quote slack absorbs the remaining changes in option values. Since $A_2(\mathbb{Q}) \leq K^2$, the inflated denominator drives $q_{2,K}$ to zero and, along the same sequence, the unconditional shortfall as well. Without a support bound, the quotes cannot exclude physical distributions that assign the crash negligible probability and negligible expected loss. Conditional depth is the ratio of these two vanishing coordinates, yet it is unaffected, because $M(\mathbb{Q})$ cancels from the ratio; its set can still respond to right-tail restrictions through the quote and forward constraints.

Remark 3.7 (Log utility is the dividing line). The collapse is not specific to $\gamma = 2$. For $\gamma \leq 1$, Jensen's inequality and the forward restriction give $M_\gamma(\mathbb{Q}) \leq 1$ for every admissible measure, so mass far in the right tail cannot inflate the denominator of the transform. For every $\gamma > 1$, with $\mathcal{M}_\infty(\mathcal{I})$ restricted to measures whose moment of order γ is finite, the construction in the proof applies with mass $H^{-(1+\gamma)/2}$ at the state H . Option values then change by $O(H^{(1-\gamma)/2})$, the moment M_γ grows without bound, and $q_{\gamma,K}$ and $\ell_{\gamma,K}$ vanish along the sequence. The dilution channel opens exactly when the investor is more risk averse than the log investor.

4 Data and Design

4.1 SPX option cross sections

The data are end-of-day bid and ask quotes on S&P 500 index options from OptionMetrics, covering January 2013 through August 2023. We retain quotes with positive strikes, nonnegative bids, positive asks, uncrossed bid–ask intervals, and positive trading volume. Quotes enter the analysis only through these intervals. The sample uses exact calendar horizons of 23, 30, and 37 days and keeps at most one cross section per horizon in a calendar week, observed on Wednesdays. It contains 1,115 horizon-specific observations covering 537 distinct calendar weeks, and each cross section carries its own discount factor and forward.

4.2 Crash event and nested quote sets

The main event is a terminal decline of at least 10 percent from each cross section’s forward, so $K = 0.90$. Appendix C examines the more extreme event $K = 0.85$. The main support condition $B = 5$ allows the terminal index level to reach five times its forward over a 23–37 day horizon. It is a loose economic bound, and its only role is to rule out states in which an arbitrarily large market payoff would dominate the kernel normalization. Section 5.4 examines tighter and looser caps and the unbounded specification.

The design expands the quote set in four steps. The first set, $\mathcal{I}^{\text{pair}}$, contains the closest retained put strike on each side of K . The second, $\mathcal{I}^{\text{local}}$, adds the six nearest additional put strikes, for eight in total. The third, $\mathcal{I}^{\text{wing}}$, contains the complete retained out-of-the-money (OTM) put wing. The final set, $\mathcal{I}^{\text{OTM}}$, which we call the full OTM set, adds all retained OTM calls and so comprises every retained OTM quote. Thus

$$\mathcal{I}^{\text{pair}} \subseteq \mathcal{I}^{\text{local}} \subseteq \mathcal{I}^{\text{wing}} \subseteq \mathcal{I}^{\text{OTM}}. \quad (21)$$

By the monotonicity of Section 3.3, every interval and every support value contracts weakly along this sequence, so the design separates information near the event threshold from information in the rest of the put wing and on the call side.

4.3 Reported statistics, verification, and inference

For each cross section we solve the programs of Section 3.4 and obtain the closed identified sets. All results condition on $\gamma = 2$, $K = 0.90$, $B = 5$, and each cross section’s fixed discounting and forward. Summaries are formed separately by horizon: the reported lower endpoint is the median of the lower endpoints across cross sections, and similarly for the upper endpoint. The main measure of information is the proportional reduction in the width of the crash probability interval. For nested quote sets $\mathcal{I}_1 \subseteq \mathcal{I}_2$ with widths $w_1 > 0$ and w_2 , the reduction within a cross section is $\Delta = 1 - w_2/w_1$, and we report its median across cross sections within each horizon. Sampling variation of these medians is assessed with a circular block bootstrap on calendar weeks, using eight-week blocks and 2,000 replications stratified by horizon (Politis and Romano, 1992; Sharipov and Wendler, 2013). Bracketed entries are percentile 95 percent intervals.

Assumption 3.3 is verifiable one cross section at a time. For each cross section we construct an admissible measure that prices every retained quote strictly inside its interval and places mass on both sides of the forward away from the strike and threshold boundaries. The construction succeeds in 1,111 of the 1,115 cross sections, and Table 4 in Appendix C reports the counts by horizon.

The programs under the full OTM set are solved to verified optimality in 1,098 cross sections, which form the samples of Table 1. Comparisons across quote sets use the 976 cross sections solved to verified optimality under all four quote sets. Cross sections that fail verification for a program are dropped from the affected statistics rather than approximated, and each table and figure reports its own counts.

Table 1: Physical crash risk and the information in SPX option quotes

<i>Panel A: Main identified sets</i>					
Horizon	Cross sections	OTM quotes	$q_{2,K}$	$\ell_{2,K}$	$s_{2,K}$
23 days	402	105.0	[0.71, 2.11]	[0.047, 0.061]	[2.44, 7.35]
30 days	372	144.0	[1.42, 3.29]	[0.086, 0.105]	[2.81, 7.00]
37 days	324	95.5	[1.97, 4.53]	[0.127, 0.152]	[2.94, 7.44]

<i>Panel B: Median widths and reductions</i>					
Horizon	Put pair	Local eight	Put wing	All OTM	Total reduction [95% CI]
23 days	7.05	1.86	1.39	1.34	79.9 [78.4, 81.7]
30 days	12.14	2.87	1.91	1.82	84.4 [82.9, 85.3]
37 days	13.80	3.42	2.50	2.36	81.8 [80.3, 83.0]

<i>Panel C: Marginal source of probability identification</i>				
Added quote layer	23 days	30 days	37 days	
Pair to local eight	69.6 [66.1, 72.0]	72.2 [69.1, 74.8]	72.5 [70.2, 75.0]	
Local eight to put wing	28.0 [25.9, 31.0]	35.2 [32.7, 36.9]	26.9 [25.3, 28.4]	
Put wing to all OTM	2.6 [2.2, 3.3]	4.2 [3.3, 5.2]	5.0 [4.1, 6.4]	

Notes: One S&P 500 cross section per week and horizon, January 2013 through August 2023. All results use $\gamma = 2$, $K = 0.90$, $B = 5$, and each cross section's fixed discounting and forward. Panel A reports the number of cross sections, median retained OTM quote counts, and medians of interval endpoints across cross sections, in percent. The $s_{2,K}$ medians use 402, 371, and 324 cross sections. Panel B reports median interval widths for the crash probability, in percentage points, and the median proportional reduction from the put pair to all OTM quotes, in percent. Panel C reports median proportional reductions relative to the preceding quote set, in percent. Brackets are 95 percent calendar-week block bootstrap intervals.

5 What Do the Traded Quotes Reveal?

5.1 Main identified sets

Table 1 reports the main results. Under all retained OTM quotes, the median identified crash probability interval rises with the horizon, reaching 1.97 to 4.53 percent at 37 days. Even with every retained OTM quote in use, the upper endpoint remains two to three times the lower endpoint, so the quotes still admit economically different assessments of the same event.

Conditional depth is identified less tightly, with median intervals spanning roughly 2.4 to 7.4 percent of additional decline below the 10 percent threshold. The contrast between the two coordinates is not specific to this event: at $K = 0.85$ the probability intervals are narrower and the depth intervals wider, with the same ordering of quote-set contributions (Appendix C).

Panel B of Table 1 shows how much ambiguity remains under each quote set. All retained OTM quotes cut the put pair widths by a median of 80 to 84 percent across horizons, and the bootstrap intervals are narrow relative to these magnitudes.

5.2 Physical crash risk over time

Figure 1 plots the probability intervals under the full OTM set by horizon. They rise in every major stress episode of the sample, and the width often rises with the level. Both endpoints move sharply upward around the COVID-19 episode at every horizon, so stress raises the implied crash risk and, at the same time, the range of crash probabilities the quotes admit.




Figure 1: Time variation in quote-compatible physical crash probabilities. Shaded bands are weekly identified intervals for the physical probability of $X \leq 0.90$, in percent, under the benchmark kernel and all retained OTM quotes. The cross sections cover exact horizons of 23, 30, and 37 days with 402, 372, and 324 cross sections. Missing weeks appear as gaps. Results use $\gamma = 2$ and $B = 5$.

5.3 Where the identifying information comes from

Figure 2 separates the level of the remaining ambiguity from the contribution of each quote layer. The largest change occurs when the six additional nearby puts join the put pair, the rest of the put wing remains informative, and OTM calls add little beyond the complete put wing (Table 1, Panel C). The identification of a 10 percent decline therefore comes mainly from put prices, with the strongest contribution near the threshold.

These are marginal contributions along the nested sequence, and a reverse exercise confirms the same concentration. Deleting the far put wing from the full quote set widens the median interval, deleting the OTM calls barely moves it, and deleting the six nearby puts leaves the median width unchanged, because neighboring wing strikes substitute for them. The quoted intervals themselves anchor the identification: doubling every bid–ask spread moves the endpoints an order of magnitude more than retaining only every other strike. Appendix C reports the magnitudes.

5.4 The unobserved right tail

The main specification places X in $[0, 5]$. Figure 3 compares that choice with tighter and looser support conditions on a fixed robustness subsample of 180 cross sections, 60 per horizon. The subsample is




Figure 2: Identification gains across nested quote sets. Panel A reports the median width of the identified crash probability interval under each quote set and horizon, in percentage points. Panel B reports the median proportional reduction relative to the preceding quote set, in percent. Results use $\gamma = 2$, $K = 0.90$, and $B = 5$.




Figure 3: Tail dependence of crash probability and conditional depth. The robustness subsample contains 180 cross sections, 60 at each horizon. The blue series reports the median ratio $q_{2,K}^-(B)/q_{2,K}^-(5)$, where $q_{2,K}^-(B)$ is the probability lower endpoint under support cap B . The gray dashed series reports the median ratio of conditional depth interval widths. Finite caps use all retained OTM quotes. Moving from $B = 2$ to $B = 10$ lowers a cross section's probability lower bound by a median of 0.0008 percentage points and by 0.0053 at the 95th percentile. The unbounded point reports the optimal values of the unbounded program: the probability lower endpoint is zero, and the largest depth width ratio relative to $B = 5$ is 1.0002. Results use $\gamma = 2$ and $K = 0.90$.

spaced evenly over the sample period and includes cross sections adjacent to the August 2015, February 2018, March 2020, and June 2022 stress dates. For finite caps between two and ten, the results are nearly unchanged.

Removing the cap acts on the probability bound alone. Without a support cap, the lower endpoint of the closed interval for $q_{2,K}$ is zero in every subsample cross section, as Proposition 3.6 predicts (Appendix A states the unbounded variant of the program). The conditional depth interval has no global-moment denominator, and its width barely moves.

6 Jointly Attainable Crash Scenarios

Each of the four marginal endpoints is attained by its own extremal measure, and nothing forces these measures to coincide, so a probability and shortfall pair drawn from the marginal rectangle may describe no single admissible distribution. This section examines the physical crash frontiers, first for selected cross sections and then for the full sample. Each plotted frontier follows the construction of Section 3.4,




Figure 4: Joint probability–shortfall sets across market states. Each panel overlays the envelopes of supporting half planes of $\Theta_{q\ell}(\mathcal{I})$ for the put pair, local eight, complete put wing, and all OTM quotes. Each successive envelope intersects all half planes accumulated through that layer, so the sets are nested by construction. Probability and unconditional shortfall are in percent, at a 30-day horizon with $\gamma = 2$, $K = 0.90$, and $B = 5$. The cross section dates are 26 August 2015, 4 December 2019, 25 March 2020, and 8 June 2022. The COVID-19 and 2022 targets were 18 March 2020 and 15 June 2022. Each panel uses its own axes.

with the direction grid refined until the envelope is accurate to within 1 percent of the cross section’s put pair scale, the marginal interval widths under the put pair (Appendix C).

Figure 4 covers four prespecified 30-day market states, each matched to the closest admissible cross section: the 2015 selloff, a calm market in late 2019, the COVID-19 crash, and the 2022 selloff. The identified sets contract in the same order in all four panels. Two puts leave a large range of probability–shortfall combinations, nearby puts remove most of that range, and the complete put wing and the OTM calls leave a much smaller set. The location of the surviving set tracks market conditions: the calm-market set lies close to the origin, and the 2015 and COVID-19 sets allow larger probability and shortfall values.

Figure 5 repeats the exercise across maturities around the COVID-19 crash. The 37-day set reaches the largest shortfall values among the three and has a broad probability range under sparse quotes. Additional puts contract the feasible region at every horizon, and the sets that remain once every OTM quote is included still differ in level and shape.

Across the 976 cross sections shared by the four quote sets, the joint restriction is strongest where shortfall is stressed. In the two directions that place positive weight on shortfall, the horizon-level median gaps between the full OTM joint set and its marginal rectangle run from 2.7 to 4.0 percent of the put pair scale (Appendix C.4). The same directional statistics show where each quote layer acts: nearby puts deliver most of the contraction in high-probability directions, while the rest of the put wing dominates in low-probability directions.




Figure 5: Joint probability–shortfall sets near the COVID-19 crash. Each panel overlays the envelopes of supporting half planes for the put pair, local eight, complete put wing, and all OTM quotes at one horizon. Probability and unconditional shortfall are in percent. The panels use the closest admissible 23-, 30-, and 37-day observations around the 18 March 2020 target, dated 11 March, 25 March, and 18 March 2020. Results use $\gamma = 2$, $K = 0.90$, and $B = 5$. The vertical scales differ across panels.

7 Conclusion

Finite option quotes are informative about short-horizon physical crash risk, and the physical crash frontier measures how informative. Mapping every quote-compatible distribution through the power utility investor’s pricing kernel separates what traded prices imply from what interpolation and tail assumptions add. In the S&P 500 cross sections studied here, most of the width of the put pair probability interval disappears once all retained OTM quotes are used, the largest contribution comes from puts near the threshold, and the sets rise and widen together in market stress.

What the quotes cannot supply is a positive floor on the physical crash probability. Under the $\gamma = 2$ benchmark that probability depends on a global second moment that traded quotes leave unrestricted far in the tail, so removing the support cap drives the lower bound to zero. The same happens for every investor more risk averse than the log investor (Remark 3.7). The denominator cancels from conditional depth, whose interval is stable across the support conditions we examine. A positive lower bound on crash probability is therefore a joint statement about traded prices and the maintained support condition.

The frontier turns this evidence into a discipline for scenario analysis. It retains only the probability and shortfall combinations that a single admissible distribution can generate or approach, and its exclusions are strongest where scenarios stress shortfall. Under the maintained kernel, event, and support conditions, any estimate that fills in the curve between and beyond the quoted strikes while pricing the same quotes within their spreads implies physical crash risk inside the frontier.

References

- Yacine Aït-Sahalia and Andrew W. Lo. Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94(1–2):9–51, 2000. doi: 10.1016/S0304-4076(99)00016-0.
- Farid Alizadeh and Donald Goldfarb. Second-order cone programming. *Mathematical Programming*, 95(1):3–51, 2003. doi: 10.1007/s10107-002-0339-5.
- David Backus, Mikhail Chernov, and Ian Martin. Disasters implied by equity index options. *Journal of Finance*, 66(6):1969–2012, 2011. doi: 10.1111/j.1540-6261.2011.01697.x.
- Shane Barratt, Jonathan Tuck, and Stephen Boyd. Convex optimization over risk-neutral probabilities. *Optimization and Engineering*, 25(1):283–299, 2024. doi: 10.1007/s11081-023-09802-2.
- Gilbert W. Bassett. Nonparametric bounds for the probability of future prices based on option values. In Yadolah Dodge, editor, *L1-Statistical Procedures and Related Topics*, volume 31 of *IMS Lecture Notes–Monograph Series*, pages 287–300. Institute of Mathematical Statistics, 1997. doi: 10.1214/lnms/1215454145.
- Arie Beresteanu and Francesca Molinari. Asymptotic properties for a class of partially identified models. *Econometrica*, 76(4):763–814, 2008. doi: 10.1111/j.1468-0262.2008.00859.x.
- Dimitris Bertsimas and Ioana Popescu. On the relation between option and stock prices: A convex optimization approach. *Operations Research*, 50(2):358–374, 2002. doi: 10.1287/opre.50.2.358.424.
- Robert R. Bliss and Nikolaos Panigirtzoglou. Option-implied risk aversion estimates. *Journal of Finance*, 59(1):407–446, 2004. doi: 10.1111/j.1540-6261.2004.00637.x.
- Tim Bollerslev and Viktor Todorov. Tails, fears, and risk premia. *Journal of Finance*, 66(6):2165–2211, 2011. doi: 10.1111/j.1540-6261.2011.01695.x.
- Tim Bollerslev, Viktor Todorov, and Lai Xu. Tail risk premia and return predictability. *Journal of Financial Economics*, 118(1):113–134, 2015. doi: 10.1016/j.jfineco.2015.02.010.
- Jaroslav Borovička, Lars Peter Hansen, and José A. Scheinkman. Misspecified recovery. *Journal of Finance*, 71(6):2493–2544, 2016. doi: 10.1111/jofi.12404.
- Douglas T. Breeden and Robert H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621–651, 1978. doi: 10.1086/296025.
- Abraham Charnes and William W. Cooper. Programming with linear fractional functionals. *Naval Research Logistics Quarterly*, 9(3–4):181–186, 1962. doi: 10.1002/nav.3800090303.
- Alexander M. G. Cox and Jan Obłój. Robust pricing and hedging of double no-touch options. *Finance and Stochastics*, 15(3):573–605, 2011. doi: 10.1007/s00780-011-0154-z.
- Mark H. A. Davis and David G. Hobson. The range of traded option prices. *Mathematical Finance*, 17(1):1–14, 2007. doi: 10.1111/j.1467-9965.2007.00291.x.
- William N. Goetzmann, Dasol Kim, and Robert J. Shiller. Crash beliefs from investor surveys. NBER Working Paper 22143, National Bureau of Economic Research, 2016.
- Lars Peter Hansen and José A. Scheinkman. Long-term risk: An operator approach. *Econometrica*, 77(1):177–234, 2009. doi: 10.3982/ECTA6761.
- Jens Carsten Jackwerth. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, 13(2):433–451, 2000. doi: 10.1093/rfs/13.2.433.

- Alan J. King, Matti Koivu, and Teemu Pennanen. Calibrated option bounds. *International Journal of Theoretical and Applied Finance*, 8(2):141–159, 2005. doi: 10.1142/S0219024905002925.
- Mark G. Krein and Adolf A. Nudelman. *The Markov Moment Problem and Extremal Problems*, volume 50 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, 1977.
- Bryan Liang, Marcel Nutz, Shunan Sheng, and Valentin Tissot-Daguette. Bid–ask martingale optimal transport, 2026. URL <https://arxiv.org/abs/2603.24605>.
- Andrew W. Lo. Semi-parametric upper bounds for option prices and expected payoffs. *Journal of Financial Economics*, 19(2):373–387, 1987. doi: 10.1016/0304-405X(87)90010-9.
- Charles F. Manski. *Partial Identification of Probability Distributions*. Springer Series in Statistics. Springer, New York, 2003. doi: 10.1007/b97478.
- Ian Martin. What is the expected return on the market? *Quarterly Journal of Economics*, 132(1):367–433, 2017. doi: 10.1093/qje/qjw034.
- Ian Martin and Ran Shi. Forecasting crashes with a smile. CEPR Discussion Paper 21236, Centre for Economic Policy Research, March 2026. URL <https://cepr.org/publications/dp21236>.
- Jun Pan. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics*, 63(1):3–50, 2002. doi: 10.1016/S0304-405X(01)00088-5.
- Dimitris N. Politis and Joseph P. Romano. A circular block-resampling procedure for stationary data. In Raoul LePage and Lynne Billard, editors, *Exploring the Limits of Bootstrap*, pages 263–270. Wiley, New York, 1992.
- R. Tyrrell Rockafellar. *Convex Analysis*. Number 28 in Princeton Mathematical Series. Princeton University Press, Princeton, 1970.
- Joshua V. Rosenberg and Robert F. Engle. Empirical pricing kernels. *Journal of Financial Economics*, 64(3):341–372, 2002. doi: 10.1016/S0304-405X(02)00128-9.
- Stephen A. Ross. The recovery theorem. *Journal of Finance*, 70(2):615–648, 2015. doi: 10.1111/jofi.12092.
- Olimjon Sh. Sharipov and Martin Wendler. Normal limits, nonnormal limits, and the bootstrap for quantiles of dependent data. *Statistics & Probability Letters*, 83(4):1028–1035, 2013. doi: 10.1016/j.spl.2012.12.017.
- Wouter J. E. C. van Eekelen. A generalized moment approach to sharp bounds for conditional expectations, 2023. URL <https://arxiv.org/abs/2401.00090>.
- Jirong Zhuang. The risk-neutral crash frontier: Sharp joint bounds on crash probability and conditional depth from option bid–ask quotes, 2026. URL <https://arxiv.org/abs/2607.25353>.

A Proofs and Complete Programs

A.1 Positive-return and threshold closure

Lemma A.1 (Approximation away from zero). *Under Assumption 3.3, every $\mathbb{Q} \in \mathcal{M}(\mathcal{I})$ is the weak limit of quote-compatible measures supported on $(0, B]$. For every fixed $r > 0$, their global and truncated r th moments converge to those of $\mathbb{Q}$.*

Proof. Form the mixture $(1 - \delta_\varepsilon)\mathbb{Q} + \delta_\varepsilon\mathbb{Q}^\circ$, where $\delta_\varepsilon \rightarrow 0$ and $\varepsilon/\delta_\varepsilon \rightarrow 0$, and move all mass at zero in the mixture to $\varepsilon > 0$. This preserves total mass and changes the mean and each retained option value by $O(\varepsilon)$. A transfer of order $O(\varepsilon)$ between U_+ and U_- restores the mean; the mass available in these intervals is of order δ_ε , so the transfer preserves nonnegativity for small ε . Strict quote slack inherited from the mixture absorbs the remaining payoff changes.

The resulting measures have strictly positive support. On $[0, B]$, the functions x^r are uniformly continuous for every $r > 0$, and the shifted zero-state contribution is $O(\varepsilon^r)$, so the global and truncated moments converge. $\square$

A.2 Proof of Theorem 3.2

Convexity and compactness of the joint set. For an admissible measure $\mathbb{Q}_j$, let $d_j = M_\gamma(\mathbb{Q}_j) > 0$, and set

$$\mathbf{n}_j = \begin{pmatrix} A_\gamma(\mathbb{Q}_j) \\ KA_\gamma(\mathbb{Q}_j) - A_{\gamma+1}(\mathbb{Q}_j) \end{pmatrix},$$

so that the probability–shortfall image of $\mathbb{Q}_j$ is $\mathbf{w}_j = \mathbf{n}_j/d_j$. Fix $\theta \in [0, 1]$ and define

$$\alpha = \frac{\theta d_1}{\theta d_1 + (1 - \theta)d_0}.$$

The mixture $\mathbb{Q}_\alpha = \alpha\mathbb{Q}_0 + (1 - \alpha)\mathbb{Q}_1$ is admissible because the quote and moment restrictions are linear in the measure, and

$$\frac{\alpha\mathbf{n}_0 + (1 - \alpha)\mathbf{n}_1}{\alpha d_0 + (1 - \alpha)d_1} = \theta\mathbf{w}_0 + (1 - \theta)\mathbf{w}_1.$$

The attainable image is therefore convex, and its closure is convex and nonempty. Since $\ell_{\gamma,K} = q_{\gamma,K}s_{\gamma,K}$ with $s_{\gamma,K} \in [0, K]$ whenever $q_{\gamma,K} > 0$, the closure is contained in $\{(q, \ell) : 0 \leq q \leq 1, 0 \leq \ell \leq Kq\}$. It is closed and bounded, hence compact. $\square$

The conditional depth interval. Let $\mathcal{M}_+ = \{\mathbb{Q} \in \mathcal{M}(\mathcal{I}) : A_\gamma(\mathbb{Q}) > 0\}$, a convex set. For $\mathbb{Q}_0, \mathbb{Q}_1 \in \mathcal{M}_+$ and $\alpha \in [0, 1]$, the mixture $\mathbb{Q}_\alpha = \alpha\mathbb{Q}_0 + (1 - \alpha)\mathbb{Q}_1$ satisfies

$$s_{\gamma,K}(\mathbb{Q}_\alpha) = \frac{\alpha A_\gamma(\mathbb{Q}_0)s_{\gamma,K}(\mathbb{Q}_0) + (1 - \alpha)A_\gamma(\mathbb{Q}_1)s_{\gamma,K}(\mathbb{Q}_1)}{\alpha A_\gamma(\mathbb{Q}_0) + (1 - \alpha)A_\gamma(\mathbb{Q}_1)}.$$

As α varies, this expression spans the interval between $s_{\gamma,K}(\mathbb{Q}_0)$ and $s_{\gamma,K}(\mathbb{Q}_1)$. The image is therefore an interval contained in $[0, K]$, and its closure is a nonempty compact interval, which is $\Theta_s^\gamma(\mathcal{I})$. $\square$

A.3 Threshold bookkeeping and the depth programs

Recall that j_K indexes the cell beginning at K . In the program of Section 3.4, the moment conditions on this cell apply to the residual moments

$$\bar{y}_{0,j_K} = y_{0,j_K} - \tilde{a}_K, \quad \bar{y}_{1,j_K} = y_{1,j_K} - K\tilde{a}_K, \quad \bar{y}_{2,j_K} = y_{2,j_K} - K^2\tilde{a}_K,$$

which satisfy (17)–(18) with $a = K$ and $b = b_{j_K+1}$, together with $\tilde{a}_K \geq 0$. The atom is part of y_{r,j_K} , so it enters the global constraints (14) and the quote values (15) without further adjustment, while the targets (19) count it as event mass. An empty sum in (15) is interpreted as zero when $k_i \geq B$. When the support cap is removed, the terminal cell becomes a half line and keeps the positive semidefinite moment condition and the lower location inequality, while the upper support inequalities are dropped.

Conditional depth uses a second normalization. When $A_2(\mathbb{Q}) > 0$, define $\lambda = A_2(\mathbb{Q})^{-1}$ and $d\widehat{\mathbb{Q}} = \lambda d\mathbb{Q}$. The mass, forward, and truncated second-moment restrictions become

$$\int d\widehat{\mathbb{Q}} = \lambda, \quad \int X d\widehat{\mathbb{Q}} = \lambda, \quad \int_{[0,K]} X^2 d\widehat{\mathbb{Q}} = 1, \quad (22)$$

each quote band is multiplied by λ , and the objective is linear:

$$s_{2,K} = \int_{[0,K]} (K - X) X^2 d\widehat{\mathbb{Q}}. \quad (23)$$

With cell moments $\widehat{y}_{r,j} = \int x^r d\widehat{\mu}_j$ taken from a local decomposition of $\widehat{\mathbb{Q}}$, and with $\widehat{a}_K \geq 0$ the scaled event atom, the finite constraints are

$$\sum_j \widehat{y}_{0j} = \lambda, \quad \sum_j \widehat{y}_{1j} = \lambda, \quad \sum_{j \in \mathcal{J}_-} \widehat{y}_{2j} + K^2 \widehat{a}_K = 1, \quad (24)$$

$$\widehat{C}_i(\widehat{\mathbf{y}}) = \sum_{j: b_j \geq k_i} (\widehat{y}_{1j} - k_i \widehat{y}_{0j}), \quad \lambda \underline{c}_i \leq \widehat{C}_i(\widehat{\mathbf{y}}) \leq \lambda \bar{c}_i, \quad (25)$$

with the same cellwise moment cones and threshold treatment as before, and the linear objective

$$s_{2,K} = \sum_{j \in \mathcal{J}_-} (K \widehat{y}_{2j} - \widehat{y}_{3j}). \quad (26)$$

No global second-moment normalization is imposed.

A.4 Proof of Corollary 3.4

Proof. The feasible region is compact. It is closed, being cut out by finitely many linear equalities and inequalities and by second-order cone constraints. For boundedness, fix a feasible point. The cellwise moment conditions are sufficient, so the cell moments, net of the designated atom on the threshold cell, are moments of nonnegative measures on the closed cells. Summing the cells and adding the atom yields a nonnegative measure on $[0, B]$ with total mass ρ , first moment ρ , and second moment one. The Cauchy–Schwarz inequality gives $\rho^2 \leq \rho \cdot 1$, so $\rho \leq 1$, and the bounded support gives $0 \leq y_{r,j} \leq B^r y_{0j} \leq B^r$ for every stored order together with $0 \leq \widetilde{a}_K \leq y_{0,j_K} \leq 1$. Every coordinate of a feasible point is therefore bounded.

The conic image of the feasible region under the linear targets (19) is therefore compact. We show that every attainable pair lies in it, so that, being closed, it contains $\Theta_{q\ell}(\mathcal{I})$, and that every point of it is a limit of attainable pairs, so that it is contained in $\Theta_{q\ell}(\mathcal{I})$.

First take any admissible $\mathbb{Q}$ and set $\rho = M(\mathbb{Q})^{-1} \in [1/B, 1]$ and $d\widetilde{\mathbb{Q}} = \rho d\mathbb{Q}$. Decompose $\widetilde{\mathbb{Q}}$ into local components, assigning the atom at K , if any, to the cell beginning at K , and let $\widetilde{a}_K$ be that scaled atom. The resulting $(\mathbf{y}, \rho, \widetilde{a}_K)$ satisfies every constraint: the global equalities are (12), the quote inequalities are the scaled quote bands, the cells below K carry genuine moment sequences of measures on closed intervals, and the threshold cell satisfies its conditions net of the designated atom. The targets (19) equal $(q_{2,K}(\mathbb{Q}), \ell_{2,K}(\mathbb{Q}))$ by (8).

Conversely, take any feasible $(\mathbf{y}, \rho, \widetilde{a}_K)$. The truncated Hausdorff conditions (16)–(18) supply, cell by cell, a nonnegative measure on the closed cell with the prescribed moments. Adding the designated atom $\widetilde{a}_K$ at K and summing the cells produces a nonnegative measure $\widetilde{\mathbb{Q}}'$ that satisfies the mass, forward, second-moment, and quote restrictions, since each restriction is an integral of a continuous function against the cell moments. Where a reconstructed cell measure places mass at a shared endpoint $b_j \neq K$, the location is unambiguous relative to the threshold, so neither the constraints nor the targets depend on which cell carries it. The support inequalities give $1 = \sum_j y_{2j} \leq B \sum_j y_{1j} = B\rho$, so $\rho \geq 1/B > 0$, and the measure $d\mathbb{Q}' = d\widetilde{\mathbb{Q}}'/\rho$ has mass one, mean one, second moment $1/\rho$, and satisfies every quote interval.

One discrepancy can remain. The reconstruction of the cell beginning at K may itself place mass at K beyond the designated atom, whereas the targets (19) treat that residual mass as lying above the threshold. Move the residual mass to $K + \varepsilon$ and repair the measure with the mixing construction of Lemma A.1: blend in weight δ_ε of the interior measure $\mathbb{Q}^\circ$, transfer order ε between U_- and U_+ to restore the mean, and use the inherited strict quote slack to keep every quote value inside its interval for small ε . The resulting measures $\mathbb{Q}_\varepsilon$ are admissible, and their probability and shortfall converge to the targets of the

feasible point. Every feasible point is therefore attained or approached, so the conic image equals $\Theta_{q\ell}(I)$, and the support values follow by linear optimization over it.

For the depth programs, let S denote the set of attainable conditional depths and V the set of objective values of feasible points. Any admissible $\mathbb{Q}$ with $A_2(\mathbb{Q}) > 0$ produces a feasible pair $(\hat{y}, \lambda)$ with $\lambda = A_2(\mathbb{Q})^{-1}$ whose objective value is $s_{2,K}(\mathbb{Q})$, so $S \subseteq V$. Conversely, every feasible point carries a finite scalar λ , which the unit constraint in (24) forces to be positive. If its threshold-cell reconstruction carries no residual mass at K , dividing $\hat{\mathbb{Q}}'$ by λ yields an admissible measure with $A_2 = 1/\lambda$ whose conditional depth equals the objective (26); in general, the same $K + \varepsilon$ construction yields admissible measures whose depths converge to the objective value. Every value in V is therefore a limit of values in S , and

$$S \subseteq V \subseteq \text{cl } S = \Theta_s(I).$$

Suprema and infima agree across this chain, so the optimal values of the two conic programs equal the endpoints of the compact interval $\Theta_s(I)$. These optimal values need not be attained by feasible points: when $A_2 \downarrow 0$ along a sequence of admissible measures, λ diverges, and no feasible point need exist whose objective value equals the endpoint. $\square$

A.5 Proof of Proposition 3.6

Proof. Let ν_- and ν_+ be the conditional distributions of $\mathbb{Q}^\circ$ on U_- and U_+ , with means $\mu_- < 1 < \mu_+$. For H above the retained strikes and U_+ , set

$$\varepsilon_H = H^{-3/2}, \quad a_H = \varepsilon_H \frac{H - \mu_+}{\mu_+ - \mu_-}, \quad b_H = \varepsilon_H \frac{H - \mu_-}{\mu_+ - \mu_-}.$$

Let δ_H denote unit mass at H , and define

$$\mathbb{Q}_H = \mathbb{Q}^\circ + \varepsilon_H \delta_H + a_H \nu_- - b_H \nu_+.$$

The coefficients satisfy $\varepsilon_H + a_H - b_H = 0$ and $\varepsilon_H H + a_H \mu_- - b_H \mu_+ = 0$, so $\mathbb{Q}_H$ has mass one and mean one. Since $b_H \rightarrow 0$ and $\mathbb{Q}^\circ(U_+) > 0$, the measure is nonnegative for all sufficiently large H .

For a retained call payoff $h_i(x) = (x - k_i)^+$,

$$\mathbb{E}_{\mathbb{Q}_H}[h_i(X)] - \mathbb{E}_{\mathbb{Q}^\circ}[h_i(X)] = \varepsilon_H h_i(H) + a_H \mathbb{E}_{\nu_-}[h_i(X)] - b_H \mathbb{E}_{\nu_+}[h_i(X)] = O(H^{-1/2}).$$

There are finitely many quotes, so strict interior slack implies that $\mathbb{Q}_H$ satisfies every quote interval for all sufficiently large H . The intervals U_- and U_+ are bounded, so

$$M(\mathbb{Q}_H) = M(\mathbb{Q}^\circ) + \varepsilon_H H^2 + O(H^{-1/2}) = M(\mathbb{Q}^\circ) + H^{1/2} + O(H^{-1/2}) \rightarrow \infty.$$

Finally, $A_2(\mathbb{Q}_H) \leq K^2$, and since $K - X \leq K$ on the event, also $KA_2(\mathbb{Q}_H) - A_3(\mathbb{Q}_H) \leq KA_2(\mathbb{Q}_H) \leq K^3$. Hence

$$0 \leq q_{2,K}(\mathbb{Q}_H) \leq \frac{K^2}{M(\mathbb{Q}_H)} \rightarrow 0, \quad 0 \leq \ell_{2,K}(\mathbb{Q}_H) \leq \frac{K^3}{M(\mathbb{Q}_H)} \rightarrow 0,$$

so the probability–shortfall pair of $\mathbb{Q}_H$ converges to the origin, which therefore lies in the closure of the joint image. $\square$

B The Log-Investor Benchmark

B.1 Mathematical specialization

When $\gamma = 1$, the forward restriction fixes the denominator of the transform (4) at $M_1(\mathbb{Q}) = \mathbb{E}_{\mathbb{Q}}[X] = 1$, so the dilution channel of Proposition 3.6 is absent and the tail acts only through the quote and forward

restrictions (Remark 3.7). The probability and unconditional shortfall are therefore linear in the unscaled measure,

$$q_{1,K}(\mathbb{Q}) = A_1(\mathbb{Q}), \quad \ell_{1,K}(\mathbb{Q}) = KA_1(\mathbb{Q}) - A_2(\mathbb{Q}), \quad (27)$$

and, when $A_1(\mathbb{Q}) > 0$, conditional depth is $s_{1,K}(\mathbb{Q}) = K - A_2(\mathbb{Q})/A_1(\mathbb{Q})$. No change of scale is needed for the joint set. Use the same cells as Section 3.4 and a local decomposition $\mathbb{Q} = \sum_j \mu_j$, with strike-endpoint atoms assigned to either adjacent cell and any atom at K assigned to the cell beginning at K . For each μ_j , define the cell mass z_j and first moment m_j , together with the second moment v_j on the cells in $\mathcal{J}_-$. The global restrictions are

$$\sum_j z_j = 1, \quad \sum_j m_j = 1. \quad (28)$$

Cells below the threshold carry the complete order-two conditions (17)–(18) in (z_j, m_j, v_j) , and cells at or above K need only $az_j \leq m_j \leq bz_j$. With j_K again indexing the cell beginning at K and a_K the threshold mass assigned to the closed event,

$$0 \leq a_K \leq z_{j_K}, \quad m_{j_K} \leq b_{j_K+1}z_{j_K} - (b_{j_K+1} - K)a_K, \quad (29)$$

each quote value is linear as in (15) with (z_j, m_j) in place of (y_{0j}, y_{1j}) and unscaled bands, and the targets are

$$q_{1,K} = \sum_{j \in \mathcal{J}_-} m_j + Ka_K, \quad \ell_{1,K} = \sum_{j \in \mathcal{J}_-} (Km_j - v_j). \quad (30)$$

Under Assumption 3.3, the closure argument of Appendix A.4 identifies this feasible region with the closed log-investor joint set. Conditional depth again uses a second normalization: with $\lambda = A_1(\mathbb{Q})^{-1}$ and $d\widehat{\mathbb{Q}} = \lambda d\mathbb{Q}$, the mass and forward integrals both equal λ , the unit constraint is $\int_{[0,K]} X d\widehat{\mathbb{Q}} = 1$, and the objective $\int_{[0,K]} (K - X)X d\widehat{\mathbb{Q}}$ is linear. The scaled quote bands, cell moment conditions, and threshold-atom treatment carry over from Appendix A.3, with first and second moments in the roles that second and third moments play there.

B.2 Comparison with the $\gamma = 2$ kernel

Proposition B.1 (Kernel comparison). *Fix $K \in (0, 1)$ and consider a common quote-compatible domain of unit-mean measures with finite second moments. For every $\mathbb{Q}$ in this domain with positive truncated moments,*

$$q_{2,K}(\mathbb{Q}) \leq Kq_{1,K}(\mathbb{Q}) \leq q_{1,K}(\mathbb{Q}), \quad \ell_{2,K}(\mathbb{Q}) \leq K\ell_{1,K}(\mathbb{Q}), \quad s_{2,K}(\mathbb{Q}) \leq s_{1,K}(\mathbb{Q}). \quad (31)$$

Proof. On the event $X \leq K$, $X^2 \leq KX$. Also $M(\mathbb{Q}) \geq M_1(\mathbb{Q})^2 = 1$. Hence

$$q_{2,K} = \frac{A_2}{M(\mathbb{Q})} \leq A_2 \leq KA_1 = Kq_{1,K} \leq q_{1,K}.$$

The same argument applied to $(K - X)X^2 \leq K(K - X)X$ gives $\ell_{2,K} \leq K\ell_{1,K}$. The truncated moment sequence is log-convex by the Cauchy–Schwarz inequality, so $A_2^2 \leq A_1A_3$. Dividing by $A_1A_2 > 0$ and subtracting from K gives $s_{2,K} \leq s_{1,K}$. $\square$

Taking extrema in (31) compares corresponding endpoints: each $\gamma = 2$ probability and shortfall endpoint lies below K times its $\gamma = 1$ counterpart, and each depth endpoint lies below its counterpart on the common positive-probability domain. The inequalities impose no ordering across different endpoints, and the optimizing measures can differ.

Table 2: Kernel comparison on common cross sections

Horizon	Cross sections	q width		ℓ width		s width	
		$\gamma = 1$	$\gamma = 2$	$\gamma = 1$	$\gamma = 2$	$\gamma = 1$	$\gamma = 2$
23 days	405	1.54	1.40	0.017	0.013	4.93	4.81
30 days	369	2.08	1.88	0.024	0.017	4.27	4.10
37 days	332	2.57	2.35	0.033	0.023	4.54	4.26

Notes: Median interval widths in percentage points, computed on the cross sections where both kernels solve, under all retained OTM quotes, $K = 0.90$, $B = 5$, and each cross section's fixed discounting and forward.

Table 3: Identified sets for a 15 percent decline

Horizon	Cross sections	Median $q_{2,K}$	q width	Depth cross sections	s width
23 days	404	[0.14, 0.70]	0.53	393	9.98
30 days	373	[0.33, 1.08]	0.76	362	7.70
37 days	330	[0.47, 1.61]	1.07	329	8.07

Notes: All retained OTM quotes, $K = 0.85$, $B = 5$. Probability entries are medians of interval endpoints across cross sections, in percent. Width entries are median interval widths, in percentage points. Depth widths use the cross sections whose closed probability interval excludes zero.

B.3 Full-sample comparison

Table 2 reports both kernels on the same cross sections, the same event $K = 0.90$, the same support cap $B = 5$, and the same full OTM quote sets. The sample keeps the cross sections where both kernels' programs solve, a criterion that differs from the verification behind Table 1, so the counts differ slightly. The $\gamma = 2$ intervals are narrower for every coordinate at every horizon. Proposition B.1 orders endpoint levels, not widths, so the width ordering is an empirical regularity, and the two kernels agree on how identification varies across horizons and coordinates.

C Supplementary Evidence

C.1 A more extreme threshold

Table 3 repeats the analysis under the full OTM set at $K = 0.85$. The probability intervals are narrower than at $K = 0.90$ and the conditional depth intervals are wider: a more extreme event can have a narrow probability range and a wide range of depth conditional on occurring. The quote layers contribute in the same order as at the main threshold. Adding the six nearby puts cuts the median put pair widths by 70 to 75 percent, and the complete put wing and the calls deliver the remaining contraction to the widths in the table.

C.2 Quote perturbations and deletions

Two exercises alter the information content of the quotes on a prespecified subset of 24 cross sections from the robustness subsample. Doubling every bid–ask spread around its midpoint changes a probability endpoint by a median of 0.79 percentage points and a conditional depth endpoint by a median of 1.84 percentage points. Retaining every other strike changes the corresponding endpoints by medians of 0.073 and 0.122 percentage points.

A reverse decomposition deletes groups of quotes from the full OTM set on the robustness subsample of 180 cross sections. Removing the put wing beyond the local eight widens the median probability interval by 0.46, 0.76, and 0.52 percentage points at the three horizons. Removing all OTM calls widens it by 0.05, 0.08, and 0.13 percentage points. Removing the six nearby puts while keeping every other quote

Table 4: Verification of the regularity condition by cross section

Horizon	Cross sections	Verified	Rate (%)	Median slack ($\times 10^{-5}$)
23 days	406	405	99.8	1.1
30 days	376	374	99.5	1.1
37 days	333	332	99.7	1.5

Notes: A cross section is verified when an admissible measure prices every retained quote strictly inside its bid–ask interval and places mass on both sides of the forward away from strike and threshold boundaries. Slack is the smallest distance from a constructed option value to its quote endpoints, in normalized price units.

leaves the median width unchanged, because the neighboring wing strikes substitute for them. Marginal contributions therefore depend on the base set: nearby puts carry most of the information relative to the sparse pair, and the wing carries information that survives even when the local strikes are dense.

C.3 Verification of the regularity condition

Table 4 summarizes the verification of Assumption 3.3 in each cross section, described in Section 4.3. The verification fails only in the 23- and 37-day cross sections of 20 January 2016 and the 30-day cross sections of 31 May 2017 and 7 September 2022.

C.4 Directional joint restrictions

Let q_{pair}^- and ℓ_{pair}^- be the put pair lower endpoints of a cross section, and let w_q^{pair} and w_ℓ^{pair} be the positive put pair widths. Joint comparisons use the normalized coordinates

$$\bar{q} = \frac{q - q_{\text{pair}}^-}{w_q^{\text{pair}}}, \quad \bar{\ell} = \frac{\ell - \ell_{\text{pair}}^-}{w_\ell^{\text{pair}}}. \quad (32)$$

Let $\bar{\Theta}$ be the normalized joint set and $\bar{\mathcal{B}}$ its marginal rectangle, with support functions $H_{\bar{\Theta}}$ and $H_{\bar{\mathcal{B}}}$ as in (11). For the four diagonal directions

$$u \in \left\{ \frac{(1,1)}{\sqrt{2}}, \frac{(1,-1)}{\sqrt{2}}, \frac{(-1,1)}{\sqrt{2}}, \frac{(-1,-1)}{\sqrt{2}} \right\},$$

the normalized support gap and the directional contraction between nested quote sets $\mathcal{I}_1 \subseteq \mathcal{I}_2$ are

$$g(u) = 100 [H_{\bar{\mathcal{B}}}(u) - H_{\bar{\Theta}}(u)], \quad c_{1 \rightarrow 2}(u) = 100 [H_{\bar{\Theta}(\mathcal{I}_1)}(u) - H_{\bar{\Theta}(\mathcal{I}_2)}(u)]. \quad (33)$$

Figure 6 summarizes these statistics on the 976 cross sections shared by the four quote sets. The gap between the joint set and its marginal rectangle concentrates in the two directions that place positive weight on shortfall, where the horizon-level median gaps run from 2.7 to 4.0 percent of the put pair scale. Pooling all four directions, the median gap under the full OTM set is 0.52 percent. The contraction statistics show where each layer acts. Moving from the pair to the local eight produces the largest contractions in high-probability directions, with medians of 33 and 51 percent of the pair scale. Moving from the local eight to the complete put wing dominates in low-probability directions, with medians of 8 and 43 percent. The median contraction from the put pair to the full OTM set across the four directions is 50 percent of the pair scale.

C.5 Accuracy of the plotted joint sets

The joint sets of Section 6 are drawn as intersections of supporting half planes on an adaptive direction grid. The grid doubles from 32 to at most 128 directions until the gap to the inner hull of support points falls below 1 percent of the put pair scale. Table 5 reports the realized accuracy for each plotted cross section: the largest normalized gap across the four quote sets is at most 0.99 percent, so every plotted envelope is visually indistinguishable from the identified set at the resolution of the figures.




Figure 6: Directional joint restrictions and quote-set contraction. Panel A reports the median normalized support gap between the full OTM joint set and its marginal rectangle in four prespecified directions. Panel B reports the median support contraction contributed by adjacent quote layers in the same directions. Both panels use each cross section's put pair marginal widths for normalization and cover the 976 cross sections shared by the four quote sets. Results use $\gamma = 2$, $K = 0.90$, and $B = 5$.

Table 5: Adaptive-grid accuracy of the plotted joint sets

Panel	Date	Horizon	Directions	Max gap (%)
2015 selloff	26 Aug 2015	30 days	128	0.80
Calm market	4 Dec 2019	30 days	128	0.93
COVID-19 crash	25 Mar 2020	30 days	128	0.90
2022 selloff	8 Jun 2022	30 days	128	0.98
COVID-19, 23 days	11 Mar 2020	23 days	128	0.76
COVID-19, 37 days	18 Mar 2020	37 days	64	0.99

Notes: For each plotted cross section, Directions is the largest adaptive grid level used across the four quote sets, and Max gap is the largest normalized support gap between the plotted outer polygon and the inner hull of its support points, in percent of the put pair scale. The 30-day COVID-19 cross section appears in both Figures 4 and 5.