

7498464-the-risk-neutral-crash-frontier-sharp-joint-bounds-on-crash-probability-and-cond

SUMMARY: The article discusses the risk-neutral crash frontier, a concept used to characterize the set of probability and loss pairs that one distribution can generate while pricing every quote inside its spread, with depth as their ratio. The authors define π and ℓ as the probability of a crash and the normalized loss of a put price, respectively. They show that the market prices the product of likelihood and depth, not the factors themselves. The authors characterize the set of probability and loss pairs that one distribution can generate and call its boundary the risk-neutral crash frontier. They use linear programming to trace the frontier and price any portfolio of digital and put payoffs sharply, each bound certified by a static super-replicating portfolio of cash, forward, and quoted options. The authors apply their method to SPX quotes sampled weekly from January 2013 through August 2023.

METHOD: The study uses weekly SPX cross sections from 2013 to 2023 to analyze the risk-neutral crash frontier. The authors define the probability and loss pairs based on the risk-neutral measure Q and the terminal index level divided by its forward. They use linear programming to trace the frontier and price any portfolio of digital and put payoffs sharply, each bound certified by a static super-replicating portfolio of cash, forward, and quoted options. They apply their method to SPX quotes sampled weekly from January 2013 through August 2023.

KEY FINDINGS:

- The risk-neutral crash frontier characterizes the set of probability and loss pairs that one distribution can generate.
- The frontier can be traced by linear programming without discretizing the state space.
- The depth frontier and its area follow in closed form.
- The method can be applied to analyze the risk-neutral crash frontier using SPX quotes sampled weekly from January 2013 through August 2023.

IMPLICATION FOR TRADING: The risk-neutral crash frontier provides a framework for understanding the relationship between crash probability and loss, which can be used by risk committees to stress test particular scenarios. Practitioners can use the method to price and hedge claims that pay only in a crash, such as put options. The method can also be applied to analyze the risk-neutral crash frontier using SPX quotes sampled weekly from January 2013 through August 2023. This can help traders and risk committees better understand the market's perception of crash risk and make more informed decisions.

The Risk-Neutral Crash Frontier:

Sharp Joint Bounds on Crash Probability and Conditional Depth from Option Bid–Ask Quotes

Jirong Zhuang

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

September 2026

Abstract

Index put prices are the market's quotes for crash insurance, and a put's value equals the probability of a crash times the expected shortfall given one. The market therefore prices the product of likelihood and depth, not the factors, and finitely many bid and ask quotes leave a range of ways to split it. Fitting one density hides that range, and bounds computed one factor at a time can combine into scenarios that no single risk-neutral distribution could produce. We characterize the set of probability and loss pairs that one distribution can generate while pricing every quote inside its spread, with depth as their ratio, and call its boundary the risk-neutral crash frontier. Partitioning the state space at the quoted strikes and the threshold, with one coordinate for mass at the threshold, makes the set the exact projection of a finite linear system, with no price grid. Linear programs trace the frontier and price any portfolio of digital and put payoffs sharply, each bound certified by a static super-replicating portfolio of cash, forward, and quoted options. In weekly SPX cross sections from 2013 to 2023, a median 37% of the scenarios that separate bounds admit are jointly infeasible. Extending the quote set from the eight strikes nearest the threshold to the complete put wing shrinks it by a further 5.4 to 18.2%.

Keywords: risk-neutral crash frontier, option-implied tail risk, partial identification, model-free bounds, conditional depth, linear programming

JEL Classification: G13, C14, C61

1 Introduction

Index put prices are the market's quotes for crash insurance. Dealers use them to price and hedge claims that pay only in a crash, and a large literature reads crash probabilities and tail expectations from them (Backus et al., 2011; Bollerslev and Todorov, 2011). A put price, however, answers two questions at once: how likely is a crash, and how deep will it be if it comes? Let X be the terminal index level divided by its forward, fix a threshold $K < 1$, and under a risk-neutral measure $\mathbb{Q}$ define

$$p_K = \mathbb{Q}(X \leq K), \quad L_K = \mathbb{E}_{\mathbb{Q}}(K - X)^+, \quad s_K = \frac{L_K}{p_K} = K - \mathbb{E}_{\mathbb{Q}}[X \mid X \leq K], \quad p_K > 0.$$

The loss L_K is the normalized put value. It averages the shortfall over all states, including those in which the index never falls below K , whereas the depth s_K conditions on the crash and asks how far below K the index lands on average. The identity $L_K = p_K s_K$ says that the market prices the product of the two and not the factors.

A finite set of strikes quoted at bid and ask pins down neither factor. A frequent shallow breach and a rare deep one can carry the same price: a put worth 0.10% of forward is consistent with a 2% crash probability and a 5% average shortfall, and equally with a 0.5% probability and a 20% shortfall. The usual response fits one smooth density through the quotes and reports one crash probability, a number

that blends what the market says with what the fitting method adds. Bounding probability and depth one at a time keeps the quotes but throws the trade-off away. The distribution that makes a crash least likely need not be the one that makes it deepest. A stress scenario that pairs the two extremes can therefore be one that the quotes rule out, and a price range for a crash-contingent claim built from separate bounds is wider than the market supports.

We therefore study the two quantities together. Let $\mathcal{M}$ be the set of distributions on $\mathbb{R}_+$ with unit mean that price every retained quote inside its interval, and let

$$\mathcal{R}_K = \overline{\{(\mathbb{Q}(X \leq K), \mathbb{E}_{\mathbb{Q}}(K - X)^+) : \mathbb{Q} \in \mathcal{M}\}}$$

be the closure of the probability and loss pairs these distributions generate. We call the boundary of $\mathcal{R}_K$ the risk-neutral crash frontier, or simply the crash frontier. We use loss rather than depth as the second coordinate because $\mathcal{R}_K$ is convex while its depth image need not be; since each pair (p, L) with $p > 0$ maps to $(p, L/p)$, the depth frontier follows exactly once the boundary of $\mathcal{R}_K$ is known.

The joint set answers questions that marginal intervals cannot. A desk that holds a digital option together with a put can ask for the sharp range of the package’s value, which is a support function of the set. Every point of the frontier is backed by a dual solution that reads as a static portfolio of cash, forward, and quoted options whose cost equals the bound and whose payoff dominates the target, so each bound is a price at which the quotes can be traded. A risk committee that wants to stress a particular scenario, say a 2% crash probability with a 20% average shortfall, can check whether any distribution that prices today’s quotes produces it; the answer is whether the point lies inside the set. Comparing the frontier across nested sets of strikes shows whether the deep out-of-the-money put wing carries information about the tail beyond what strikes near K already say. None of this requires a smooth density or a parametric tail.

The technical difficulty is that the crash indicator jumps at K while every traded payoff is continuous there. Once the observed strikes and the threshold partition the state space into cells, the probability and first moment of each cell price every retained option exactly, but they cannot tell mass at K from mass just above it, and the indicator treats the two differently. We add one coordinate that records the mass at K . Under a strict interior condition on the quote constraints, the resulting finite system of linear inequalities projects exactly onto $\mathcal{R}_K$, so the frontier can be traced by linear programming without discretizing the state space. An adaptive algorithm recovers the polygon from finitely many support directions, and the depth frontier and its area then follow in closed form.

We apply the method to SPX quotes sampled weekly from January 2013 through August 2023, expanding the quote set from two strikes near the threshold to the complete put wing. To measure what joint identification adds, we compare the joint set with the region that separate marginal bounds allow once the universal restriction $0 \leq L \leq Kp$ is imposed. In the primary sample the median ratio of the joint area to the area of that region is 0.6340, so about a third of the scenarios that pass the marginal bounds cannot be produced by any single distribution. Moving from the eight strikes nearest the threshold to the complete put wing shrinks the median area of the joint set, measured in probability and depth coordinates, by a further 5.4 to 18.2%. The shrinkage comes from requiring one distribution to fit the whole wing, not from a better fit near K . A comparison with price-grid methods shows what the threshold coordinate adds: median grid errors vanish under refinement, but a grid misplaces any atom at the threshold, and the probability error then equals the mass misplaced.

Section 2 places the paper in the literature. The closest methods either discretize the state space (Barratt et al., 2024) or select a single distribution (De Vries, 2026); we keep the original quote intervals, characterize the exact joint region, and handle the threshold without discretization. Sections 3–5 define the identified object, derive its finite representation, and describe the reconstruction algorithm. Sections 6–8 present the data, the identification gains, and the price-grid comparison. Section 9 concludes, and the appendix collects proofs and computational details.

2 Related literature

Merton (1973) derives general no-arbitrage restrictions on option prices, and Breeden and Litzenberger

(1978) link the curvature of call prices across strikes to state prices. In this tradition, nonparametric recovery selects one distribution or one arbitrage-free price surface by optimization or smoothing (Jackwerth and Rubinstein, 1996; Aït-Sahalia and Lo, 1998; Fengler, 2009). Bid–ask information enters the same way. Monnier (2013) selects the smoothest density whose put prices lie inside the quoted intervals, Wizman et al. (2026) combine arbitrage filtering with a smoothness and entropy criterion for short-dated quotes, and Qin et al. (2026), starting from discrete arbitrage-free prices, construct a marginal that reprices every input exactly. We instead retain every distribution compatible with the intervals. Where a selection criterion resolves observational equivalence, we measure it.

The mathematical foundation is the model-free bounds literature. Bertsimas and Popescu (2002) treat sharp option-price bounds under moment information as generalized moment problems, King et al. (2005) develop calibrated option bounds, and Davis and Hobson (2007) and Cousot (2007) characterize when finitely many prices are consistent with an arbitrage-free model. Bertsimas and Bushueva (2006) reduce bounds for continuous piecewise-linear payoffs to linear programs, Bassett (1997) derives probability bounds from finitely many option prices, and Cohen et al. (2020) repair option quote data by sparse linear programming. Discontinuous payoffs behave differently. Cox and Obłój (2011) document nonattainment and sensitivity to boundary mass under finite strike information, and Liang et al. (2026) show that digital bounds respond to vanishing bid–ask spreads at a square-root rate. Our threshold coordinate carries this boundary behavior into the finite-dimensional joint projection.

The closest computational work is Barratt et al. (2024), who bound the CDF, conditional probabilities, VaR, and CVaR over risk-neutral probabilities on a discretized price grid. De Vries (2026) construct projection estimators for general payoffs, obtain a pricing-consistent CDF from indicator payoffs, and prove a finite-sample approximation bound. Their method returns one estimator under a chosen projection weight, whereas our target is the full feasible set under the observed intervals. van Eekelen (2023) derives sharp generalized moment bounds for conditional expectations, and our object identifies the conditioning probability jointly with its shortfall numerator. Neufeld et al. (2023) compute ε -optimal bounds for continuous piecewise-affine payoffs, a class that excludes digitals without prior approximation. On the computational side, enumeration of projected polyhedra by repeated support programs is longstanding in multiobjective optimization (Ehrgott et al., 2012), and support-function representations of convex identified sets are established in partial identification (Beresteanu and Molinari, 2008; Löhne and Weißing, 2016). What is new here is the target: a vanilla representation with explicit threshold mass whose projection is the entire sharp joint region, mapped analytically into conditional depth.

A related empirical literature extracts disaster and tail information from index options (Backus et al., 2011; Bollerslev and Todorov, 2011). Vilkov and Xiao (2013) fit parametric tails motivated by extreme value theory, and O’Sullivan and Wang (2018) estimate risk-neutral tail probabilities and tail expectations beyond the observed strike range. Bounds on physical crash probabilities require restrictions linking physical and risk-neutral measures; Martin and Shi (2026) impose a power utility stochastic discount factor and bound over dependence structures consistent with option-implied marginals. Our analysis stays entirely under $\mathbb{Q}$.

Finally, the fixed-threshold depth s_K differs from quantile-based tail measures. When K is the relevant quantile, the total normalized loss $(1 - K) + s_K = 1 - \mathbb{E}_{\mathbb{Q}}[X \mid X \leq K]$ is the CVaR of the loss $1 - X$, and with mass at the quantile the tail-average convention splits the boundary atom (Rockafellar and Uryasev, 2000, 2002). Our object fixes K and retains the full joint set implied by the quote intervals.

3 Market setup and the joint tail object

3.1 Normalization and quote intervals

Fix a date and a maturity T . Let S_T be the terminal SPX level, $D_{0,T}$ the discount factor, and $F_{0,T}$ the forward, and set

$$X = \frac{S_T}{F_{0,T}}, \quad X \geq 0.$$

Under a risk-neutral measure with the correct forward,

$$\mathbb{E}_{\mathbb{Q}}[X] = 1.$$

A strike κ in index points has normalized strike $k = \kappa/F_{0,T}$, and the time-0 call price $C(\kappa, T)$ has normalized value

$$c_{\mathbb{Q}}(k) = \mathbb{E}_{\mathbb{Q}}[(X - k)^+] = \frac{C(\kappa, T)}{D_{0,T}F_{0,T}}.$$

Let $\mathcal{I}$ index the retained quotes. Quote $i \in \mathcal{I}$ supplies the interval

$$\underline{c}_i \leq c_{\mathbb{Q}}(k_i) \leq \bar{c}_i.$$

Put quotes with bid $\underline{P}_i$ and ask $\bar{P}_i$ enter through parity:

$$[\underline{c}_i, \bar{c}_i] = \left[\frac{\underline{P}_i}{D_{0,T}F_{0,T}} + 1 - k_i, \frac{\bar{P}_i}{D_{0,T}F_{0,T}} + 1 - k_i \right].$$

We impose no midpoint, and when several retained quotes share a strike, each interval remains a separate constraint.

Let $\mathcal{P}(\mathbb{R}_+)$ denote the probability measures on $\mathbb{R}_+$. The quote-compatible measures are

$$\mathcal{M}(\mathcal{I}) = \left\{ \mathbb{Q} \in \mathcal{P}(\mathbb{R}_+) \left| \begin{array}{l} \mathbb{E}_{\mathbb{Q}}[X] = 1, \\ \underline{c}_i \leq \mathbb{E}_{\mathbb{Q}}(X - k_i)^+ \leq \bar{c}_i, \quad i \in \mathcal{I} \end{array} \right. \right\}.$$

When these restrictions are jointly feasible we call the cross section feasible, and $\mathcal{M}(\mathcal{I})$ is the identified set of distributions under the original intervals. Section 4.4 treats inconsistent cross sections.

3.2 Probability, loss, and conditional depth

Fix $K \in (0, 1)$. For $\mathbb{Q} \in \mathcal{M}(\mathcal{I})$, define

$$p_K(\mathbb{Q}) = \mathbb{E}_{\mathbb{Q}}[\mathbf{1}\{X \leq K\}], \tag{1}$$

$$L_K(\mathbb{Q}) = \mathbb{E}_{\mathbb{Q}}[(K - X)^+], \tag{2}$$

$$s_K(\mathbb{Q}) = \frac{L_K(\mathbb{Q})}{p_K(\mathbb{Q})}, \quad p_K(\mathbb{Q}) > 0. \tag{3}$$

The identity

$$s_K(\mathbb{Q}) = K - \mathbb{E}_{\mathbb{Q}}[X \mid X \leq K]$$

shows that $s_K \in [0, K]$: it is the average forward-normalized shortfall below K , conditional on the threshold event. The total conditional decline from the forward is $(1 - K) + s_K$, so the depth coordinate isolates the shortfall beyond the chosen threshold.

The attainable image, its closure, and the conditional depth image are

$$\mathcal{R}_K^{\text{att}}(\mathcal{I}) = \{(p_K(\mathbb{Q}), L_K(\mathbb{Q})) : \mathbb{Q} \in \mathcal{M}(\mathcal{I})\}, \tag{4}$$

$$\mathcal{R}_K(\mathcal{I}) = \overline{\mathcal{R}_K^{\text{att}}(\mathcal{I})}, \tag{5}$$

$$\mathcal{T}_K(\mathcal{I}) = \{(p, L/p) : (p, L) \in \mathcal{R}_K(\mathcal{I}), p > 0\}. \tag{6}$$

We call the boundary of $\mathcal{R}_K(\mathcal{I})$ the risk-neutral crash frontier of the cross section; since every object in this paper lives under $\mathbb{Q}$, we usually drop the qualifier.

Every pair satisfies the universal inequalities $0 \leq L_K \leq Kp_K$, a wedge in the (p, L) plane. The joint set is typically a proper subset of the product of its marginal intervals intersected with this wedge, because both coordinates must come from one quote-compatible measure. Section 6 turns this comparison into the marginal benchmark used in the empirical analysis.

4 A finite representation with explicit threshold mass

The finite representation has three ingredients: cell masses and first moments that price every retained vanilla payoff, one added coordinate that assigns threshold mass to the event, and the quote intervals that select the admissible moment vectors. Theorem 4.3 shows that together they recover the joint object exactly.

4.1 Cell probabilities and first moments

Let

$$0 = b_0 < b_1 < \dots < b_J$$

be the sorted union of all retained normalized strikes and the threshold K . The bounded cells are $I_j = [b_j, b_{j+1})$, $j = 0, \dots, J-1$, and the terminal cell is $I_J = [b_J, \infty)$. Define

$$z_j = \mathbb{Q}(X \in I_j), \quad m_j = \mathbb{E}_{\mathbb{Q}}[X \mathbf{1}\{X \in I_j\}],$$

and collect them as

$$z = (z_0, \dots, z_J)^\top \in \mathbb{R}_+^{J+1}, \quad m = (m_0, \dots, m_J)^\top \in \mathbb{R}_+^{J+1}.$$

They satisfy

$$z_j \geq 0, \quad m_j \geq 0, \quad \sum_{j=0}^J z_j = 1, \quad \sum_{j=0}^J m_j = 1, \quad (7)$$

$$b_j z_j \leq m_j \leq b_{j+1} z_j, \quad j = 0, \dots, J-1, \quad (8)$$

$$m_J \geq b_J z_J. \quad (9)$$

Every retained call payoff is affine on each cell, because its only kink is at a strike and every strike is a cell boundary, so its price depends on the measure only through (z, m) . For an observed strike $k = b_\ell$,

$$c_{\mathbb{Q}}(k) = \sum_{j=\ell}^J (m_j - k z_j). \quad (10)$$

We write $c_i(z, m)$ for the corresponding linear form of quote i , so every bid–ask restriction is linear in (z, m) .

Lemma A.1 in Appendix A.1 proves the converse: every vector satisfying (7)–(9) is a limit of cell vectors of unit-mean measures, and when every nonzero conditional cell mean m_j/z_j lies in the interior of its cell, placing mass z_j at m_j/z_j attains it exactly. Limits are needed only at shared cell endpoints and on the unbounded tail. Thus (10) is exact without discretizing the state space.

Remark 4.1 (Order-one truncated moment problem). Conditions (7)–(9) are the feasibility conditions of the truncated moment problem of order one on each cell (Krein and Nudelman, 1977). Piecewise-affine payoffs price through cell masses and first moments alone, which is why linear constraints suffice. Targets built from higher moments of the state would require higher cell moments and semidefinite conditions.

4.2 The threshold atom

Let j_K be the index with $K = b_{j_K}$, and introduce

$$a_K = \mathbb{Q}(X = K).$$

In the finite system, a_K is the coordinate that allocates threshold mass to the event. It satisfies

$$0 \leq a_K \leq z_{j_K}. \quad (11)$$

The residual mass $z_{j_K} - a_K$ lies strictly to the right of K in the threshold cell, so its first moment $m_{j_K} - Ka_K$ is at least $K(z_{j_K} - a_K)$, which reduces to $m_{j_K} \geq Kz_{j_K}$ and is already imposed by (8) or (9). When $j_K < J$, the same residual mass lies below b_{j_K+1} , which gives

$$m_{j_K} \leq b_{j_K+1}z_{j_K} - (b_{j_K+1} - K)a_K. \quad (12)$$

When $j_K = J$, (12) is omitted. Together these restrictions separate mass at K from residual mass to its right. The tail coordinates are then

$$p_K = \sum_{j < j_K} z_j + a_K, \quad (13)$$

$$L_K = \sum_{j < j_K} (Kz_j - m_j). \quad (14)$$

The atom counts toward the probability and adds nothing to the loss, as the definitions of p_K and L_K require.

Let $\mathcal{F}_K(I)$ denote the set of vectors (z, m, a_K) that satisfy (7)–(9), (11)–(12) (with (12) omitted when $j_K = J$), and every quote interval through (10). Let G_K denote the linear map (13)–(14) from (z, m, a_K) to (p_K, L_K) .

Assumption 4.2 (Strict quote interior (a Slater condition)). *Every retained quote interval has positive width, and $\mathcal{F}_K(I)$ contains a vector whose modeled prices lie strictly inside their respective intervals.*

This constraint qualification (Rockafellar, 1970) ensures that closing the set of attainable cell vectors commutes with intersecting it with the closed quote intervals. It is checked by one auxiliary linear program that maximizes the smallest relative margin between modeled prices and quote endpoints; given positive widths, the assumption holds exactly when the optimal margin is positive (Appendix A.2.1). A strictly interior complete quote set qualifies each of its nested subsets.

Theorem 4.3 (Sharp projection with explicit threshold mass). *Under Assumption 4.2,*

$$\mathcal{R}_K(I) = G_K(\mathcal{F}_K(I)).$$

Every sharp linear bound on (p_K, L_K) is therefore the value of a finite-dimensional linear program.

Remark 4.4 (Extremal scenarios are simple). The measures constructed in the proof of Theorem 4.3 are discrete, with at most one atom per cell, one atom at the threshold, and two auxiliary atoms. Every point of $\mathcal{R}_K(I)$ is therefore generated, or approached, by distributions with at most $J + 4$ support points, and the scenarios that trace the crash frontier are small collections of atoms.

4.3 Why a price grid can fail at the threshold

Without a_K , moving mass from $K - \varepsilon$ to $K + \varepsilon$ changes vanilla prices continuously but changes p_K discontinuously. A price-grid method, which replaces $\mathbb{Q}$ by probabilities on fixed nodes, inherits this problem: it can match every continuous payoff closely and still assign threshold mass to the wrong side of the event.

Proposition 4.5 (Threshold nonuniformity). *Fix $\alpha \in (0, 1)$. For every sufficiently small $\varepsilon > 0$ there are two unit-mean distributions on $\mathbb{R}_+$, one placing mass α at $K - \varepsilon$ and the other at $K + \varepsilon$, whose call prices differ by at most $4\alpha\varepsilon$ at every strike. Their values of $\mathbb{Q}(X \leq K)$ differ by α .*

The map from vanilla prices to sharp bounds on the distribution function is thus discontinuous at any law with an atom at the threshold, and refining a grid does not restore uniform accuracy. This is why the finite system carries the threshold mass explicitly. Section 8 quantifies the effect on real and synthetic cross sections.

4.4 Inconsistent cross sections

Some cross sections admit no distribution consistent with every quote interval. For these cases we relax each interval by nonnegative slacks, penalize the slacks through fixed nonnegative weights, and impose a common budget on the weighted slack sum. The budget is fixed before the main analysis and shared by every nested quote set within a cross section, so the resulting regions remain nested. The relaxed system is again a finite linear system, and the reconstruction of Section 5 applies without change. Appendix A.2.2 states the relaxed system, the weights, and the relaxed region under the budget, written $\mathcal{R}_{K,\delta}(I)$. Sharp results under the original intervals use feasible cases satisfying Assumption 4.2, and relaxed results are reported separately.

5 Frontier reconstruction

5.1 Support programs

Collect the decision variables in $y = (z^\top, m^\top, a_K)^\top$ and write $\mathcal{Y}$ for the feasible polytope: $\mathcal{Y} = \mathcal{F}_K(I)$ under the original intervals, and its slack-augmented analogue under the common budget (Appendix A.2.2), in which case G_K acts on the (z, m, a_K) block of y . The mass and first-moment normalizations, together with the budget on any slacks, make $\mathcal{Y}$ bounded, so its projection is a compact polygon. For a direction $u = (u_p, u_L) \in \mathbb{R}^2$, the support program is the linear program

$$h(u) = \max_{y \in \mathcal{Y}} u^\top G_K y. \quad (15)$$

Corollary 5.1 (Sharp portfolio bounds). *Let Assumption 4.2 hold and let $\mathcal{Y} = \mathcal{F}_K(I)$. For every $u = (u_p, u_L) \in \mathbb{R}^2$,*

$$h(u) = \sup_{Q \in \mathcal{M}(I)} \mathbb{E}_Q[u_p \mathbf{1}\{X \leq K\} + u_L (K - X)^+],$$

and the corresponding sharp lower value is $-h(-u)$.

The corollary restates Theorem 4.3 through support functions, which coincide for a convex set and its closure. Every portfolio of the digital and put payoffs is priced sharply by one linear program, and once the polygon is known no further program is needed. Appendix A.2 describes the vertex selection used for polygon enumeration and the primal–dual accuracy checks.

5.2 Adaptive hull algorithm

The hull procedure adapts standard support-function and outer-approximation methods (Ehrgott et al., 2012; Löhne and Weißing, 2016). Let V_t be the support points collected after iteration t and $H_t = \text{conv}(V_t)$. The reconstruction initializes V_0 from eight evenly spaced directions, queries the outward normal of each current edge, and adds any support point outside the current hull. In exact arithmetic this process recovers a point, a line segment, or the full polygon after finitely many support solves; Proposition A.2 in Appendix A.1 gives the formal result, and Appendix A.2 the implementation. Dual support halfspaces also define an outer polygon, and the inner–outer area difference measures reconstruction error.

5.3 The conditional depth image

For $p > 0$, define $\phi(p, L) = (p, L/p)$. The map preserves the first coordinate but not convexity, so the depth image is recovered from the polygon edge by edge. To measure joint ambiguity away from zero probability, fix a floor $p_0 > 0$. For a nonempty compact convex $C \subset \mathbb{R}^2$, let $C(p) = \{L : (p, L) \in C\}$, and when the section is nonempty write $L_C^+(p) = \max C(p)$ and $L_C^-(p) = \min C(p)$. Define

$$\mathcal{A}_{p_0}(C) = \int_{\{p \geq p_0 : C(p) \neq \emptyset\}} \frac{L_C^+(p) - L_C^-(p)}{p} dp, \quad (16)$$

the transformed area of C .

Proposition 5.2 (Depth image and transformed area). *On $\{p > 0\}$, ϕ is a diffeomorphism with $|\det D\phi(p, L)| = 1/p$. A nonvertical polygon edge with equation $L = \alpha_e + \beta_e p$ maps to the arc $s(p) = \beta_e + \alpha_e/p$, and a vertical edge maps to a vertical segment. For every nonempty compact convex C and every $p_0 > 0$,*

$$\mathcal{A}_{p_0}(C) = \text{Leb}_2(\phi(C \cap \{p \geq p_0\})),$$

where Leb_2 is planar Lebesgue measure.

The edge formulas are direct substitution, the area identity is the change-of-variables formula, and for polygonal C the integral in (16) has the closed form given in Appendix A.2. On any domain $p \geq p_0 > 0$, the relative boundary of $\mathcal{T}_K(\mathcal{I})$ is therefore a finite collection of vertical segments and arcs of the form $\beta_e + \alpha_e/p$. The frontier itself is defined for every $p > 0$; the floor p_0 enters only the area summary. Because area depends on the chosen coordinates, the full frontier remains the primary identified object.

5.4 Dual certificates

The dual of (15) attaches multipliers to the bid, ask, mass, first-moment, cell, and threshold-atom constraints, and strong duality gives the same support value in both programs (Appendix A.2.4). For each exposed boundary point, an optimal dual solution decomposes the bound into cash and forward terms, a signed portfolio of retained quote endpoints, and nonnegative structural cell terms. Quote constraints with nonzero multipliers identify which intervals produce the bound; dual optima need not be unique, and the certificate reports one optimal solution. We call this decomposition a dual certificate.

Read as a portfolio, the certificate is a static super-replication of the target payoff: cash and forward positions plus retained quotes executed at bid or ask, whose combined payoff dominates the digital and put target state by state, with the structural terms certifying the domination on each cell (Davis and Hobson, 2007; Cox and Obłój, 2011). The bound is the cost of that portfolio, lower bounds correspond to sub-replication, and each reported bound is therefore enforceable by trading the quotes the certificate names. For relaxed cross sections the dual carries an additional budget multiplier, so the portfolio reading is exact only under the original intervals.

6 Data and empirical design

6.1 SPX cross sections

The application uses end-of-day SPX call and put bid–ask quotes from OptionMetrics, from 2 January 2013 through 31 August 2023. Normalization uses the SPX level and a forward estimate for the selected date and maturity, and robustness checks apply multiplicative shifts of ± 20 basis points to the forward.

The sampling rule selects one eligible cross section per calendar week, preferring midweek dates. Each selected date must offer a short expiry of 7–21 days closest to 14 and a medium expiry of 22–60 days closest to 45, with at least 18 distinct put strikes in log-moneyness $[-0.35, 0.05]$. The sample has 557 dates from 2 January 2013 to 30 August 2023, 98.92% of them Wednesdays. Appendix A.3 records the weekday fallback order and the tie rules.

We retain positive strikes, nonnegative bids, positive asks, noncrossed bid–ask pairs, positive trading volume, and expiries between 7 and 120 days. The primary quote sets use puts converted by parity, and a call–put cross section is reserved for robustness. The complete put wing comprises all retained puts with log-moneyness in $[-0.35, 0.05]$ for the selected date and expiry, and multiple retained observations at the same strike remain separate constraints.

6.2 Prespecified estimands and quote sets

The main thresholds are $K \in \{0.85, 0.90\}$, forward-normalized declines of 15% and 10%. Within each cross section we compare four objects:

- (i) the marginal benchmark: the wedge-adjusted product of separate marginal bounds, defined below;

- (ii) the put pair: the nearest puts at or below and at or above K ;
- (iii) the local eight: the put pair plus the six nearest remaining put strikes;
- (iv) the complete put wing $\mathcal{I}^C$: every eligible put in the cross section.

Tie rules for the put pair are recorded in Appendix A.3, and the slack weights used in the relaxation analysis in Appendix A.2.2.

For a nonempty compact convex $C \subset \mathbb{R}^2$, the wedge-adjusted marginal benchmark is

$$\mathcal{B}_K(C) = \left([\min_C p, \max_C p] \times [\min_C L, \max_C L] \right) \cap \{(p, L) : 0 \leq L \leq Kp\},$$

with extrema over $(p, L) \in C$, so that $\mathcal{R}_K(\mathcal{I}) \subseteq \mathcal{B}_K(\mathcal{R}_K(\mathcal{I}))$. The reported area ratio is

$$\frac{\mathcal{A}_{p_0}(C)}{\mathcal{A}_{p_0}(\mathcal{B}_K(C))},$$

the share of the transformed benchmark area that survives joint identification. By Proposition 5.2, both areas are measured in probability and depth coordinates. Smaller ratios indicate tighter joint identification, and a ratio of one means that joint restrictions add nothing beyond the marginal intervals and the universal wedge on $p \geq p_0$. The ratio is reported for two-dimensional joint sets with positive benchmark area; point and line sets are classified as degenerate. The main specification uses $p_0 = 10^{-4}$, with 10^{-5} and 10^{-3} as robustness values.

For a quote-based baseline A , write $\mathcal{R}_K^A = \mathcal{R}_K(\mathcal{I}^A)$ under the original intervals and $\mathcal{R}_{K,\bar{\delta}}^A = \mathcal{R}_{K,\bar{\delta}}(\mathcal{I}^A)$ under the common budget. Because the quote sets are nested within each cross section, $\mathcal{I}^A \subseteq \mathcal{I}^C$ implies $\mathcal{R}_K^C \subseteq \mathcal{R}_K^A$ and $\mathcal{R}_{K,\bar{\delta}}^C \subseteq \mathcal{R}_{K,\bar{\delta}}^A$. When the baseline area is positive, the identification gain from set A to the complete put wing is

$$\Delta_A = 1 - \frac{\mathcal{A}_{p_0}(\mathcal{R}_K^C)}{\mathcal{A}_{p_0}(\mathcal{R}_K^A)},$$

the fractional contraction in transformed area from the additional quotes. It satisfies $0 \leq \Delta_A \leq 1$ in exact arithmetic. Sensitivity calculations use the same formula with $\mathcal{R}_{K,\bar{\delta}}^A$ in place of $\mathcal{R}_K^A$.

Feasibility under the original bid–ask constraints is the primary criterion and holds in 93.90% of the 2,228 cases with the complete put wing. Assumption 4.2 holds in 2,076 of these cases. Cross sections with ambiguous contract identity are excluded, leaving 1,966 cases in the primary empirical sample. Sixteen feasible cases fail Assumption 4.2, and 136 cases are infeasible without relaxation and enter the separately reported common-budget analysis.

6.3 Sampling variation and robustness

We form percentile intervals for median identification gains from 2,000 draws of a circular block bootstrap. Blocks contain eight chronologically ordered dates, the cube-root rule applied to the number of sampled dates in each primary cell. Resampling dates rather than individual quotes preserves market cross sections and allows short-range serial dependence (Lahiri, 2003).

Robustness checks target four design choices: the probability floor, forward and spread inputs, strike and duplicate selection, and the maturity and threshold grid. Probability-floor results reuse the reconstructed polygons, and the remaining checks use selected date subsamples. Appendix A.3 summarizes the specifications.

7 Empirical results

7.1 Main identification gains

Table 1 summarizes identification gains in the primary sample. Across the four combinations of horizon and threshold, the complete put wing contracts the area of the local eight by 5.35% to 18.16% at the

Table 1: Incremental identification from the complete put wing

Horizon	K	Put pair		Local eight	
		Median	95% interval	Median	95% interval
Short	0.85	74.33%	[70.15,77.59]	5.35%	[0.52,13.92]
Short	0.90	83.10%	[81.42,84.27]	15.03%	[10.38,22.75]
Medium	0.85	84.75%	[83.28,86.25]	8.12%	[2.73,13.71]
Medium	0.90	89.03%	[87.63,90.39]	18.16%	[13.19,21.56]

Notes: Each entry is $1 - \mathcal{A}_{p_0}(\mathcal{R}_K^C) / \mathcal{A}_{p_0}(\mathcal{R}_K^{\text{baseline}})$, with $p_0 = 10^{-4}$. Larger entries indicate greater contraction relative to the stated baseline. For each threshold, there are 535 primary short-horizon cases and 448 primary medium-horizon cases, giving 1,966 cases in total. Intervals use 2,000 circular block bootstrap draws with eight chronologically ordered dates satisfying the primary-sample conditions per block. Point estimates use the midpoint of the numerical area interval, and the intervals summarize sampling variation only; numerical accuracy is verified separately (Appendix A.3).




Figure 1: Sharp joint frontiers and primary-sample identification gains. The representative case is also used for the dual certificate in Figure 3. Filled regions in panels (a) and (b) correspond to the put pair, local eight, and complete put wing. Panel (a) plots the probability-loss set with loss in basis points of forward. Panel (b) zooms the depth range on a log- p axis and its inset restores the full range. Panel (c) reports primary-sample median area reductions and circular block bootstrap intervals for the four main specifications.

median. Quotes beyond the eight nearest strikes therefore eliminate between one-twentieth and one-fifth of the joint ambiguity that those strikes leave open.

The contraction is global rather than local in origin. Quotes away from K restrict where a common risk-neutral distribution can place compensating probability and first moment while keeping every modeled price inside its interval and satisfying $\mathbb{E}_{\mathbb{Q}}[X] = 1$. A reallocation that fits the nearby strikes can therefore fail elsewhere on the wing. The complete put wing disciplines the feasible probability-depth combinations rather than refining the interpolation at K .

Across primary-sample cases with the complete put wing, the median ratio of joint area to wedge-adjusted marginal benchmark area is 0.6340. In the median case, 36.6% of the transformed benchmark area is unattainable: separate sharp marginals, even combined with the inequality $L \leq Kp$, admit many probability-depth pairs that no single quote-compatible distribution generates.

Median probability widths under the complete put wing are 0.0045 and 0.011 at the short horizon and 0.019 and 0.037 at the medium horizon, for $K = 0.85$ and 0.90. Marginal probability uncertainty therefore remains visible even where joint restrictions sharply reduce the probability-depth region.

Figure 1 displays the representative case used for the dual certificate, the short-horizon cross section dated 5 July 2017 with maturity 19 July 2017 and $K = 0.90$. It satisfies the original quote constraints and Assumption 4.2.

7.2 Stability and robustness

Appendix Table A1 reports subperiod summaries and results that exclude the calibration dates, both under the primary sample definition. The median area ratio varies little across subperiods.

Median identification gains remain positive under prespecified perturbations to forwards, quotes, strike coverage, horizons, and thresholds. Appendix Table A2 reports the range of median ratios for each specification. The denominator is itself a moving benchmark, so doubled spreads can lower the ratio even as absolute probability widths grow.

8 Price-grid benchmark and dual evidence

8.1 Median convergence and worst-case persistence

The benchmark follows the price-grid approach of Barratt et al. (2024). We repeat the support analysis on price grids with 51, 101, 201, and 401 nodes spanning $[0, 2]$ or $[0, 3]$ (support caps 2 and 3), with the threshold either inserted as a node or centered in a cell. Every grid solve is compared with a continuous-state solve in the same 64 directions and under the same quote loss budget, so support errors compare the two state representations at a common relaxation level. Each specification contains 36 cases: 32 empirical combinations of cross section and threshold, and four synthetic atomic cases. Appendix A.3 details the grid construction, the shared budget, and the error definitions.

The budget required for grid feasibility is not negligible at 51 nodes, with median increments over the continuous-state minimum of $(1.22\text{--}1.63) \times 10^{-4}$ at cap 2, and it falls to the numerical floor of 10^{-10} at 201 and 401 nodes. Because the budget varies with the grid, changes across grids combine refinement with changes in the relaxation level. Table A3 reports the values.

Figure 2 separates typical from worst-case error. Median error declines strongly and is numerically zero with K inserted and 401 nodes. The maximum, however, remains 0.15 with K inserted, and centering K in a cell lowers the plateau to about 0.075. Both plateaus come from the synthetic 0.15 threshold atom, as Proposition 4.5 predicts: the inserted-node plateau equals the atom mass, and the centered-cell plateau is about half as large. Grid density improves the approximation of continuous payoffs, while the event convention governs atomic cases. An inserted node assigns all of its mass to $\{X \leq K\}$, and representing limiting mass immediately to its right would require a separate one-sided threshold state. On the real SPX cross sections at support cap 2, the maximum error falls from 0.0751 to 0.0100 for centered grids and from 0.0784 to 0.0138 with K inserted. The persistent maxima therefore come from the atomic counterexample, and errors on typical cross sections decline with refinement.

Computation time rises with grid density, as panel (c) shows. The two methods address different computational objectives: the continuous-state method reconstructs an entire frontier, whereas the benchmark evaluates a fixed set of directions on each grid.

8.2 Primal–dual verification and a dual certificate

For the representative case, primal–dual gaps and KKT residuals are below 2×10^{-14} in the directions that maximize probability, loss, and the probability–loss trade-off. At a 10^{-8} multiplier threshold, the certificate for each direction puts weight on two quote endpoints, and the remaining terms enforce the mass, cell-moment, and threshold-atom restrictions. Figure 3 displays an optimal trade-off dual certificate.

9 Conclusion

Option-implied crash probability and conditional depth are two readings of the same put value, tied together by the risk-neutral distribution that generates both. Marginal intervals lose that tie, and the risk-neutral crash frontier is built to keep it. Under a verifiable interior condition the joint set is the exact projection of a finite system of linear inequalities with an explicit coordinate for threshold mass, linear




Figure 2: Price-grid accuracy and computation time. Panel (a) first takes the median support error over directions within each case, then reports the median of those values across cases. In panel (b), lines report the 95th percentile of each case’s maximum support error over directions, and crosses report sample maxima. Panel (c) reports median computation time. Synthetic threshold-atom designs generate the horizontal worst-case plateau. Every plotted grid is compared with a continuous-state solve under the same quote loss budget for that grid.




Figure 3: Representative trade-off dual certificate for direction $u^* = (1, -1)$ in the representative frontier, shown at its original scale. Axis labels identify each quote by strike and endpoint. Only quote weights above the 10^{-8} threshold are displayed.

programs recover its boundary, and every point on the boundary comes with a dual certificate that names the individual quotes behind it.

In more than ten years of weekly SPX cross sections, the complete put wing contracts the transformed area by 5.4–18.2% relative to the eight nearest strikes and by 74.3–89.0% relative to the nearest pair. Strikes far from the threshold therefore carry real information about which joint tail scenarios remain compatible with the market. Price grids do not substitute for the threshold coordinate: median grid errors fall with refinement, but worst-case errors at atomic thresholds persist because the event indicator is discontinuous. The continuous-state frontier handles that boundary directly and turns a finite option cross section into valuation bounds, feasible stress scenarios, comparisons across strike sets, and certificates that trace each bound back to observed quotes.

The analysis prices one maturity at a time and stays under the risk-neutral measure. Partitions that couple several maturities, and preference restrictions that map the frontier into bounds on physical crash probabilities, are natural extensions.

A Appendix: Proofs, computation, and supplementary evidence

A.1 Proofs and auxiliary results

A.1.1 Cell representation

Lemma A.1 (Cell-moment representation for retained payoffs). *Let each retained payoff g_ℓ be continuous and affine on every cell, $g_\ell(x) = \alpha_{\ell j} + \beta_{\ell j}x$ for $x \in I_j$. The cell vector (z, m) of any unit-mean probability measure $\mathbb{Q}$ satisfies (7)–(9), and*

$$\mathbb{E}_{\mathbb{Q}}[g_\ell(X)] = \sum_{j=0}^J (\alpha_{\ell j} z_j + \beta_{\ell j} m_j) \quad \text{for every } \ell.$$

Conversely, for every (z, m) satisfying (7)–(9) there are unit-mean probability measures $\mathbb{Q}_n$ with

$$\mathbb{E}_{\mathbb{Q}_n}[g_\ell(X)] \longrightarrow \sum_{j=0}^J (\alpha_{\ell j} z_j + \beta_{\ell j} m_j) \quad \text{for every } \ell.$$

Proof. Integrating $b_j \leq x < b_{j+1}$ over a bounded cell and $x \geq b_J$ over the terminal cell gives (8)–(9), and (7) is unit mass and unit mean. Since g_ℓ is affine on I_j ,

$$\mathbb{E}_{\mathbb{Q}}[g_\ell(X)\mathbf{1}\{X \in I_j\}] = \alpha_{\ell j} z_j + \beta_{\ell j} m_j,$$

and summing over j gives the pricing formula.

For the converse, the closures of the attainable mass and first-moment cones on a bounded cell and on the terminal cell are, respectively,

$$\{(z, m) : z \geq 0, b_j z \leq m \leq b_{j+1} z\}, \quad \{(z, m) : z \geq 0, m \geq b_J z\}.$$

For $z > 0$, mass z at m/z attains (z, m) unless $m/z = b_{j+1}$, the excluded right endpoint of a bounded cell, which is approached from the left. On the terminal cell, $(0, m)$ with $m > 0$ is approached by mass $1/n$ at nm . Applying these constructions cellwise gives finite measures μ_n whose payoff vectors converge to the prescribed vector and whose masses M_n and first moments U_n converge to one. Let X_n have law μ_n/M_n and set $\widehat{X}_n = (M_n/U_n)X_n$, so that $\mathbb{E}[\widehat{X}_n] = 1$. Each g_ℓ is Lipschitz, hence

$$\left| \mathbb{E}g_\ell(\widehat{X}_n) - \mathbb{E}g_\ell(X_n) \right| \leq \text{Lip}(g_\ell) \left| \frac{M_n}{U_n} - 1 \right| \frac{U_n}{M_n} \longrightarrow 0.$$

Since $M_n \rightarrow 1$, the payoff vectors of the laws of X_n , and therefore of $\widehat{X}_n$, converge to the prescribed vector. $\square$

A.1.2 Proof of Theorem 4.3

Proof. Forward inclusion. For $\mathbb{Q} \in \mathcal{M}(I)$, set $a_K = \mathbb{Q}(X = K)$. The residual threshold-cell pair is $(z_{j_K} - a_K, m_{j_K} - Ka_K)$. It satisfies (11), and when $j_K < J$ its upper cell bound is

$$m_{j_K} - Ka_K \leq b_{j_K+1}(z_{j_K} - a_K),$$

which is (12). Splitting the cells below K from the atom at K gives (13)–(14). Therefore $\mathcal{R}_K^{\text{att}}(I) \subseteq G_K(\mathcal{F}_K(I))$. The latter is the linear image of a compact polyhedron and is closed, so

$$\mathcal{R}_K(I) \subseteq G_K(\mathcal{F}_K(I)).$$

Reverse inclusion: cell realization. Fix $y = (z^\top, m^\top, a_K)^\top \in \mathcal{F}_K(I)$. Apply the cellwise constructions of Lemma A.1, before global normalization, to every nonthreshold pair and to $(z_{j_K} - a_K, m_{j_K} - Ka_K)$, and

place mass a_K at K . The structural constraints put the residual pair in the moment cone of $[K, b_{j_K+1})$ or $[K, \infty)$. A residual mean equal to K is approached from the right, which keeps that mass outside $\{X \leq K\}$ while all continuous-payoff values converge, and Lemma A.1 handles the remaining cell boundaries. This gives finite nonnegative measures $\tilde{\mu}_n$ with mass $1 + A_n$ and first moment $1 + B_n$, where $A_n, B_n \rightarrow 0$.

Mass and mean correction. Fix $x_{\text{buf}} > \max\{b_J, 1\}$ and the unit-mass, unit-mean buffer

$$\mu_{\text{buf}} = \left(1 - \frac{1}{x_{\text{buf}}}\right) \delta_0 + \frac{1}{x_{\text{buf}}} \delta_{x_{\text{buf}}}.$$

For sufficiently large n , set $\varepsilon_n = (|A_n| + |B_n|)^{1/2} + n^{-1} < 1$ and $\bar{\mu}_n = (1 - \varepsilon_n)\tilde{\mu}_n + \varepsilon_n\mu_{\text{buf}}$. If $\bar{A}_n$ and $\bar{B}_n$ are its mass and moment errors, set

$$\mu_n = \bar{\mu}_n + \left(-\bar{A}_n + \frac{\bar{B}_n}{x_{\text{buf}}}\right) \delta_0 - \frac{\bar{B}_n}{x_{\text{buf}}} \delta_{x_{\text{buf}}}.$$

Then μ_n has exactly unit mass and unit mean. The correction has total absolute size $O(|A_n| + |B_n|)$, which is $o(\varepsilon_n)$ by the choice of ε_n , while the buffer atoms at 0 and x_{buf} have mass of order ε_n ; hence μ_n is nonnegative for all large n . The correction changes the target coordinates and every retained vanilla value by $o(1)$, preserves support in $\mathbb{R}_+$, and does not move the atom at K . Thus the projections converge to $G_K(y)$.

Quote compatibility. Let y° be the vector from Assumption 4.2 whose modeled prices lie strictly inside all quote intervals, and put $y_\theta = (1 - \theta)y + \theta y^\circ$. For every $0 < \theta \leq 1$, its modeled prices have positive distance from all quote endpoints, so the preceding approximation can be chosen inside every quote interval. A diagonal sequence with $\theta \downarrow 0$ is therefore quote-compatible and has projection converging to $G_K(y)$. Hence

$$G_K(\mathcal{F}_K(I)) \subseteq \mathcal{R}_K(I),$$

which proves the equality. $\square$

A.1.3 Finite termination of adaptive hull reconstruction

Proposition A.2 (Finite termination of adaptive hull reconstruction). *Let $P \subset \mathbb{R}^2$ be a nonempty compact convex polygon, possibly a point or a segment. For $u \in \mathbb{R}^2 \setminus \{0\}$, define*

$$h_P(u) = \max_{x \in P} u^\top x, \quad x(u) \in \arg \max_{x \in P} u^\top x,$$

with $x(u)$ selected as an endpoint of the exposed face. At iteration t , write $H_t = \text{conv}(V_t)$. For an edge e of H_t , let n_e be its outward unit normal and v_e either endpoint, and define

$$\gamma_t(e) = h_P(n_e) - n_e^\top v_e.$$

In exact arithmetic, the dimension probes of Appendix A.2 identify a point or a line segment. If P is two-dimensional, repeatedly adding $x(n_e)$ whenever $\gamma_t(e) > 0$ terminates after finitely many support queries with $H_t = P$.

Proof. Every collected support point belongs to P , so $H_t \subseteq P$. If all initial points equal q , the opposing coordinate supports imply $x_i = q_i$ for $i = 1, 2$ and every $x \in P$, hence $P = \{q\}$. If the initial points span a line $\ell = \{x : n^\top x = c\}$, the probes in directions n and $-n$ establish $P \subseteq \ell$ exactly when their support values are c and $-c$; otherwise one probe returns a noncollinear point. In the collinear case, the opposing coordinate probes select the endpoints of a nontrivial segment.

Suppose now that P is two-dimensional. If $\gamma_t(e) > 0$, then

$$n_e^\top x(n_e) = h_P(n_e) > n_e^\top v_e,$$

so $x(n_e) \notin H_t$, and endpoint selection makes $x(n_e)$ a vertex of P . Every expansion therefore adds a previously absent vertex. Since P has finitely many vertices and each intermediate hull has finitely many edges, only finitely many support queries are made.

At termination every edge gap is zero, and therefore

$$H_t \subseteq P \subseteq \bigcap_{e \in \mathcal{E}(H_t)} \{x : n_e^\top x \leq h_P(n_e)\} = \bigcap_{e \in \mathcal{E}(H_t)} \{x : n_e^\top x \leq n_e^\top v_e\} = H_t.$$

Hence $H_t = P$. □

A.1.4 Proof of Proposition 4.5

Proof. Choose $0 < \beta < 1 - \alpha$ and, for sufficiently small $\varepsilon > 0$, set

$$\mathbb{Q}_\pm = (1 - \alpha - \beta)\delta_0 + \alpha\delta_{K \pm \varepsilon} + \beta\delta_{x_\pm}, \quad x_\pm = \frac{1 - \alpha(K \pm \varepsilon)}{\beta}.$$

Because $K < 1$, these are nonnegative unit-mass, unit-mean distributions with $x_\pm > K$ for all sufficiently small ε , and their probabilities of $\{X \leq K\}$ differ by α .

Couple the threshold atoms with each other and the high atoms with each other. The transport costs are $2\alpha\varepsilon$ and

$$\beta|x_+ - x_-| = 2\alpha\varepsilon,$$

so the Wasserstein-1 distance satisfies $W_1(\mathbb{Q}_+, \mathbb{Q}_-) \leq 4\alpha\varepsilon$. Every call payoff is 1-Lipschitz, so each call-price difference is at most $4\alpha\varepsilon$. The vanilla price vectors converge while the threshold probabilities remain α apart. □

A.2 Computational method and validation

A.2.1 Verification of the strict interior condition

Let $\mathcal{F}_K^0$ collect the structural cell and threshold restrictions (7)–(9) and (11)–(12) before the quote inequalities. Solve

$$\begin{aligned} \eta^\star &= \max_{z, m, a_K, \eta} \eta \\ \text{subject to } & (z, m, a_K) \in \mathcal{F}_K^0, \quad 0 \leq \eta \leq \frac{1}{2}, \\ & \underline{c}_i + \eta(\bar{c}_i - \underline{c}_i) \leq c_i(z, m), \quad i \in \mathcal{I}, \\ & c_i(z, m) \leq \bar{c}_i - \eta(\bar{c}_i - \underline{c}_i), \quad i \in \mathcal{I}. \end{aligned}$$

Given positive widths for all retained quote intervals, Assumption 4.2 is equivalent to $\eta^\star > 0$. The numerical acceptance threshold is stated below.

A.2.2 Relaxed system and common budget

For inconsistent cross sections, introduce interval slacks $\xi_i^-, \xi_i^+ \geq 0$ and replace each quote constraint by

$$\underline{c}_i - \xi_i^- \leq c_i(z, m) \leq \bar{c}_i + \xi_i^+.$$

Call $(z, m, a_K, \xi^-, \xi^+)$ feasible for the relaxed system if (z, m, a_K) satisfies (7)–(9) and (11)–(12) (with (12) omitted when $j_K = J$), the slacks are nonnegative, and the displayed inequalities hold for every $i \in \mathcal{I}$.

Weights are set on the complete put wing $\mathcal{I}^C$. If it has N_C distinct strikes and $n_C(k_i)$ observations at strike k_i ,

$$w_i = \frac{1}{N_C n_C(k_i)}, \quad \sum_{i \in \mathcal{I}^C} w_i = 1,$$

so each distinct strike receives equal total weight, divided equally across observations at that strike. Each nested set inherits the weights on its retained constraints, which avoids renormalizing common constraints. The minimum weighted quote loss is

$$\begin{aligned} \delta^\star(\mathcal{I}) &= \min_{z, m, a_K, \xi^-, \xi^+} \sum_{i \in \mathcal{I}} w_i (\xi_i^- + \xi_i^+) \\ \text{subject to } & (z, m, a_K, \xi^-, \xi^+) \text{ is feasible for the relaxed system.} \end{aligned}$$

Under the prespecified common budget $\bar{\delta} \geq 0$, the relaxed region is

$$\mathcal{R}_{K, \bar{\delta}}(\mathcal{I}) = \left\{ G_K(z, m, a_K) \left| \begin{array}{l} \exists \xi^-, \xi^+ \text{ such that} \\ (z, m, a_K, \xi^-, \xi^+) \text{ is feasible for the relaxed system,} \\ \sum_{i \in \mathcal{I}} w_i (\xi_i^- + \xi_i^+) \leq \bar{\delta} \end{array} \right. \right\}.$$

The region is an outer relaxation: it contains every pair generated by a distribution whose weighted quote loss is within the budget, while boundary points of the finite system need not be generated by such a distribution. Because every nested quote set within a cross section is compared under the same $\bar{\delta}$, the resulting regions remain nested. Appendix A.3 describes the empirical choice of $\bar{\delta}$.

A.2.3 Adaptive hull algorithm

After constructing the applicable finite feasible set and checking Assumption 4.2 when required, the reconstruction uses the following support-function oracle.

Setup.

A feasible polytope $\mathcal{Y}$, the projection G_K , and the marginal box

$$M_K = [\underline{p}, \bar{p}] \times [\underline{L}, \bar{L}],$$

whose endpoints come from the four coordinate support bounds. The initial direction set is

$$D_0 = \left\{ \left(\cos \frac{\pi j}{4}, \sin \frac{\pi j}{4} \right) : j = 0, \dots, 7 \right\}.$$

The numerical parameters are the degenerate-width tolerance ε_{deg} , point tolerance ε_{pt} , edge tolerance $\varepsilon_{\text{edge}}$, area tolerances $(\varepsilon_{\text{area}}^{\text{abs}}, \varepsilon_{\text{area}}^{\text{rel}})$, and the maximum number of support evaluations N_{max} . If either marginal range is at most ε_{deg} , the image is classified as numerically degenerate and represented by a point or line segment. Otherwise each coordinate is mapped to its marginal range,

$$\tilde{p} = \frac{p - \underline{p}}{\bar{p} - \underline{p}}, \quad \tilde{L} = \frac{L - \underline{L}}{\bar{L} - \underline{L}},$$

so that the transformed marginal box is $[0, 1]^2$. In the remainder of this description, x , V_t , H_t , and O_t are in these normalized coordinates. A normalized direction $u = (u_p, u_L)$ corresponds to the original support direction

$$a(u) = \left(\frac{u_p}{\bar{p} - \underline{p}}, \frac{u_L}{\bar{L} - \underline{L}} \right),$$

and all inner products with the raw image $G_K y$ use this direction. The support program yields

$$(x(u), \bar{h}(u)),$$

where $x(u)$ is the normalized image of $G_K y_u^*$ and y_u^* is selected by a secondary program: for nonzero u and $v = (-u_L, u_p)$, maximize $a(v)^\top G_K y$ over $y \in \mathcal{Y}$ subject to $a(u)^\top G_K y = \bar{h}_{\text{raw}}(u)$, where $\bar{h}_{\text{raw}}(u)$ is the value of the support program (15) along $a(u)$ in the original coordinates. The secondary program leaves the support value unchanged and, in exact arithmetic, returns an endpoint of the projected exposed face. The primal–dual accuracy check uses the primary solve. The normalized bound is

$$\bar{h}(u) = \bar{h}_{\text{raw}}(u) - \frac{u_p \underline{p}}{\bar{p} - \underline{p}} - \frac{u_L \underline{L}}{\bar{L} - \underline{L}}.$$

In exact arithmetic, $\bar{h}(u) = \max u^\top x$ over the normalized image of $\mathcal{Y}$, and the primal–dual gap is zero.

Initialization.

Evaluate the support program for each $u \in D_0$, set $V_0 = \{x(u) : u \in D_0\}$ after identifying points within ε_{pt} , and let $d = \dim \text{aff}(V_0)$. If $d = 0$, return the identified point. If $d = 1$, let $n \in \mathbb{R}^2$ be a unit normal to the candidate line and solve in directions n and $-n$. Add the two returned support points and recompute the affine dimension using ε_{pt} . If it remains one, return the two extreme endpoints; otherwise continue with $d = 2$.

Refinement.

At iteration t , form $H_t = \text{conv}(V_t)$. For every edge $e \in \mathcal{E}(H_t)$ with outward unit normal n_e and endpoint v_e , solve in direction n_e and compute the numerical edge gap

$$\widehat{\gamma}_t(e) = \max\{0, \bar{h}(n_e) - n_e^\top v_e\}, \quad \Gamma_t = \max_{e \in \mathcal{E}(H_t)} \widehat{\gamma}_t(e).$$

Add every $x(n_e)$ satisfying $\widehat{\gamma}_t(e) > \varepsilon_{\text{edge}}$ that is distinct within ε_{pt} , and recompute the hull. To preserve numerical nesting, replace each collected upper bound by

$$\bar{h}_t^{\text{cons}}(u) = \max \left\{ \bar{h}(u), \max_{v \in H_t} u^\top v \right\}.$$

Using all direction and upper-bound pairs collected through iteration t , form

$$O_t = \text{conv} \left(\left[[0, 1]^2 \cap \bigcap_{\text{collected } (u, \bar{h})} \{x : u^\top x \leq \bar{h}_t^{\text{cons}}(u)\} \right] \cup H_t \right).$$

The conservative upper bounds and the final convexification enlarge the numerical outer approximation and enforce inner–outer nesting.

Convergence.

The edge criterion is met when $\Gamma_t \leq \varepsilon_{\text{edge}}$. Under relaxation with the common budget, refinement may also end when

$$\mathcal{A}_{p_0}(O_t) - \mathcal{A}_{p_0}(H_t) \leq \varepsilon_{\text{area}}^{\text{abs}} + \varepsilon_{\text{area}}^{\text{rel}} \mathcal{A}_{p_0}(O_t),$$

where the area functional is evaluated after mapping H_t and O_t back to the original (p, L) coordinates. If the limit on support evaluations is reached first, or if a sweep with a positive gap adds no distinct point, the procedure ends without asserting the edge criterion. Otherwise set V_{t+1} to the expanded support set and repeat.

Transformation.

For a nondegenerate return, map H_t and O_t back to the original (p, L) coordinates as H and O . Map their edges by $\phi(p, L) = (p, L/p)$ and evaluate (16).

Primal and dual feasibility tolerances are 10^{-10} . The strict-interior cutoff is 10^{-8} , and the degenerate-width, point, and normalized edge tolerances are 10^{-10} , 10^{-9} , and 10^{-8} , respectively. Relaxed cases may also stop when the absolute area gap is below 10^{-12} plus 0.2% of outer area. The edge criterion alone establishes the stated edge tolerance.

A.2.4 Dual support program

Stack the mass and first-moment normalizations as $Ey = \mathbf{q}_E$ with $\mathbf{q}_E = (1, 1)^\top$, and collect every remaining linear inequality, apart from nonnegativity, as $Ay \leq \mathbf{q}_A$. Under the original intervals these are the cell-moment, threshold-atom, and bid–ask quote inequalities. Under relaxation, the quote rows carry

slacks and the common budget adds one row. Let G be the two-row matrix that implements G_K on y , with zero columns for any slack variables. For direction u , the dual of (15) is

$$\begin{aligned} \min_{\lambda, \nu} \quad & \mathbf{q}_A^\top \lambda + \mathbf{q}_E^\top \nu \\ \text{subject to} \quad & A^\top \lambda + E^\top \nu \geq G^\top u, \\ & \lambda \geq 0, \quad \nu \text{ free.} \end{aligned}$$

Quote, cell, atom, and, when present, common budget constraints contribute distinct blocks of λ .

A.2.5 Transformed area calculation

For a polygonal set C , both section boundaries are affine between successive vertex probabilities. If $L_C^\pm(p) = \alpha_\pm + \beta_\pm p$ on $[p_1, p_2]$, that interval contributes

$$(\alpha_+ - \alpha_-) \log \frac{p_2}{p_1} + (\beta_+ - \beta_-)(p_2 - p_1)$$

to (16). Summing these contributions evaluates the transformed area without numerical quadrature.

A.2.6 Synthetic validation

Five zero-width synthetic designs are evaluated at both $K = 0.85$ and $K = 0.90$. For each $x \in \{K, K - 10^{-6}, K + 10^{-6}\}$, one design assigns weights $(0.15, 0.70, 0.15)$ to $(x, 1, 2 - x)$. One design places a small weight far in the tail, assigning $(0.40, 0.59, 0.01)$ to $(0.9, 1, 5)$, and the bimodal design assigns equal weights to $(0.82, 1.18)$. Each distribution has mean one. Across all ten combinations of a design and a threshold, the LP remains feasible when both known coordinates are imposed, and the known (p_K, L_K) lies in the numerical outer envelope. Point and line projections are classified as degenerate. The zero-width designs validate containment and numerical accuracy, while the positive-width empirical cases use Assumption 4.2.

A.3 Supplementary empirical evidence

A.3.1 Sampling details

Within each calendar week, Wednesday is preferred. If it is absent or ineligible, the fallback order is Tuesday, Monday, Thursday, Friday. Ties between eligible expiries are resolved in favor of the earlier expiry. For the put pair, if both selections coincide with a quote exactly at K , the pair retains that strike and the next-nearest distinct strike by absolute distance.

A.3.2 Qualification and numerical verification

Before contract-identity screening, 2,076 feasible cases satisfy the strict interior cutoff, with a smallest relative margin of 0.02. The remaining 16 feasible cases have zero maximum margin, below the numerical cutoff of 10^{-8} .

The common relaxation budget is the larger of the 95th percentile of $\delta^\star$ in the calibration sample (zero) and the prespecified floor based on normalized quote tick size. This gives $\bar{\delta} = 5 \times 10^{-5}$, fixed before analysis of the main sample. Every nondegenerate case under the original constraints satisfies the normalized edge criterion. Replacing midpoint areas by the endpoints of their numerical envelopes leaves the reported medians unchanged at displayed precision.

A.3.3 Subsample stability and robustness

Table A1: Primary-sample stability of area ratios under the complete put wing

Subsample	Nondegenerate cases	Median area ratio
All primary cases	1,966	0.6340
Excluding calibration dates	1,946	0.6338
2013–2015	578	0.6233
2016–2018	542	0.6336
2019–2021	546	0.6432
2022–August 2023	300	0.6367

Notes: All rows use the primary sample definition. Case counts report observations entering the area-ratio median. Calibration dates are those used to set the common budget before analysis of the main sample. Lower ratios indicate tighter joint sets relative to the marginal benchmark, and similar values across rows indicate stability over time and sample composition.

Table A2: Range of median area ratios across horizon and threshold cells

Specification	Minimum median	Maximum median
Forward –20 bp	0.6200	0.6436
Forward +20 bp	0.6225	0.6458
Doubled bid–ask spread	0.6053	0.6458
50% strike thinning	0.6162	0.6378
Narrowest quoted interval per strike	0.6237	0.6462
$p_0 = 10^{-5}$	0.6225	0.6441
$p_0 = 10^{-3}$	0.6373	0.6639
Long maturity	0.6336	0.6527
Threshold $K = 0.80$	0.6128	0.6316
Threshold $K = 0.95$	0.6396	0.6556
Full call–put cross section	0.6301	0.6926

Notes: These sensitivity summaries pool feasible and relaxed cases after contract screening. Strike thinning retains every other distinct strike after sorting from lowest to highest. The long-maturity check uses an eligible expiry in 61–120 days closest to 90 days. Each ratio compares a specification with its own marginal benchmark, so the rows measure stability rather than relative performance.

A.3.4 Price-grid benchmark

For support cap $x_{\max} \in \{2, 3\}$ and N_{grid} nodes, the base grid is $x_j = jx_{\max}/(N_{\text{grid}} - 1)$ for $j = 0, \dots, N_{\text{grid}} - 1$. The inserted-node design replaces the nearest interior node by K and reorders the grid, resolving an equal-distance tie in favor of the smaller node. The centered-cell design uses interior nodes $K + (\ell + \frac{1}{2})x_{\max}/(N_{\text{grid}} - 1)$ for $\ell \in \mathbb{Z}$, retaining the $N_{\text{grid}} - 2$ nearest points in $(0, x_{\max})$ and adding 0 and $x_{\max}$.

Let c index a case and g a grid specification. Let $\delta_{\text{grid},g,c}^*$ be the minimum weighted quote loss for the grid formulation, and let $\delta_{\text{cont},c}^*$ be the continuous-state minimum computed before imposing the support cap. The shared budget is

$$\delta_{g,c} = \max\{\delta_{\text{grid},g,c}^*, \delta_{\text{cont},c}^*\} + 10^{-10}. \quad (17)$$

Once the shared budget in (17) is chosen from the grid and uncapped continuous minima, the matched continuous support program imposes $0 \leq X \leq x_{\max}$, equivalently $m_J \leq x_{\max}z_J$ in the terminal cell.

Directions are

$$D_{64} = \left\{ \left(\cos \frac{\pi \ell}{32}, \sin \frac{\pi \ell}{32} \right) : \ell = 0, \dots, 63 \right\}.$$

For case c , grid g , and $u \in D_{64}$, let $h_{g,c}(u)$ and $h_{\text{cont},c}(u)$ be the support values from the grid and matched continuous-state programs, both under the shared budget $\delta_{g,c}$. The reported absolute error is

$$e_{g,c}(u) = |h_{g,c}(u) - h_{\text{cont},c}(u)|,$$

$$e_{g,c}^{\max} = \max_{u \in D_{64}} e_{g,c}(u).$$

For each case, the median over directions is $\text{med}_{u \in D_{64}} e_{g,c}(u)$; figures and tables report its median across cases, and the tails of the case-level distribution summarize $e_{g,c}^{\max}$. The four benchmark synthetic cases at $K = 0.90$ use the three threshold locations and the bimodal law defined in the synthetic validation subsection.

Table A3: Selected price-grid errors and computation times

Nodes	Cap	Threshold placement	Median error	Maximum error	Grid budget	Seconds
51	2	Centered cell	0.01136	0.10134	1.63×10^{-4}	0.138
51	2	Inserted node	0.00956	0.15000	1.22×10^{-4}	0.152
201	2	Centered cell	0.00122	0.07739	1.0×10^{-10}	0.419
201	2	Inserted node	1.52×10^{-10}	0.15000	1.0×10^{-10}	0.382
401	2	Centered cell	0.00026	0.07597	1.0×10^{-10}	0.781
401	2	Inserted node	$< 10^{-14}$	0.14999	1.0×10^{-10}	0.808

Notes: Grid budget is the median additional weighted quote loss beyond the uncapped continuous-state minimum, so the reported error includes this feasibility adjustment, especially at 51 nodes. Lower errors indicate closer agreement with the continuous-state solution, and a lower grid budget indicates less additional quote relaxation. Maximum error is the largest $e_{g,c}^{\max}$ across cases. Seconds report median computation time for each grid specification.

References

- Yacine Aït-Sahalia and Andrew W. Lo. Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, 53(2):499–547, 1998. doi: 10.1111/0022-1082.215228.
- David Backus, Mikhail Chernov, and Ian Martin. Disasters implied by equity index options. *The 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 *Institute of Mathematical Statistics 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 Natasha Bushueva. Option pricing without price dynamics: A probabilistic approach, 2006. arXiv:math/0612075.
- 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.
- Tim Bollerslev and Viktor Todorov. Tails, fears, and risk premia. *The Journal of Finance*, 66(6): 2165–2211, 2011. doi: 10.1111/j.1540-6261.2011.01695.x.
- Douglas T. Breeden and Robert H. Litzenberger. Prices of state-contingent claims implicit in option prices. *The Journal of Business*, 51(4):621–651, 1978. doi: 10.1086/296025.

- Samuel N. Cohen, Christoph Reisinger, and Sheng Wang. Detecting and repairing arbitrage in traded option prices. *Applied Mathematical Finance*, 27(5):345–373, 2020. doi: 10.1080/1350486X.2020.1846573.
- Laurent Cousot. Conditions on option prices for absence of arbitrage and exact calibration. *Journal of Banking & Finance*, 31(11):3377–3397, 2007. doi: 10.1016/j.jbankfin.2007.04.006.
- 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.
- Tjeerd De Vries. Recovering risk-neutral moments from options, 2026. arXiv:2601.14852.
- Matthias Ehrgott, Andreas Löhne, and Lizhen Shao. A dual variant of Benson’s outer approximation algorithm for multiple objective linear programming. *Journal of Global Optimization*, 52(4):757–778, 2012. doi: 10.1007/s10898-011-9709-y.
- Matthias R. Fengler. Arbitrage-free smoothing of the implied volatility surface. *Quantitative Finance*, 9(4):417–428, 2009. doi: 10.1080/14697680802595585.
- Jens Carsten Jackwerth and Mark Rubinstein. Recovering probability distributions from option prices. *The Journal of Finance*, 51(5):1611–1631, 1996. doi: 10.1111/j.1540-6261.1996.tb05219.x.
- 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.
- S. N. Lahiri. *Resampling Methods for Dependent Data*. Springer Series in Statistics. Springer, New York, 2003. doi: 10.1007/978-1-4757-3803-2.
- Bryan Liang, Marcel Nutz, Shunan Sheng, and Valentin Tissot-Daguette. Bid–ask martingale optimal transport, 2026. arXiv:2603.24605.
- Andreas Löhne and Benjamin Weißing. Equivalence between polyhedral projection, multiple objective linear programming and vector linear programming. *Mathematical Methods of Operations Research*, 84(2):411–426, 2016. doi: 10.1007/s00186-016-0554-0.
- Ian W. R. Martin and Ran Shi. Forecasting crashes with a smile. CEPR Discussion Paper DP21236, Centre for Economic Policy Research, 2026. URL <https://cepr.org/publications/dp21236>.
- Robert C. Merton. Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1):141–183, 1973. doi: 10.2307/3003143.
- Jean-Baptiste Monnier. Risk-neutral density recovery via spectral analysis. *SIAM Journal on Financial Mathematics*, 4(1):650–667, 2013. doi: 10.1137/110840340.
- Ariel Neufeld, Antonis Papapantoleon, and Qikun Xiang. Model-free bounds for multi-asset options using option-implied information and their exact computation. *Management Science*, 69(4):2051–2068, 2023. doi: 10.1287/mnsc.2022.4456.
- Conall O’Sullivan and Yan Wang. Nonparametric option implied tail risk and market returns. Research Paper 18-8, Michael J. Brennan Irish Finance Working Paper Series, 2018. SSRN 3251803.
- Hao Qin, Ruozhong Yang, Charlie Che, and Liming Feng. Arbitrage-free multi-maturity risk-neutral marginals, 2026. arXiv:2607.06204.

- R. Tyrrell Rockafellar. *Convex Analysis*. Number 28 in Princeton Mathematical Series. Princeton University Press, Princeton, 1970.
- R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *The Journal of Risk*, 2(3):21–41, 2000. doi: 10.21314/JOR.2000.038.
- R. Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471, 2002. doi: 10.1016/S0378-4266(02)00271-6.
- Wouter J. E. C. van Eekelen. A generalized moment approach to sharp bounds for conditional expectations, 2023. arXiv:2401.00090.
- Grigory Vilkov and Yan Xiao. Option-implied information and predictability of extreme returns. SAFE Working Paper 5, Sustainable Architecture for Finance in Europe, 2013. SSRN 2209654.
- Aaron Wizman, Gabriel Turinici, and Gregory Merran. From arbitrage removal to density extraction: A model-free framework for short-dated options, 2026. arXiv:2605.22792.