

American Option Pricing Around Scheduled Earnings Announcements: A Kou Jump-Diffusion with an Event Jump, an IMEX Finite-Difference Solver, and Evidence from 2020–2025 Equity Option Chains

Gleb Alikhver*

October 15, 2026

Abstract

A scheduled earnings announcement concentrates much of a short-dated option’s variance at one known instant. I price American equity options under a Kou jump-diffusion with a compensated double-exponential jump on the announcement date (Kou-SEJ), solved by an IMEX-CNAB finite-difference scheme that treats the announcement as a jump condition and restarts after it. With space and time refined together, the scheme converged at observed order two against Fourier-cosine prices and matched published American put values to within 2.3×10^{-5} . In a study of 9,996 announcements by 1,281 firms in 2020–2025, pre-registered in a plan file (no external registry), with a development sample through 2021 and a holdout priced only after a recorded specification freeze, I compare nine models on the 6,893 announcements on which all are defined (development: 2021 only). Black–Scholes with historical volatility missed pre-event mids by 11.29 (development) and 13.71 (holdout) volatility points, root-mean-square and vega-scaled as pre-registered. With option-implied announcement variance, Kou-SEJ removed 0.80 [0.79, 0.81] of the benchmark’s holdout dollar error, against 0.42 [0.40, 0.44] with the law fitted to past announcements, but the flat at-the-money market volatility removed at least as much (0.82). Away from the money a double-exponential law with the asymmetry of the firm’s past moves fitted worse than a Gaussian jump, and four of the eight pre-registered predictions failed in both samples. Implied announcement variance exceeded realized variance (ratios 1.70 [1.60, 1.82] and 1.14 [1.09, 1.19]), a difference between two laws, not mispricing; with maturities in trading days the holdout ratio was 1.06 [0.99, 1.13], announcement weeks resampled. On 141 later announcements by 28 large firms, priced before the freeze and evaluated on an expiration outside the calibration, that shape again did not help.

Keywords: American options; jump-diffusion; Kou model; earnings announcements; IMEX time stepping; linear complementarity problem; finite differences.

MSC2020: 91G20, 91G60, 65M06, 65M12, 45K05. **JEL:** C63, G13, C58.

1 Introduction

Four times a year a listed firm reports its earnings on a date that it announces weeks in advance. For an option that expires a few days or weeks after that date, the return on the announcement is a large part of the variance the option has to price. [Patell and Wolfson \(1979\)](#) follow Black–Scholes implied variances through annual earnings disclosures and find systematic shifts around the release; [Patell and Wolfson \(1981\)](#) extend the evidence to quarterly announcements, where

*School of Law, University of Georgia. Email: gleb@alikhver.com. Code: <https://github.com/GlebSoftware/American-Option-Pricing-Around-Scheduled-Earnings-Announcements>; data: see the data and code availability statement at the end.

implied volatility rises into the announcement and falls sharply after it, and [Whaley and Cheung \(1982\)](#) test whether option prices anticipate quarterly announcements efficiently. [Ederington and Lee \(1996\)](#) document the same pattern for scheduled macroeconomic releases in Treasury bond, Eurodollar and Deutsche mark options, and [Donders et al. \(2000\)](#) find it in individual stock options around earnings announcements, together with higher trading volume and a build-up of open interest before the event. The announcement return itself behaves like a jump: the jumps that the nonparametric test of [Lee and Mykland \(2008\)](#) detects in individual United States stocks coincide with scheduled earnings announcements and firm-specific news.

The example that opens [Dubinsky and Johannes \(2004\)](#) shows the magnitudes. Before Intel's July 1997 announcement the at-the-money July call traded at a Black–Scholes implied volatility of 71 percent while the August call traded at 45 percent; the day after the announcement the July call's implied volatility was 43 percent. In a diffusion with constant volatility, variance accrues in proportion to calendar time, so the two maturities would have to share one implied volatility; a deterministic time-varying volatility can match both only by concentrating the variance in the event window, which is the announcement jump in disguise and still makes the announcement return Gaussian. The announcement adds a lump of variance at a known time, so total variance to expiry is the diffusive variance plus a term that does not shrink with the time to expiry ([Dubinsky and Johannes, 2004](#)). A Black–Scholes price with an inflated volatility can match one quote but not the pre-event smile, whose wings reflect the fat tails and asymmetry of the announcement return, and it misstates the time decay of every position. A Poisson jump-diffusion ([Merton, 1976](#); [Kou, 2002](#)) has the right tail shape but the wrong timing: it spreads the event over the option's life, assigns positive probability to no move and to several moves, and ties the jump variance to the time to expiry. Section 5 compares the three specifications at equal total variance.

A second strand of the literature measures how the market prices announcement risk, and it forces a distinction that I keep throughout. [Dubinsky and Johannes \(2004\)](#) estimate the risk-neutral announcement volatility from the term structure of implied volatility before the announcement and from its change across it, and find that it exceeds the volatility of realized announcement returns, which they read as a volatility risk premium. [Dubinsky et al. \(2019\)](#) develop this into reduced-form models and estimators that separate announcement uncertainty from day-to-day volatility, and find that the anticipated uncertainty is large, varies over time and forecasts future return volatility. [Barth and So \(2014\)](#) find that the premiums in option prices around announcements are associated with the non-diversifiable component of announcement volatility risk and are concentrated among bellwether firms. [Kelly et al. \(2016\)](#) show that options whose lives span scheduled political events are more expensive than neighboring options, with a steeper smirk and a larger variance risk premium. [Gao et al. \(2018\)](#) find that single-stock straddles lose money on average but that at-the-money straddles bought three days before an earnings announcement and held to the announcement date return 3.34 percent on average, which they read as investors underestimating announcement uncertainty. [Alexiou et al. \(2025\)](#) find that short-dated implied volatility curves often turn concave before earnings announcements, which points to a bimodal risk-neutral law of the announcement return, and that firms with concave curves have larger announcement moves. The lesson for a pricing study is that the difference between the option-implied and the historical announcement variance is a difference between two laws: a risk premium in the reading of [Dubinsky and Johannes \(2004\)](#) and [Barth and So \(2014\)](#), underestimated uncertainty in the reading of [Gao et al. \(2018\)](#). In either reading it is not an error of the pricing model, up to estimation error. A model whose announcement law is estimated from past returns and priced against market quotes therefore measures premium plus model error; a model whose announcement variance is calibrated to the same chain can be judged only on the strikes and maturities it was not fitted to. I report both and say which object each comparison measures.

Models. Merton (1976) prices options when returns mix a diffusion with a compound Poisson jump component whose sizes are lognormal, and Kou (2002) replaces the lognormal by the double exponential, which matches the leptokurtosis of returns and the volatility smile while keeping closed-form prices for vanilla and several path-dependent options; Kou and Wang (2004) extend the tractability to first-passage, barrier, lookback and perpetual American problems and give an analytic approximation for finite-horizon American options in the manner of Barone-Adesi and Whaley (see also Cont and Tankov, 2004). Two features of the double exponential matter here. Its memoryless tails give the integral operator closed-form weights and an evaluation whose cost is linear in the grid size (Toivanen, 2008), and its shape, a sign and an exponentially distributed size on each side, describes an announcement return whose direction and magnitude are both uncertain, with fatter tails than a Gaussian of the same variance. Its density peaks at a zero move, however, while the concave pre-announcement curves of Alexiou et al. (2025) suggest a law with two modes. Jumps at deterministic times have a history in the yield-curve literature: Piazzesi (2001, 2005) model the Federal Reserve’s target as a jump process whose intensity depends on the calendar of Federal Open Market Committee meetings, and Beber and Brandt (2006) compare option-implied state-price densities of Treasury bond prices shortly before and after scheduled macroeconomic announcements. For equities, Dubinsky and Johannes (2004) build a no-arbitrage model with a jump of random size at the known announcement time, note that absence of arbitrage requires the expected price jump at a deterministic time to be zero, and report that adding the announcement jump to a stochastic-volatility model roughly halves its pricing errors around earnings dates. The specification closest to mine is Leung and Santoli (2014), who combine Kou dynamics with a double-exponential (or Gaussian) jump on the announcement date, derive pricing formulas and implied-volatility asymptotics, calibrate to the pre-event surface and compare the risk-neutral jump law with its historical counterpart. For the American put before the announcement they set up two linear complementarity problems linked at the announcement by the jump integral, solve them backward in time by the Fourier space time stepping of Jackson et al. (2008), show numerically that the exercise boundary falls steeply as the announcement approaches, note by Jensen’s inequality that an imminent announcement leaves no exercise region, and give a Barone-Adesi–Whaley-type approximation (Leung and Santoli, 2014, Section 6). I adopt their specification under the name Kou jump-diffusion with a scheduled earnings jump (Kou-SEJ), with the jump compensated so that it is mean-preserving for the stock price, and their two-stage formulation of the American problem.

Numerical methods. Pricing under Kou-SEJ means solving a partial integro-differential equation (PIDE) with a jump condition at the announcement, and, for American options, a linear complementarity problem (LCP). The finite-difference treatment of the jump-diffusion PIDE is mature. Andersen and Andreasen (2000) pair an alternating-direction implicit scheme with FFT evaluation of the integral; d’Halluin et al. (2005) prove stability and convergence of an implicit scheme with a fixed-point iteration for the integral, computed as a correlation by FFT with wrap-around control; Almendral and Oosterlee (2005) use the second-order backward differentiation formula with an iterative splitting of the system matrix; Cont and Voltchkova (2005) analyze an explicit-implicit scheme, its localization error and kernels of infinite activity; Briani et al. (2007) apply implicit-explicit Runge–Kutta methods; Feng and Linetsky (2008) extrapolate the IMEX Euler scheme to high order; Kwon and Lee (2011a) prove stability and second-order convergence of a three-time-level scheme with tridiagonal systems; and Toivanen (2008) gives the algorithm of optimal cost for the Kou integral that I use, together with European and American reference prices. The implicit-explicit multistep construction itself, with the differential operator implicit by Crank–Nicolson and the integral explicit by second-order Adams–Bashforth (IMEX-CNAB), goes back to Ascher et al. (1995); Frank et al. (1997) analyze the linear stability of its popular second-order members, and Salmi and Toivanen (2014) study IMEX-midpoint, IMEX-CNAB and IMEX-BDF2 families for the jump-diffusion PIDE by Fourier analysis and,

under their assumptions and time-step restrictions, find the IMEX-CNAB and IMEX-BDF2 families absolutely stable for every value of the split parameter and the IMEX-midpoint family only when the term $-\lambda u$ is wholly implicit. For the early-exercise constraint, Jaillet et al. (1990) give the variational-inequality formulation and justify the algorithm of Brennan and Schwartz (1977); Cryer (1971) solves the discrete LCP by projected over-relaxation; Forsyth and Vetzal (2002) show that a penalty method attains quadratic convergence only with a time-step selector, and d’Halluin et al. (2004) extend it to jump diffusions; Ikonen and Toivanen (2004) introduce the operator-splitting treatment of the constraint; and Kwon and Lee (2011b), Salmi and Toivanen (2011) and Salmi and Toivanen (2012) develop, compare and survey these solvers for finite-activity jump diffusions. Because the payoff is not smooth, Crank–Nicolson needs the damped start of Rannacher (1984), whose effect on prices and derivatives Giles and Carter (2006) analyze; Pooley et al. (2003) show that payoff smoothing, grid shifting and projection help only together with a suitable time-stepping scheme. Lattices for jump diffusions exist (Amin, 1993; Hilliard and Schwartz, 2005), as do transform methods: the Fourier-cosine expansion of Fang and Oosterlee (2008), my European benchmark, and its extension to early exercise (Fang and Oosterlee, 2009), and the Fourier space time-stepping of Jackson et al. (2008), which also prices Bermudan and American options; Toivanen (2010) gives a high-order front-tracking finite-difference method for American options under jump-diffusion models, with Kou dynamics among its examples. A condition at a known date inside a finite-difference march is standard for discrete dividends: the option value is carried across the ex-dividend date by the deterministic shift $S \rightarrow S - D$, interpolated on the grid, and the exercise check follows (Leentvaar and Oosterlee, 2006); see also Haug et al. (2003) on the discrete-dividend problem. The announcement replaces that shift by a jump of random size, an integral operator applied once, and Leung and Santoli (2014) already carry the American problem across it in Fourier time stepping. In a finite-difference multistep scheme the jump does more: it invalidates the Adams–Bashforth history and the smoothing of the Rannacher start, which must be restarted after it, and the projection has to act on the post-jump value. That step and its verification are the numerical content of this paper.

Estimation. The parameters of the Kou model are estimated from daily returns by maximum likelihood in Ramezani and Zeng (2007), who fit index and individual-stock data and compare the fit with lognormal jump-diffusion and geometric Brownian motion by an information criterion. Jump-diffusion likelihoods are treacherous: the likelihood of a mixture of normals is unbounded as the variance of one component tends to zero, and Honoré (1998) documents this unboundedness for the likelihood of a jump-diffusion, whose unconstrained maximization is therefore unreliable. Aït-Sahalia (2004) shows that Brownian noise can be disentangled from jumps as the sampling interval shrinks, Mancini (2009) separates them by a threshold rule, and Lee and Mykland (2008) test for jumps at intraday frequency. None of this resolves the identification problem that matters here. From a few years of daily returns the intensity, the sign probability and the two tail parameters of a Poisson jump component trade off against the diffusion coefficient, and a threshold fixed at a multiple of the sample standard deviation flags a fixed fraction of observations under any continuous law and so does not identify an intensity. The scheduled jump is different: its timing is known, so the announcement returns themselves, one per quarter, are a sample from its law. The sample is small: among the holdout events with at least 4 prior announcements of the same firm, the median event had 11, and no event had more than 21. I estimate the law from each firm’s own past announcement returns, measure by Monte Carlo what such histories can reveal, compare the law with the one implied by the pre-event option chain through the two-maturity decomposition of Dubinsky and Johannes (2004), and estimate the background diffusion from returns outside the announcement windows.

Findings. Black–Scholes with a background volatility estimated outside past announcement windows (BS-HV) was far too cheap before announcements, with mean errors of -10.25 and

–12.79 implied-volatility points in the development and holdout samples (on the common sample of the nine models, whose development part holds the announcements of 2021 only). Adding the firm’s historical announcement variance, which uses no option price, cut BS-HV’s root-mean-square error in the holdout, under the pre-registered metric, by 38.9% in DJ-hist, a Gaussian announcement law with the firm’s mean squared past move as its variance, and by 36.1% in Kou-SEJ-P, a double-exponential law whose variance is that second moment net of the diffusion in the announcement window; the two differ in the variance as well as in the shape of the law. The option-implied announcement variance cut it by 69.1% in Kou-SEJ-Q, the Kou-SEJ model calibrated to the at-the-money quotes of the two nearest expirations that contain the announcement, but most of that advantage is the at-the-money level of the event’s own chain, to which these models are fitted: the market’s at-the-money volatility, held flat across strikes, did at least as well. Away from the at-the-money strike of the two calibration expirations, which hold almost all quotes, the double-exponential shape did worse than a Gaussian announcement law with the same level. On the few quotes of other expirations (624 holdout quotes in 50 events) it was slightly better, though not significantly, and on the later announcements of Section 8.10, evaluated on an expiration that neither calibrated model uses, slightly worse. The market smile was skewed toward low strikes. The Kou-SEJ-Q smile, whose shape comes from the asymmetry and tail ratio of the firm’s few past moves, lay above the flat Gaussian level on both sides and most of all at high strikes. The option-implied announcement variance exceeded the realized one by factors of 1.70 in the development sample and 1.14 in the holdout, a difference between the risk-neutral and the realized law that I keep apart from pricing error throughout; with maturities counted in trading days the holdout ratio was 1.06, within clock and term-structure uncertainty of one (Section 6.7). The American price helped for calls with an ex-dividend date in their life and not for puts. Of the eight pre-registered predictions, P5 and P7 held in both samples, P3 and P8 in the holdout only, and P1, P2, P4 and P6 failed in both; Section 8.9 reports every component and Section 9 what the failures teach.

Contributions.

- A finite-difference realization of the two-stage American problem of [Leung and Santoli \(2014\)](#) (Section 4). It uses the IMEX-CNAB scheme of [Salmi and Toivanen \(2014\)](#), with the whole differential operator implicit and only the integral explicit, the evaluation of the Kou integral of [Toivanen \(2008\)](#), whose cost is linear in the number of grid points, and the LCP solvers of Section 4.5; it applies the announcement once as a discrete operator with exact quadrature weights for the shifted announcement density and closed-form tails, restarts the Adams–Bashforth history and the Rannacher start after it, and applies the exercise decision to the post-jump value. A proof, for every convex payoff that is not affine and every announcement law of full support, of the no-exercise property at the instant before a compensated announcement that [Leung and Santoli \(2014\)](#) note for the put by Jensen’s inequality, of the nonnegativity of the announcement optionality, and of the collapse of the put’s exercise boundary as the announcement approaches, whose steep fall they show numerically (Section 3.3). A verification against Fourier-cosine prices, an exact benchmark for the scheduled jump built from the product of the Kou and announcement characteristic functions, the American put values that [Salmi and Toivanen \(2011\)](#) list for the test case of [Toivanen \(2008\)](#), and the published American prices with an announcement of [Leung and Santoli \(2014\)](#), with the orders as observed, the temporal order of the American scheme measured on its own, and a list of what the suite does not test (Section 5).
- The estimators of both announcement laws, examined at the sample sizes of the data (Section 6). A firm’s historical announcement variance is unbiased, but at $n = 12$, the Monte Carlo history length nearest the holdout median of 11 prior announcements, its standard deviation is 0.66 times its true value; its double-exponential shape is not identified

firm by firm; and a background Poisson intensity is not identified from a trailing window of returns, so the models keep their jump content in the announcement, whose timing is known. The option-implied announcement variance of the two-expiry decomposition (Dubinsky and Johannes, 2004) is noisy event by event and unbiased on average, and its ratio to the realized variance measures the earnings variance premium of the sample.

- A pre-registered event-level pricing-error study (a plan file of the replication package, not deposited with an external registry; Section 8) of 9,996 announcements by 1,281 firms between 22 January 2020 and 22 April 2025 (Sections 7 and 8): nine models priced with the information of the last close before each announcement, a development sample through 2021 on which the specification was frozen, a holdout from 2022 priced afterwards with the frozen code and priced again with it after a recorded data correction (deviation (viii) of Section 8.1), every pre-registered test reported with its verdict, and every deviation from the plan labeled. A post-freeze analysis, labeled as such, recalibrates the American model to the American quotes (Appendix F.5).
- An out-of-calibration check on 141 announcements by 28 large firms between 1 May 2025 and 10 October 2026, after the frozen database ends, evaluated on an expiration that neither calibrated model uses (Section 8.10), on a sample priced before the freeze that supplied the primary Kou-SEJ-Q specification, and model prices for all 24 announcements scheduled after the snapshot of 12 October 2026, recorded with a hash before the events for evaluation in a later version (Appendix E).

The rest of the paper is organized as follows. Section 2 sets out the model, its constraints, the compensators and the European price used for verification. Section 3 states the pricing problem (PIDE, jump condition, LCP and far-field conditions) and the results on exercise around the announcement. Section 4 describes the discretization, the IMEX-CNAB time stepping with the announcement step, and the LCP solvers, and Section 5 reports the verification, including the equal-variance comparison of the scheduled, Poisson and Gaussian specifications. Section 6 studies the estimators of the announcement laws and of the background and measures the earnings variance premium. Section 7 describes the option database, the event sample and the filters. Section 8 reports the pre-registered study and, in Section 8.10, the recent announcements. Section 9 discusses the results and their limitations, and Section 10 concludes. Appendices A to D hold the proofs, the implementation details, the derivations and Monte Carlo designs of the estimation, and the data details; Appendix E records the ex-ante model prices, and Appendix F holds the secondary tables of the empirical study, the post-freeze American calibration and the recent results by firm.

2 Model

This section defines Kou-SEJ, derives its compensators and gives the European price that Section 5 uses to verify the solver. Throughout, $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{Q})$ is a filtered probability space carrying the pricing measure $\mathbb{Q}$, and $(\mathcal{F}_t)$ is the filtration generated by the Brownian motion W , the compound Poisson process $\sum_{k \leq N_t} Y_k$ and the process $Z \mathbf{1}_{\{t \geq t_E\}}$ introduced below, augmented by the $\mathbb{Q}$ -null sets. The risk-free rate $r \geq 0$ and the dividend yield $q \geq 0$ are constants, the volatility σ is positive, every rate is per year, and time is measured in trading years. The option matures at $T > 0$; the firm announces at a date t_E that is known at time 0. The announcement is in the option's life when $0 < t_E < T$. An announcement at $t_E \leq 0$ has already happened, and an option that expires at $t_E = T$ is treated as expiring just before the announcement; in both cases the model is that of Kou (2002) (Remark 2.2). Paths are right-continuous with left limits, and S_{t-} denotes the left limit of the price at t .

2.1 Kou dynamics away from the announcement

Away from t_E the ex-dividend price follows the double-exponential jump-diffusion of [Kou \(2002\)](#),

$$\frac{dS_t}{S_{t-}} = (r - q - \lambda\kappa) dt + \sigma dW_t + d\left(\sum_{k=1}^{N_t} (e^{Y_k} - 1)\right), \quad t \neq t_E, \quad (1)$$

where W is a standard $\mathbb{Q}$ -Brownian motion, $(N_t)_{t \geq 0}$ is a Poisson process with intensity $\lambda \geq 0$ per year, and the log jump sizes $Y_1, Y_2, \dots$ are independent with the double-exponential density

$$f_Y(y) = p\eta_1 e^{-\eta_1 y} \mathbf{1}_{\{y \geq 0\}} + (1-p)\eta_2 e^{\eta_2 y} \mathbf{1}_{\{y < 0\}}, \quad p \in [0, 1], \quad \eta_1 > 1, \quad \eta_2 > 0. \quad (2)$$

W , (N_t) and (Y_k) are mutually independent. The constraint $\eta_2 > 0$ makes (2) a density; the constraint $\eta_1 > 1$ makes $\mathbb{E}[e^Y]$ finite when $p > 0$, which is needed for S to have a finite mean (for $p = 0$ the up-jump rate plays no role). The compensator is

$$\kappa = \mathbb{E}[e^Y] - 1 = \frac{p\eta_1}{\eta_1 - 1} + \frac{(1-p)\eta_2}{\eta_2 + 1} - 1 = \frac{p}{\eta_1 - 1} - \frac{1-p}{\eta_2 + 1} \quad (3)$$

(Appendix A.1). With this κ the compensated jump part of (1) is a martingale, so $e^{-(r-q)t}S_t$ is a $\mathbb{Q}$ -martingale between announcements. For $\lambda = 0$ the dynamics are those of [Black and Scholes \(1973\)](#).

2.2 The scheduled announcement jump

Earnings are released at a date that is known in advance. The natural way to put this into the model is one jump at a deterministic time, not a higher Poisson intensity around the event. At t_E the price moves once,

$$S_{t_E} = S_{t_E-} \exp(Z - \kappa_E), \quad Z \sim \text{DE}(p_E, \eta_{1,E}, \eta_{2,E}), \quad \eta_{1,E} > 1, \quad \eta_{2,E} > 0, \quad p_E \in [0, 1], \quad (4)$$

where Z is independent of $(W, (N_t), (Y_k))$, and hence of $\mathcal{F}_{t_E-}$, and has the density f_E of the form (2) with parameters $(p_E, \eta_{1,E}, \eta_{2,E})$, and where

$$\kappa_E = \ln \mathbb{E}[e^Z] = \ln \left(1 + \frac{p_E}{\eta_{1,E} - 1} - \frac{1 - p_E}{\eta_{2,E} + 1} \right) \quad (5)$$

is the *log* compensator of the announcement. No calendar time elapses across the announcement, so the martingale property of the discounted price cannot be restored by a drift term as in (1); it has to be imposed on the jump itself:

$$\mathbb{E}[S_{t_E} | \mathcal{F}_{t_E-}] = S_{t_E-} \mathbb{E}[e^{Z - \kappa_E}] = S_{t_E-}. \quad (6)$$

Equation (6) is the whole content of the compensator: the announcement redistributes probability mass without adding drift. Without it the forward price is wrong by the factor $\mathbb{E}[e^Z] = e^{\kappa_E}$ and put-call parity fails by $S_0 e^{-qT} (e^{\kappa_E} - 1)$; Section 5 shows this failure numerically. The compensator κ_E has the same algebraic form as Kou's κ but enters through a logarithm: the announcement multiplies the price by e^Z exactly once, whereas Poisson jumps are compensated in expectation per unit time. By Jensen's inequality $\mathbb{E}[Z] < \kappa_E$ for every nondegenerate law of Z , so under $\mathbb{Q}$ the expected log return of the announcement, $\mathbb{E}[Z] - \kappa_E$, is negative; under the physical measure the announcement return may carry a premium, a point taken up in Section 6.

Definition 2.1 (Kou jump-diffusion with a scheduled earnings jump). The *Kou-SEJ* model is the price process defined by (1)–(5), with background parameters $\Theta = (\sigma, \lambda, p, \eta_1, \eta_2)$ and announcement parameters $\Theta_E = (t_E, p_E, \eta_{1,E}, \eta_{2,E})$.

The specification is not mine. A jump of random size at a deterministic announcement date was introduced for equity options by [Dubinsky and Johannes \(2004\)](#) (see also [Dubinsky et al., 2019](#)), who placed a Gaussian jump on Black–Scholes dynamics and derived the term-structure estimator of the announcement variance that Section 6 uses. [Leung and Santoli \(2014\)](#) combined Kou dynamics with a double-exponential (or Gaussian) announcement jump, derived closed-form European prices and pre-announcement implied-volatility asymptotics, and wrote the American put before the announcement as two linear complementarity problems linked by the jump integral, which they solved by Fourier space time stepping. Definition 2.1 is their model, which has no dividend yield, written with $q \geq 0$ and in the notation of [Kou \(2002\)](#): their Kou rates (λ_1, λ_2) are (η_1, η_2) here, their intensity κ is λ , their compensator m is κ , and their announcement law $\text{DE}(u, \eta_1, \eta_2)$ is $\text{DE}(p_E, \eta_{1,E}, \eta_{2,E})$.

Remark 2.2 (Nested models). (i) If $t_E \notin (0, T)$ the announcement never enters the option’s life and the model is that of [Kou \(2002\)](#). (ii) If in addition $\lambda = 0$ it is [Black and Scholes \(1973\)](#). (iii) If $\lambda = 0$ and the law of Z is Gaussian, with $\kappa_E = \ln \mathbb{E}[e^Z]$ computed from that law, the model is the event-jump specification of [Dubinsky and Johannes \(2004\)](#); only the shape of f_E changes, and every statement of Sections 3 and 4 that does not use the exponential tails of f_E holds verbatim. In particular the Gaussian law has full support, the condition under which Theorem 3.3 is strict. (iv) A Poisson model of the same event, Kou dynamics with the intensity raised so that the expected number of jumps over the option’s life is one, has the same expected count but a different terminal law: the count is random, so it spends mass on “no move” and on “two or more moves”. Section 5 compares the two specifications at equal total variance (Table 4).

2.3 Terminal distribution and European prices

Let $\varrho_T = \ln(S_T/S_0)$ and write

$$\mu = r - q - \lambda\kappa - \frac{1}{2}\sigma^2 \quad (7)$$

for the drift of the log price between jumps. The characteristic functions of Y and Z are

$$\phi_Y(v) = \frac{p\eta_1}{\eta_1 - iv} + \frac{(1-p)\eta_2}{\eta_2 + iv}, \quad \phi_Z(v) = \frac{p_E\eta_{1,E}}{\eta_{1,E} - iv} + \frac{(1-p_E)\eta_{2,E}}{\eta_{2,E} + iv}, \quad v \in \mathbb{R}. \quad (8)$$

Proposition 2.3 (Terminal law and European prices). *Assume $t_E \in (0, T)$.*

(i) *The characteristic function of ϱ_T is the product*

$$\phi_{\varrho_T}(v) = \exp \left\{ T \left[iv\mu - \frac{1}{2}\sigma^2 v^2 + \lambda(\phi_Y(v) - 1) \right] \right\} \cdot e^{-iv\kappa_E} \phi_Z(v), \quad (9)$$

and $\phi_{\varrho_T}(-i) = e^{(r-q)T}$. For $t_E \notin (0, T)$ the second factor is absent.

(ii) *The cumulants of ϱ_T are the sums of the Kou cumulants and the cumulants of $Z - \kappa_E$,*

$$\begin{aligned} c_1 &= T \left(\mu + \lambda \left(\frac{p}{\eta_1} - \frac{1-p}{\eta_2} \right) \right) + m_1 - \kappa_E, & c_2 &= T \left(\sigma^2 + 2\lambda \left(\frac{p}{\eta_1^2} + \frac{1-p}{\eta_2^2} \right) \right) + m_2 - m_1^2, \\ c_4 &= 24\lambda T \left(\frac{p}{\eta_1^4} + \frac{1-p}{\eta_2^4} \right) + m_4 - 4m_3m_1 - 3m_2^2 + 12m_2m_1^2 - 6m_1^4, \end{aligned} \quad (10)$$

where $m_n = \mathbb{E}[Z^n] = n! (p_E \eta_{1,E}^{-n} + (1-p_E)(-1)^n \eta_{2,E}^{-n})$.

(iii) *(Mixture representation.) If $\lambda = 0$, the time-0 price of a European claim with payoff $\varphi(S_T)$ is the f_E -mixture of Black–Scholes prices,*

$$V^{\text{Eur}}(S_0, 0) = \int_{\mathbb{R}} V^{\text{BS}}(S_0 e^{z - \kappa_E}, 0) f_E(z) dz, \quad (11)$$

where $V^{\text{BS}}(s, 0)$ is the Black–Scholes price of the same claim with spot s , volatility σ and maturity T . For $\lambda > 0$ the same holds with the Kou European price in place of V^{BS} .

(iv) (Put-call parity.) $C^{\text{Eur}}(S_0, 0) - P^{\text{Eur}}(S_0, 0) = S_0 e^{-qT} - K e^{-rT}$.

(v) The law of S_T , and hence the price of every European claim on S_T , does not depend on where t_E lies in $(0, T)$.

The proof is in Appendix A.2. Part (i) feeds the cosine-expansion method of Fang and Oosterlee (2008), whose truncation range is set from the cumulants in (ii); part (iii) gives an independent quadrature benchmark; together they verify the European side of the solver (Section 5). Part (v) fails for American options: the holder of an American option cares about when the uncertainty is resolved, and Section 3 shows how.

3 Pricing problem

This section states the American pricing problem with the announcement as a jump condition, and proves that a put or call is not exercised at the instant before a compensated announcement whose law has full support.

3.1 Coordinates, PIDE and jump condition

Let $V(S, t)$ be the option value. I work in the log price $x = \ln S$ and in time to maturity $\tau = T - t$, and write $u(x, \tau) = V(e^x, T - \tau)$, $\tau_E = T - t_E$, and φ for the payoff as a function of S , so that $u(x, 0) = \varphi(e^x)$. Away from the announcement, u solves the partial integro-differential equation

$$\begin{aligned} \partial_\tau u &= \mathcal{L}u, & \tau \neq \tau_E, \\ \mathcal{L}u &= \frac{1}{2}\sigma^2 \partial_{xx}u + (r - q - \lambda\kappa - \frac{1}{2}\sigma^2)\partial_x u - (r + \lambda)u + \lambda \int_{\mathbb{R}} u(x + y, \tau) f_Y(y) dy. \end{aligned} \quad (12)$$

Equation (12) is the backward equation of the Kou process written in (x, τ) (Cont and Tankov, 2004, Chapter 12): the substitutions $SV_S = u_x$ and $S^2V_{SS} = u_{xx} - u_x$ produce the $-\frac{1}{2}\sigma^2$ in the drift coefficient, and the term $-\lambda u$ that comes with the integral is grouped with the discounting. Two features matter for the discretization. First, the integral is a *correlation*, $\int u(x + y) f_Y(y) dy$: an up-jump $y > 0$ reads the value at a higher log price. Second, the drift coefficient is μ of (7): the compensator $\lambda\kappa$ and the Itô correction $\frac{1}{2}\sigma^2$ both enter it.

The announcement enters as a condition linking the solution on the two sides of τ_E . By the tower property and (4), no discounting takes place across the announcement and $V(S, t_E^-) = \mathbb{E}[V(S e^{Z - \kappa_E}, t_E)]$ for a European claim; in (x, τ) ,

$$\begin{aligned} u(x, \tau_E^+) &= (\mathcal{P}_E u(\cdot, \tau_E^-))(x), \\ (\mathcal{P}_E w)(x) &= \int_{\mathbb{R}} w(x + z - \kappa_E) f_E(z) dz = \int_{\mathbb{R}} w(x + y) f_E(y + \kappa_E) dy. \end{aligned} \quad (13)$$

The operator $\mathcal{P}_E$ is a correlation with the *shifted* density $f_E(\cdot + \kappa_E)$; $\mathcal{P}_E \mathbf{1} = 1$ and $\mathcal{P}_E e^x = e^x$ by (6). For an American option the holder may exercise immediately before the announcement, and the condition becomes $u(x, \tau_E^+) = \max\{\varphi(e^x), (\mathcal{P}_E u(\cdot, \tau_E^-))(x)\}$; Theorem 3.3 shows that for convex φ the maximum is the continuation value $(\mathcal{P}_E u(\cdot, \tau_E^-))(x)$. If $\tau_E \notin (0, T)$ the condition is not applied and the problem is that of Kou (2002).

3.2 The American problem

For $t \in [0, T]$ and $S > 0$ the American value is the optimal-stopping value

$$V(S, t) = \sup_{\vartheta \in \mathcal{T}_{t, T}} \mathbb{E}[e^{-r(\vartheta - t)} \varphi(S_\vartheta) \mid S_t = S], \quad (14)$$

the supremum running over the set $\mathcal{T}_{t,T}$ of $(\mathcal{F}_t)$ -stopping times with values in $[t, T]$; with the right-continuous convention the holder who stops at t_E receives $\varphi(S_{t_E})$ with the post-announcement price. In (x, τ) the value solves the linear complementarity problem

$$\partial_\tau u - \mathcal{L}u \geq 0, \quad u \geq \varphi(e^x), \quad (\partial_\tau u - \mathcal{L}u)(u - \varphi(e^x)) = 0 \quad \text{on } \mathbb{R} \times ((0, T] \setminus \{\tau_E\}), \quad (15)$$

together with the announcement condition (13) with the maximum. Away from the announcement this is the variational-inequality characterization of the American price for jump-diffusions (Pham, 1997), the analog of Jaillet et al. (1990) for diffusions; the condition at τ_E is the dynamic-programming step across a deterministic time, the link between the two complementarity problems of Leung and Santoli (2014, eqs. (6.2)–(6.4)). In calendar time (15) reads

$$\partial_t V + \mathcal{L}_S V - rV \leq 0, \quad V \geq \varphi, \quad (\partial_t V + \mathcal{L}_S V - rV)(V - \varphi) = 0, \quad (16)$$

with $\mathcal{L}_S$ the generator of the Kou process in S (Appendix A.3), so that the operator $\mathcal{L}$ of (12) is $\mathcal{L}_S - r$ written in (x, τ) . The direction of the inequality can be checked inside the exercise region of a put, where $V = K - S$ and $\partial_t V = 0$; there

$$\partial_t V + \mathcal{L}_S V - rV = qS - rK + \lambda \int_{\mathbb{R}} [V(Se^y) - (K - Se^y)] f_Y(y) dy.$$

The integral is nonnegative because $V \geq \varphi$; it vanishes for $\lambda = 0$, and it tends to zero as $S \rightarrow 0$, because only up-jumps beyond the exercise boundary contribute and $V(s) - (K - s) \leq s$. For $q = 0$ and $r > 0$ the expression is therefore $-rK < 0$ without jumps and negative for small S with them. The inequality written the other way around would exclude every small S from the exercise region, which for a put is an interval $(0, S^*]$ (Corollary 3.4), and so would empty it.

In the far field the option behaves like its payoff or its discounted forward. For the put,

$$u(x, \tau) \rightarrow Ke^{-r\tau} - e^{x-q\tau} \quad (x \rightarrow -\infty), \quad u(x, \tau) \rightarrow 0 \quad (x \rightarrow +\infty), \quad (17)$$

with $Ke^{-r\tau} - e^{x-q\tau}$ replaced by $\max\{Ke^{-r\tau} - e^{x-q\tau}, K - e^x\}$ for the American put. For the call the roles of the two ends are exchanged: the far field is zero on the left and, on the right, $e^{x-q\tau} - Ke^{-r\tau}$ for the European call and $\max\{e^{x-q\tau} - Ke^{-r\tau}, e^x - K\}$ for the American call.

Remark 3.1 (Localization). The computation takes place on $[x_c - L, x_c + L]$ with Dirichlet data from (17). For Lévy models whose Lévy measure has exponential tails, as (2) and f_E have, Cont and Voltchkova (2005) show that the error from this truncation decays exponentially in L . Unlike a pure diffusion, the PIDE also needs the solution *outside* the computational domain, because the integral in (12) reaches beyond any finite interval; Section 4 supplies it by the far-field forms (17) integrated in closed form against the tails of f_Y and f_E .

3.3 Exercise around the announcement

Exercise “just before” the announcement is the dynamic-programming step that the numerical scheme implements. With $V(\cdot, t_E) = V(\cdot, t_E^+)$ the American value immediately after the announcement, define the continuation value across the announcement and the pre-announcement value by

$$C_E(S) = \mathbb{E}[V(Se^{Z-\kappa_E}, t_E)], \quad V(S, t_E^-) = \max\{\varphi(S), C_E(S)\}, \quad (18)$$

and the pre-announcement exercise region by $\mathcal{E}_{t_E^-} = \{S > 0 : C_E(S) \leq \varphi(S)\}$. The results below use the compensation (6); the fact that the log price has independent increments that do not depend on its level; the standing assumptions $\sigma > 0$, $r \geq 0$ and $q \geq 0$ of Section 2; and, where stated, a condition on the support of the law of Z .

Lemma 3.2 (Convexity preservation). *Let $\varphi : (0, \infty) \rightarrow [0, \infty)$ be convex with $\varphi(s) \leq a + bs$ for constants $a, b \geq 0$. In the Kou-SEJ model $S \mapsto V(S, t)$ is convex for every $t \in [0, T]$, $S \mapsto C_E(S)$ is convex, and the same holds for the European value.*

The proof (Appendix A.4) writes (14) as a supremum over stopping times of expectations of convex functions of S , which is possible because the log price has independent increments. Convexity preservation is classical (Bergman et al., 1996; El Karoui et al., 1998), and Ekström and Tysk (2007) show that with jumps it can fail when the jump law depends on the state; in exponential Lévy and additive models it is elementary, and Lemma 3.2 is included because Theorem 3.3(ii) and Corollary 3.4 use it.

Theorem 3.3 (No exercise at the announcement instant). *Let φ be as in Lemma 3.2 and let the announcement jump be compensated as in (4)–(6).*

(i) *For every $S > 0$,*

$$C_E(S) \geq \mathbb{E}[\varphi(Se^{Z-\kappa_E})] \geq \varphi(S), \quad (19)$$

and hence $V(\cdot, t_E^-) = C_E$. If, in addition, the law of Z has full support, as $\text{DE}(p_E, \eta_{1,E}, \eta_{2,E})$ with $0 < p_E < 1$ and the Gaussian law do, and φ is not affine on $(0, \infty)$, in particular for the put $(K - s)^+$, the call $(s - K)^+$ and every nonnegative combination of them, then the second inequality is strict for every $S > 0$ and $\mathcal{E}_{t_E^-} = \emptyset$.

(ii) *Moreover $C_E(S) \geq V(S, t_E)$ for every $S > 0$: the announcement optionality $V(S, t_E^-) - V(S, t_E)$ is nonnegative at every spot.*

The proof (Appendix A.5) is Jensen's inequality: $\mathbb{E}[Se^{Z-\kappa_E}] = S$ by (6), and $V(\cdot, t_E) \geq \varphi$ because immediate exercise is feasible; part (i) uses only the convexity of the payoff, part (ii) Lemma 3.2, and the strictness comes from the full support of the law of Z . Without full support the strict inequality can fail. For $p_E = 0$ the announcement can only lower the price, and when $r > 0$ exercising the put immediately before it is optimal at every spot $S \leq e^{\kappa_E} S^*(t_E)$, where $S^*(t_E)$ is the put's exercise boundary just after the announcement: the post-announcement price $Se^{Z-\kappa_E}$ then lies in the exercise region with probability one, so $C_E(S) = K - S$. For $p_E = 1$ the same happens to the call when $q > 0$. Every announcement law that Sections 5–8 estimate, calibrate or use for verification has $0 < p_E < 1$: the sign splits of Section 6 are shares of up moves in samples with moves of both signs, the firm's split falls back to the symmetric law when one sign is missing, and the pre-registered family and the maximum-likelihood fit bound p_E away from 0 and 1. Economically, the holder of a put never exercises into such an announcement, because the announcement is a free, mean-preserving spread of the price and the put is convex in it; the only motive for exercising a put shortly before the announcement is the interest on the strike over the remaining interval, and Corollary 3.4 shows that this motive loses as the interval shrinks. Numerically, the projection in the announcement step of Section 4 is never active for vanilla payoffs, by the weak inequalities in (i), a consistency check that costs nothing (Section 5).

Corollary 3.4 (Collapse of the put boundary before the announcement). *Consider the American put, $\varphi(s) = (K - s)^+$, with $r \geq 0$, $q \geq 0$ and $p_E > 0$. For $t < T$ let $\mathcal{E}_t = \{S > 0 : V(S, t) = (K - S)^+\}$ and $S^*(t) = \sup \mathcal{E}_t$, with $S^*(t) = 0$ if $\mathcal{E}_t = \emptyset$. Then $\mathcal{E}_t$ is either empty or an interval $(0, S^*(t)]$, and*

$$\lim_{t \uparrow t_E} S^*(t) = 0. \quad (20)$$

In the proof (Appendix A.6), for fixed $S \in (0, K)$ the strategy “hold through the announcement and exercise immediately after it” is worth $e^{-r(t_E-t)}\mathbb{E}[(K - S_{t_E})^+]$, which converges to $\mathbb{E}[(K - Se^{Z-\kappa_E})^+] > K - S$ as $t \uparrow t_E$, the inequality being strict because $p_E > 0$ lets $Se^{Z-\kappa_E}$ exceed K with positive probability, so S leaves the exercise region before t_E ; the proof gives no rate. At t_E itself the boundary is that of the Kou model without announcement, so S^* falls to zero as $t \uparrow t_E$ and jumps back to the no-event boundary at t_E . Figure 2 in Section 5.3 shows this shape; on a finite grid the collapse appears as the exercise region reaching the lower edge of the grid and then emptying. The same argument, with $p_E < 1$ in place of $p_E > 0$, shows for the call with $q > 0$ that every fixed S leaves the exercise region as $t \uparrow t_E$, so that the call boundary

$S^{**}(t) = \inf \mathcal{E}_t$ tends to infinity; the region is an interval $[S^{**}(t), \infty)$ because $V - (S - K)$ is convex, nonnegative and bounded by K , hence nonincreasing in S .

Proposition 3.5 (Monotonicity in the announcement date; [Leung and Santoli, 2014](#)). *Let φ be as in Lemma 3.2 and fix $(\Theta, p_E, \eta_{1,E}, \eta_{2,E})$. The time-0 American price is nonincreasing in the announcement date: if $0 < t'_E < t_E < T$, then $V^{t'_E}(S, 0) \geq V^{t_E}(S, 0)$ for every $S > 0$, where the superscript indicates the model's announcement date. The European price is the same for both dates (Proposition 2.3(v)).*

The result is due to [Leung and Santoli \(2014, Proposition 6.1\)](#), who prove it for the American put in the model without dividends; Proposition 3.5 states it for $q \geq 0$ and for every convex payoff of linear growth. I state it because it organizes the numerical results of Section 5 (Table 3), and Appendix A.7 gives a short proof that couples the two models on one probability space. Together with Proposition 2.3(v) it says that the announcement date affects the option only through the early-exercise premium.

Remark 3.6 (Relation to [Leung and Santoli \(2014\)](#)). The model, the closed-form European price behind Proposition 2.3(i), Proposition 3.5 and the two-stage formulation of the American problem, which they solve by Fourier space time stepping, are theirs ([Leung and Santoli, 2014, Section 6](#)), as is the remark, by Jensen's inequality, that an imminent announcement leaves the put no exercise region; Theorem 3.3 states it for every convex payoff that is not affine and every announcement law of full support and adds part (ii), and Corollary 3.4 proves the collapse of the put boundary that their numerical boundaries show. What this paper adds on the pricing side is the finite-difference realization of Section 4 and its verification (Section 5).

4 Numerical method

This section discretizes the problem of Section 3: the grid and the discrete operators, the jump quadrature, the IMEX-CNAB time stepping, the announcement step and the solvers of the complementarity problem.

4.1 Operator splitting, grid and discrete operators

Following [Salmi and Toivanen \(2014\)](#), I split $\mathcal{L} = \mathcal{D} + \mathcal{J}$ into the whole differential operator, which is treated implicitly, and the integral, which is treated explicitly:

$$\mathcal{D}u = \frac{1}{2}\sigma^2 \partial_{xx}u + \mu \partial_x u - (r + \lambda)u, \quad \mathcal{J}u = \lambda \int_{\mathbb{R}} u(x + y, \tau) f_Y(y) dy, \quad (21)$$

with μ the drift (7). The implicit part thus contains the drift, the discounting and the term $-\lambda u$; only the nonlocal part is explicit. [Salmi and Toivanen \(2014\)](#) analyze the multistep IMEX schemes of [Ascher et al. \(1995\)](#) and [Frank et al. \(1997\)](#) for this PIDE in families whose parameter divides the term $-\lambda u$ between the implicit and the explicit part; (21) is the member with $-\lambda u$ wholly implicit. [Kwon and Lee \(2011a\)](#) use a different second-order scheme, a three-time-level IMEX-midpoint scheme in which the integral and the zeroth-order term are explicit.

The grid is uniform in x ,

$$x_i = x_c + (i - \frac{M}{2})\Delta x, \quad i = 0, \dots, M, \quad x_c = \ln K, \quad L = \frac{1}{2}M \Delta x, \quad (22)$$

with M even, so that $\ln K$ is the center node and the kink of the payoff lies on a node. The price at S_0 is read off the node when $\ln S_0$ is one and is otherwise a cubic spline in x through the grid values. The verification ladders at $S_0 \neq K$ center the grid at $x_c = \ln S_0$ instead; then Δx is adjusted from $2L/M$ so that $\ln K$ is also a node (Appendix B.1). The half-width L is the sum of a diffusive reach proportional to $\sigma\sqrt{T}$ and of jump reaches proportional to $1/\min(\eta_1, \eta_2)$

and $1/\min(\eta_{1,E}, \eta_{2,E})$; the constants are design parameters whose effect Section 5 reports. With $S_i = e^{x_i}$ and $\varphi_i = \varphi(S_i)$, the second-order central discretization of $\mathcal{D}$ is the tridiagonal operator

$$(DV)_i = \frac{\sigma^2}{2\Delta x^2}(V_{i+1} - 2V_i + V_{i-1}) + \frac{\mu}{2\Delta x}(V_{i+1} - V_{i-1}) - (r + \lambda)V_i, \quad 1 \leq i \leq M-1, \quad (23)$$

and the matrices of the time stepping are

$$A = I - \frac{\Delta\tau}{2}D, \quad B = I + \frac{\Delta\tau}{2}D, \quad (24)$$

where the rows $i = 0$ and $i = M$ of A are rows of the identity and the corresponding entries of the right-hand side carry the far-field values (17) at the new time level.

Proposition 4.1 (M-matrix). *Let $\sigma > 0$ and μ be as in (7). If*

$$\Delta x \leq \frac{\sigma^2}{|\mu|} \quad (\text{no condition when } \mu = 0), \quad (25)$$

then for every $\Delta\tau > 0$ the matrix A has a positive diagonal, non-positive off-diagonal entries and is strictly diagonally dominant by rows, with row excess $a_{ii} - \sum_{j \neq i} |a_{ij}| = 1 + \frac{\Delta\tau}{2}(r + \lambda) \geq 1$ in the interior rows. Consequently A is a nonsingular M-matrix and $A^{-1} \geq 0$ entrywise.

The proof is in Appendix A.8. Condition (25) says that the cell Péclet number $2|\mu|\Delta x/\sigma^2$ is at most two; it is independent of $\Delta\tau$ and, for the parameter ranges of this paper, far from binding. Three consequences follow. First, a nonsingular M-matrix is a P-matrix, so the discrete complementarity problem (37) below has exactly one solution for every right-hand side (Cottle et al., 1992). Second, the convergence theorem for projected successive over-relaxation of Cryer (1971) covers symmetric positive definite matrices, and A is nonsymmetric whenever $\mu \neq 0$. When (25) holds strictly, however, a positive diagonal scaling $\mathbf{V} = \Pi \mathbf{y}$ makes the interior block of A symmetric and strictly diagonally dominant, hence positive definite, and leaves both the complementarity problem and the projected SOR iterates unchanged (Appendix A.8); the theorem of Cryer (1971) then gives convergence for every relaxation parameter in $(0, 2)$. Third, the sign structure is the one under which the elimination algorithm of Brennan and Schwartz (1977) returns the exact solution of the put's complementarity problem when its contact set is an interval at the in-the-money end (Jaillet et al., 1990). I use the splitting of Ikonen and Toivanen (2004) as the default solver and the other two, which solve (37) itself, as cross-checks (Section 4.5).

4.2 The jump integral

The integral in (21) is approximated by replacing u on the grid by its piecewise-linear interpolant $I_h \mathbf{V}$, by the far-field form (17) outside the grid, and by integrating both exactly against f_Y . Let $\psi_k(y) = (1 - |y - k\Delta x|/\Delta x)^+$ denote the hat function centered at $k\Delta x$.

Lemma 4.2 (Hat-function weights for the double-exponential kernel). *Define*

$$\begin{aligned} w_k &= \int_{\mathbb{R}} \psi_k(y) f_Y(y) dy, \\ w_k^- &= \int_{(k-1)\Delta x}^{k\Delta x} \frac{y - (k-1)\Delta x}{\Delta x} f_Y(y) dy, \quad w_k^+ = \int_{k\Delta x}^{(k+1)\Delta x} \frac{(k+1)\Delta x - y}{\Delta x} f_Y(y) dy, \end{aligned} \quad (26)$$

so that $w_k = w_k^- + w_k^+$, $k \in \mathbb{Z}$. Then:

- (i) (Exact mass.) $\sum_{k \in \mathbb{Z}} w_k = 1$, and $\int_{\mathbb{R}} I_h \mathbf{V}(x_i + y) f_Y(y) dy = \sum_j w_{j-i} V_j$ for every grid function whose interpolant is defined on the whole line.

(ii) (Geometric structure.) With $\rho_1 = e^{-\eta_1 \Delta x}$ and $\rho_2 = e^{-\eta_2 \Delta x}$,

$$\begin{aligned} w_k &= A_1 \rho_1^k \quad (k \geq 1), & w_k &= A_2 \rho_2^{|k|} \quad (k \leq -1), \\ A_1 &= \frac{4p \sinh^2(\eta_1 \Delta x/2)}{\eta_1 \Delta x}, & A_2 &= \frac{4(1-p) \sinh^2(\eta_2 \Delta x/2)}{\eta_2 \Delta x}, \end{aligned} \quad (27)$$

$$w_0 = p \left(1 - \frac{1 - e^{-\eta_1 \Delta x}}{\eta_1 \Delta x}\right) + (1-p) \left(1 - \frac{1 - e^{-\eta_2 \Delta x}}{\eta_2 \Delta x}\right), \quad (28)$$

and the half weights $w_k^\pm$ are geometric with the same ratios on each side. In terms of the partial moments $G_0(a, b) = \int_a^b f_Y$ and $G_1(a, b) = \int_a^b y f_Y(y) dy$, with $y_k = k \Delta x$,

$$w_k^- = \frac{G_1(y_{k-1}, y_k) - y_{k-1} G_0(y_{k-1}, y_k)}{\Delta x}, \quad w_k^+ = \frac{y_{k+1} G_0(y_k, y_{k+1}) - G_1(y_k, y_{k+1})}{\Delta x}, \quad (29)$$

whose primitives are elementary (Appendix A.9); this form also covers a shifted density.

(iii) (Trapezoid weights.) The point-sampled weights $w_k^{\text{tr}} = \Delta x f_Y(k \Delta x)$ for $k \neq 0$ and $w_0^{\text{tr}} = \frac{\Delta x}{2} (f_Y(0^+) + f_Y(0^-))$ have total mass

$$\begin{aligned} \sum_k w_k^{\text{tr}} &= p \frac{\eta_1 \Delta x}{2} \coth \frac{\eta_1 \Delta x}{2} + (1-p) \frac{\eta_2 \Delta x}{2} \coth \frac{\eta_2 \Delta x}{2} \\ &= 1 + \frac{p(\eta_1 \Delta x)^2 + (1-p)(\eta_2 \Delta x)^2}{12} + O((\eta \Delta x)^4). \end{aligned} \quad (30)$$

The proof is in Appendix A.9. Part (i) holds because the hats form a partition of unity: constants are integrated exactly, whatever Δx , so the discrete jump operator conserves probability mass to rounding error and no resolution condition of the form $\eta \Delta x \ll 1$ is needed for the quadrature itself. Part (iii) shows what the condition would be with sampled weights: the excess mass ε in (30) enters the discrete operator as $\lambda \varepsilon u$, a spurious growth rate $\lambda \varepsilon$ per unit time, harmless when $\eta \Delta x$ is small and fatal when it is not (Appendix B.14 demonstrates both regimes, Figure 4). Part (ii) is the structure that makes the evaluation cheap: the weights are a pure exponential on each side of the origin because the kernel is, so the sums over $k \geq 1$ and $k \leq -1$ are geometric and can be accumulated by two first-order recursions, as Toivanen (2008) observed for his quadrature.

Edges and tails. On the grid, $I_h \mathbf{V}$ is defined on $[x_0, x_M]$ only; the hats of the two edge nodes are cut in half there, so the edge columns carry the half weights w_{-i}^+ (column $j = 0$) and w_{M-i}^- (column $j = M$). Outside the grid u is replaced by the far-field form (17), and the corresponding part of the integral is evaluated in closed form. For the put, with $a_i = x_0 - x_i \leq 0$,

$$\text{tail}_i^{\text{put}}(\tau) = \int_{-\infty}^{a_i} (A(\tau) - B(\tau) e^{x_i+y}) f_Y(y) dy = (1-p) e^{\eta_2 a_i} \left(A(\tau) - B(\tau) \frac{\eta_2}{\eta_2 + 1} e^{x_0} \right), \quad (31)$$

with $(A, B) = (K e^{-r\tau}, e^{-q\tau})$ for the European put; the upper tail of the put vanishes. The far field of the American put, $\max\{K e^{-r\tau} - e^{-q\tau}, K - e^x\}$, equals its payoff branch exactly for $x \leq x_\tau^* = \ln(K(1 - e^{-r\tau})/(1 - e^{-q\tau}))$, with $x_\tau^* = +\infty$ for $q = 0 < r$. The solver uses $(A, B) = (K, 1)$ when $x_0 < x_\tau^*$, which is always the case for $q \leq r$, $r > 0$, and is then exact on the whole lower tail; otherwise it uses the European pair. In the second case, possible only for $q > r$, the integrand is too small on $(-\infty, x_\tau^*)$ by at most $K(1 - e^{-r\tau})$, so $\text{tail}_i^{\text{put}}$ is too small by at most $K(1 - e^{-r\tau})(1-p)e^{\eta_2(x_\tau^* - x_i)}$; I do not correct this term. For the call the lower tail vanishes and, with $b_i = x_M - x_i \geq 0$, $\text{tail}_i^{\text{call}}(\tau) = p e^{-\eta_1 b_i} (B(\tau) \frac{\eta_1}{\eta_1 - 1} e^{x_M} - A(\tau))$, with the analogous switch for the American call (Appendix A.10).

Orientation. Collecting terms, the discrete jump operator is

$$(J\mathbf{V})_i = \lambda \sum_{j=0}^M \tilde{w}_{ij} V_j + \lambda \text{tail}_i(\tau), \quad \tilde{w}_{ij} = w_{j-i} \ (0 < j < M), \quad \tilde{w}_{i0} = w_{-i}^+, \quad \tilde{w}_{iM} = w_{M-i}^-. \quad (32)$$

The index is $j - i$, not $i - j$: $(J\mathbf{V})_i$ reads V_j at the lag $y = (j - i)\Delta x$, so an up-jump reads a higher node. A convolution, w_{i-j} , would implement jumps with the law of $-Y$ while the drift keeps the compensator of Y , so that the discounted price would no longer be a martingale. A one-line check of the orientation is $Je^x = \lambda(1 + \kappa)e^x$ up to the interpolation error of e^x (with the tails of the extension e^x added), whereas the reflected operator gives $\lambda\mathbb{E}[e^{-Y}]e^x$; Appendix B.10 reports both numbers.

Evaluation. The on-grid sum in (32) can be evaluated in three ways that give the same vector: as a dense matrix–vector product, $O(M^2)$ per step; by the two first-order recursions that the geometric weights (27) allow, $O(M)$ per step (equation (56) of Appendix B.4); or as a linear convolution with the reversed kernel by zero-padded FFT, $O(M \log M)$ per step (equation (57)). In the FFT form the reversal implements the correlation, the padding to at least $3M + 1$ points prevents the deep in-the-money values from wrapping around onto the out-of-the-money end (d’Halluin et al., 2005), and no factor Δx appears because the weights are integrals, not samples. Appendix B.10 reports the agreement of the three forms and Table 22 in Appendix B.20 their run times.

4.3 Time stepping: IMEX-CNAB with Rannacher start

Let $\tau_n = n\Delta\tau$ and $\mathbf{V}^n \approx u(\cdot, \tau_n)$. The scheme is Crank–Nicolson for $\mathcal{D}$ and second-order Adams–Bashforth for $\mathcal{J}$,

$$A\mathbf{V}^{n+1} = B\mathbf{V}^n + \Delta\tau \left(\frac{3}{2} J\mathbf{V}^n - \frac{1}{2} J\mathbf{V}^{n-1} \right), \quad n \geq 2, \quad (33)$$

the IMEX-CNAB scheme in the terminology of Ascher et al. (1995), applied to this PIDE by Salmi and Toivanen (2014). The tails in J are evaluated at the time level at which $\mathbf{V}$ is taken. Each step costs one evaluation of J (the previous one is reused) and one tridiagonal solve. The scheme is started with four IMEX-Euler half steps,

$$\left(I - \frac{\Delta\tau}{2} D \right) \mathbf{V}^{k+\frac{1}{2}} = \mathbf{V}^k + \frac{\Delta\tau}{2} J\mathbf{V}^k, \quad k = 0, \frac{1}{2}, 1, \frac{3}{2}, \quad (34)$$

which deliver $\mathbf{V}^1$ and $\mathbf{V}^2$ and the history $J\mathbf{V}^1$ needed by (33). This is the start-up of Rannacher (1984), analyzed for option pricing by Pooley et al. (2003) and Giles and Carter (2006): Crank–Nicolson is not strongly damping, so the high-frequency content excited by the kink of the payoff persists and shows up as an erratic temporal order and as oscillations in the gamma near the strike; the implicit-Euler half steps damp it. Section 5 reports the errors with and without the start-up, and Figure 5 in Appendix B.16 the gamma. The same start-up is applied after the announcement step, because the Adams–Bashforth history is not valid across the condition.

The stability of the explicit treatment of the integral is governed by the following scalar model, in which a stands for the symbol of the implicit part and θ for the normalized symbol of the jump operator; the continuous symbol is $\theta = \phi_Y(\xi)$, the discrete one is $\theta = \sum_k w_k e^{ik\xi\Delta x}$, and both satisfy $|\theta| \leq 1$ because the weights are nonnegative with unit sum.

Lemma 4.3 (Stability of CNAB for the scalar jump model). *Consider $v' = -(a + \lambda)v + \lambda\theta v$ with $a \geq 0$ real and $\theta \in \mathbb{C}$, $|\theta| \leq 1$, discretized by (33) with $D \mapsto -(a + \lambda)$ and $J \mapsto \lambda\theta$. The amplification factors are the roots of*

$$\left(1 + \frac{\alpha}{2} \right) z^2 - \left(1 - \frac{\alpha}{2} + \frac{3\beta}{2} \right) z + \frac{\beta}{2} = 0, \quad \alpha = (a + \lambda)\Delta\tau, \quad \beta = \lambda\theta\Delta\tau. \quad (35)$$

Both roots lie in the closed unit disc for all $a \geq 0$ and all $|\theta| \leq 1$ if and only if $\lambda\Delta\tau \leq 1$. For $\lambda\Delta\tau < 1$ they lie strictly inside the disc except for the neutral mode $a = 0$, $\theta = 1$, which has the root $z = 1$. For $a = 0$ and $\theta = -1$ the roots are real and of opposite sign, and the negative root z_- has modulus $|z_-| = (2\lambda\Delta\tau - 1 + \sqrt{5(\lambda\Delta\tau)^2 - 2\lambda\Delta\tau + 1})/(2 + \lambda\Delta\tau)$, which equals one at $\lambda\Delta\tau = 1$, exceeds one if and only if $\lambda\Delta\tau > 1$, and is the larger of the two root moduli if and only if $\lambda\Delta\tau \geq \frac{1}{2}$.

The proof (Appendix A.11) is the root-locus form of the Schur–Cohn criterion. Three comments. First, Salmi and Toivanen (2014) prove absolute stability of their IMEX-CNAB family for $\lambda\Delta\tau < \frac{1}{2}$ and every value of their split parameter, with the differential part of the scalar model set to zero. They also state, without a proof, that the member used here, with $-\lambda u$ wholly implicit, is stable for $\lambda\Delta\tau < 1$. Lemma 4.3 proves that bound for every $a \geq 0$ and every $|\theta| \leq 1$, including the endpoint $\lambda\Delta\tau = 1$, and shows that it is sharp. Second, the bound cannot be relaxed by keeping θ away from -1 . At $a = 0$ and $\lambda\Delta\tau = 1$, (35) factors as $(3z - 1)(z - \theta) = 0$, so the root $z = \theta$ lies on the unit circle whenever $|\theta| = 1$; and, for $a = 0$, $\lambda\Delta\tau > 1$ and $\theta = e^{i\omega}$, the root that equals one at $\theta = 1$ has $|z|^2 = 1 + \lambda\Delta\tau((\lambda\Delta\tau)^2 - 1)\omega^2 + O(\omega^3)$, so it leaves the disc for every $\theta \neq 1$ on the circle close to one. A particular kernel on a particular grid can do better only through $|\theta(\xi)| < 1$ for $\xi \neq 0$ together with the damping in the symbol $a(\xi)$ of the implicit part, and whether this carries the scheme past $\lambda\Delta\tau = 1$ depends on the parameters and the grid. Appendix B.10 reports the amplification computed with the discrete symbols of D and J (including the imaginary part that the drift adds to a , which the lemma does not cover) for the parameters scanned there, and Appendix B.14 a demonstration with $\lambda\Delta\tau$ far above one; I claim nothing beyond the parameters scanned. Third, in every production run of this paper $\lambda\Delta\tau$ is orders of magnitude below one; the time step is set by accuracy, not by stability.

4.4 The announcement step

The time levels must contain τ_E . I therefore run two legs on the same spatial grid, $[0, \tau_E]$ with N_A steps and $[\tau_E, T]$ with N_B steps, $N_A + N_B = N$ and $N_A/N_B \approx \tau_E/(T - \tau_E)$, each with its own $\Delta\tau$ and its own Rannacher start. At τ_E the condition (13) is applied once as the discrete operator

$$(\mathcal{P}_E^h \mathbf{V})_i = \sum_{j=0}^M \tilde{w}_{ij}^E V_j + \text{tail}_i^E(\tau_E), \quad w_k^E = \int_{\mathbb{R}} \psi_k(y) f_E(y + \kappa_E) dy, \quad (36)$$

built exactly as (32) but from the *shifted* density $f_E(\cdot + \kappa_E)$, so that the shift by κ_E is exact rather than snapped to the grid; the partial moments of the shifted density are those of f_E evaluated at $y + \kappa_E$ (Appendix A.9), and the tails use the far-field forms (17) integrated against the shifted tails of f_E (Appendix A.10). Two identities check the construction: $\mathcal{P}_E^h \mathbf{1} = \mathbf{1}$ to rounding error, by Lemma 4.2(i) applied to $f_E(\cdot + \kappa_E)$, and $\mathcal{P}_E^h e^x = e^x$ up to the interpolation error of e^x on the grid, which is second order in Δx , by (6). The second identity also checks the orientation, through the compensator: the reflected (convolution) operator built from the same shifted weights implements $x \mapsto x - Z + \kappa_E$ and returns $e^x \mathbb{E}[e^Z] \mathbb{E}[e^{-Z}] = e^{x+\kappa_E} \mathbb{E}[e^{-Z}]$, a factor of about $1 + \text{Var } Z$ that does not vanish as $\Delta x \rightarrow 0$, for a symmetric law as for an asymmetric one. Since $\ln(\mathbb{E}[e^Z] \mathbb{E}[e^{-Z}])$ involves only the even cumulants of Z , this identity cannot see the asymmetry of the law; the skew of Figure 6 checks that. The operator is applied once, so I evaluate it as a dense matrix–vector product; the geometric structure of Lemma 4.2(ii) holds on each side of the cell containing $-\kappa_E$, so the recursion (56) applies equally with the two cells adjacent to the kink treated separately. For the American option the projection $\mathbf{V} \leftarrow \max(\mathcal{P}_E^h \mathbf{V}, \boldsymbol{\varphi})$ follows, the multiplier of the splitting below is reset to zero, and the second leg restarts with (34). By Theorem 3.3 the projection is inactive for vanilla payoffs up to the interpolation error of $\mathcal{P}_E^h$; the solver records $\min_i ((\mathcal{P}_E^h \mathbf{V})_i - \varphi_i)$ and Section 5 reports it.

4.5 The American constraint

With $\mathbf{b}^n$ the right-hand side of (33) (or of (34), with $\Delta\tau/2$ in place of $\Delta\tau$), the time step of the American option is the discrete complementarity problem

$$A\mathbf{V}^{n+1} \geq \mathbf{b}^n, \quad \mathbf{V}^{n+1} \geq \boldsymbol{\varphi}, \quad (A\mathbf{V}^{n+1} - \mathbf{b}^n)^\top (\mathbf{V}^{n+1} - \boldsymbol{\varphi}) = 0, \quad (37)$$

with the Dirichlet rows fixed at $\max\{\text{far-field value}, \boldsymbol{\varphi}\}$. By Proposition 4.1 it has a unique solution. The Brennan–Schwartz elimination and projected SOR below compute that solution, the first exactly under the interval condition stated with it and the second to its stopping tolerance. The default, the operator splitting, solves a nearby problem instead, in which the multiplier lags by one step.

Operator splitting (Ikonen and Toivanen, 2004, 2009). With a multiplier $\boldsymbol{\mu}^n \geq 0$ that approximates $(A\mathbf{V}^{n+1} - \mathbf{b}^n)/\Delta\tau$,

$$\begin{aligned} A\tilde{\mathbf{V}}^{n+1} &= \mathbf{b}^n + \Delta\tau \boldsymbol{\mu}^n, & \mathbf{V}^{n+1} &= \max(\tilde{\mathbf{V}}^{n+1} - \Delta\tau \boldsymbol{\mu}^n, \boldsymbol{\varphi}), \\ \boldsymbol{\mu}^{n+1} &= \boldsymbol{\mu}^n + \frac{\mathbf{V}^{n+1} - \tilde{\mathbf{V}}^{n+1}}{\Delta\tau} \geq 0, \end{aligned} \quad (38)$$

one tridiagonal solve and one projection per step, no iteration; the multiplier is nonnegative by construction, and $\boldsymbol{\mu}^0 = \mathbf{0}$. This is the default. The pair $(\mathbf{V}^{n+1}, \boldsymbol{\mu}^{n+1})$ satisfies $\mathbf{V}^{n+1} \geq \boldsymbol{\varphi}$, $\boldsymbol{\mu}^{n+1} \geq 0$ and $(\mathbf{V}^{n+1} - \boldsymbol{\varphi})^\top \boldsymbol{\mu}^{n+1} = 0$ exactly, but $A\mathbf{V}^{n+1} - \mathbf{b}^n = \Delta\tau \boldsymbol{\mu}^{n+1} + \Delta\tau(A - I)(\boldsymbol{\mu}^{n+1} - \boldsymbol{\mu}^n)$. The splitting therefore returns the solution of (37) with the right-hand side $\mathbf{b}^n$ replaced by $\mathbf{b}^n - \frac{1}{2}\Delta\tau^2 D(\boldsymbol{\mu}^{n+1} - \boldsymbol{\mu}^n)$, and differs from the solution of (37) by a splitting error that vanishes as $\Delta\tau \rightarrow 0$ on a fixed grid; Panel C of Table 2 measures it against the other two solvers.

Brennan–Schwartz elimination (Brennan and Schwartz, 1977). For the put, factor A by Gaussian elimination starting from the out-of-the-money end $i = M$, then back-substitute from the in-the-money end,

$$V_i = \max\left\{\varphi_i, \frac{\hat{b}_i - \hat{a}_{i,i-1}V_{i-1}}{\hat{a}_{ii}}\right\}, \quad i = 1, 2, \dots, M-1, \quad (39)$$

where hats denote the eliminated system. Jaillet et al. (1990) prove that, for a tridiagonal M-matrix and a contact set that is an interval at the in-the-money end, this returns the exact solution of (37). Lemma 3.2 and Corollary 3.4 give that structure for the exercise region of the continuous problem, not for the contact set of (37), whose right-hand side carries the explicit jump term; I do not prove the discrete analogue. I assume it, and check it a posteriori by the agreement of the elimination with projected SOR, which does not need it (Table 2). The call is handled by mirror image.

Projected SOR (Cryer, 1971). With relaxation parameter $\omega_{\text{SOR}} \in (0, 2)$,

$$V_i^{(k+1)} = \max\left\{\varphi_i, V_i^{(k)} + \omega_{\text{SOR}}(\hat{y}_i - V_i^{(k)})\right\}, \quad \hat{y}_i = \frac{b_i - a_{i,i-1}V_{i-1}^{(k+1)} - a_{i,i+1}V_{i+1}^{(k)}}{a_{ii}}, \quad (40)$$

stopped when $\|\mathbf{V}^{(k+1)} - \mathbf{V}^{(k)}\|_\infty \leq \varepsilon \max\{1, \|\mathbf{V}^{(k+1)}\|_\infty\}$. The test is on the change of the iterate; a test on the amount by which the projection moved the iterate is zero whenever no node is projected and positive wherever exercise is optimal, so it stops too early in the first case and never in the second. PSOR converges here for every $\omega_{\text{SOR}} \in (0, 2)$ by the symmetrization argument after Proposition 4.1; I use it as the cross-check of the two non-iterative methods (Table 2).

For the European option the step is the tridiagonal solve $A\mathbf{V}^{n+1} = \mathbf{b}^n$ alone. In all cases the jump integral of the American solution uses the American far field in (31).

Algorithm 1: IMEX-CNAB solver for the Kou-SEJ option (American put shown; European and call analogous)

- 1: **Input:** grid (22) with $\ln K$ on a node; N_A, N_B with $\Delta\tau_A = \tau_E/N_A, \Delta\tau_B = (T - \tau_E)/N_B$; parameters Θ, Θ_E ; payoff φ .
- 2: Assemble $A_\ell = I - \frac{\Delta\tau_\ell}{2}D, B_\ell = I + \frac{\Delta\tau_\ell}{2}D$ for $\ell \in \{A, B\}$ (23)–(24); the weights $w_k, w_k^\pm$ and tail coefficients (Lemma 4.2, (31)); the announcement weights w_k^E and tails (36).
- 3: $\mathbf{V}^0 \leftarrow \varphi; \boldsymbol{\mu} \leftarrow \mathbf{0}; \tau \leftarrow 0$.
- 4: **for** leg $\ell = A$ (up to τ_E) and then $\ell = B$ (up to T):
- 5: Rannacher start: four IMEX-Euler half steps (34), each with the Dirichlet rows from (17) and the step (38) with $\Delta\tau_\ell/2$; keep $\mathbf{V}^{n-1}, \mathbf{V}^n, J\mathbf{V}^{n-1}$.
- 6: **for** the remaining steps of the leg:
- 7: $\mathbf{b} \leftarrow B_\ell \mathbf{V}^n + \Delta\tau_\ell(\frac{3}{2}J\mathbf{V}^n - \frac{1}{2}J\mathbf{V}^{n-1})$ with J by (56) or (57) and tails (31); Dirichlet rows at $\tau + \Delta\tau_\ell$.
- 8: $(\mathbf{V}^{n+1}, \boldsymbol{\mu}) \leftarrow$ the step (38) [alternatives: (39), (40)]; European: solve $A_\ell \mathbf{V}^{n+1} = \mathbf{b}$.
- 9: $\tau \leftarrow \tau + \Delta\tau_\ell$.
- 10: **if** $\ell = A$ and $\tau_E \in (0, T)$: $\mathbf{V} \leftarrow \mathcal{P}_E^h \mathbf{V}$ (36); record $\min_i (V_i - \varphi_i)$; $\mathbf{V} \leftarrow \max(\mathbf{V}, \varphi)$; $\boldsymbol{\mu} \leftarrow \mathbf{0}$.
- 11: **Output:** $u(\ln S_0, T)$, the node value if $\ln S_0$ is a node and otherwise the cubic spline in x through $\mathbf{V}$; delta and gamma by central differences in x ; the exercise boundary at each level from $\{i : V_i = \varphi_i\}$.

4.6 Complexity and algorithm

Per time step the cost is one tridiagonal solve, $O(M)$, plus one jump evaluation, $O(M)$ with (56), $O(M \log M)$ with (57) or $O(M^2)$ with the dense matrix, plus $O(M)$ for the splitting or the elimination and $O(M \times \text{sweeps})$ for PSOR. The announcement step costs $O(M^2)$ once in the dense form. The whole solve is therefore $O(NM)$ operations and $O(M)$ memory with the recursion, against $O(NM^2)$ for the dense evaluation; which one is faster at a given M is a question of constants that Table 22 in Appendix B.20 answers. Algorithm 1 summarizes the method.

Appendix B.15 lists the implementation errors that give a scheme that runs and returns plausible-looking numbers (Remark B.1), with the diagnostics that expose them.

5 Verification

The solver of Section 4 is tested against benchmarks that are exact or published, each of which isolates one part of the scheme, and against consistency checks that the model supplies at no cost. Every number quoted in this section is written by the verification scripts together with the tables and figures (Appendix B.22), except those of Section 5.6, which come from the grid study of the pricing script. I report the convergence orders as observed, including the places where they are not two. Appendix B gives the implementation, the commands and the secondary checks with their tables and figures, and Section 5.7 lists what the suite does not test. The fixed-step temporal ladders, the binomial reference of Table 2, the American calls, the Greeks with the announcement and the comparison with the published prices with an announcement were added after the specification freeze of Section 8; they test the frozen code and changed nothing in it.

5.1 Design of the suite

What each benchmark isolates. With $\lambda = 0$ and no announcement the PIDE (12) is the Black–Scholes equation, so the closed form of Black and Scholes (1973) tests the differential operator D of (23), the Crank–Nicolson step with its Rannacher start, the Dirichlet data (17) and the log grid, and nothing else; with the exercise constraint added, a converged binomial

reference, the binomial Black–Scholes method with Richardson extrapolation of [Broadie and Detemple \(1996\)](#), tests the three LCP solvers of Section 4.5 without jumps. With $\lambda > 0$ the cosine expansion of [Fang and Oosterlee \(2008\)](#) applied to the Kou characteristic function (9) is an independent European price, and because D is already verified the comparison isolates the jump operator J : the hat weights, the tails, the orientation of the correlation and the explicit Adams–Bashforth treatment. The American put reference values of the same test case, as [Salmi and Toivanen \(2011\)](#) list them, then test the complementarity problem with jumps. With an announcement inside the option’s life, the product characteristic function of Proposition 2.3(i) gives a European price by the same cosine method, and the mixture formula of Proposition 2.3(iii) gives the same price by adaptive quadrature of Black–Scholes values, a computation that shares no code with the cosine expansion; this comparison isolates the announcement step (36) together with the restart and the two-leg bookkeeping of Section 4.4. For the American option with an announcement there is no closed form. The one published benchmark, the Fourier-transform prices of [Leung and Santoli \(2014\)](#) for Kou dynamics with a double-exponential announcement jump, is printed to two decimals, and Section 5.4 compares the solver with it at that resolution. Beyond it I rely on self-convergence, on the agreement of the three LCP solvers, and on the properties proved in Section 3: the continuation value at the announcement never falls below the payoff (Theorem 3.3), the European price does not depend on t_E (Proposition 2.3(v)), the American price is nonincreasing in t_E (Proposition 3.5), and the put’s exercise boundary falls to zero as t increases to t_E (Corollary 3.4), so that on a finite grid the exercise region empties before the announcement.

Conventions. A convergence ladder refines M and N together with $N = M/2$, so the observed order is a joint order in $(\Delta x, \Delta \tau)$. How much of the error is temporal is measured separately, by refining time alone at fixed Δx (Section 5.5). At this ratio the error is almost entirely spatial for the European Kou put, whose temporal part is 0.1 percent of the error on the (800, 400) grid, and for the Kou American put with the splitting (1.1 percent there), but not with the announcement: on the (1600, 800) grid the temporal part is 10 percent of the American error and 22 percent of the error of the early-exercise premium. Where an exact reference exists the error is $|V_{\text{FD}}(S_0) - V_{\text{ref}}(S_0)|$ at a grid node and the order is $\log_2$ of the ratio of successive errors; where none exists I report the difference Δ between successive grid values and the ratio of successive differences, which is four for second order. Richardson extrapolation uses the two finest grids and assumes order two in Δx and $\Delta \tau$ jointly, an assumption that the American ladders satisfy in space but not in time (Section 5.5). Prices are in currency units with $K = 100$.

The Black–Scholes limit. On four grids from $(M, N) = (100, 50)$ to $(800, 400)$ the at-the-money European put with $\sigma = 0.3$, $r = 0.05$ and $T = 0.5$ converges at observed order 2.00 at every level, its error falling from 1.50×10^{-2} to 2.34×10^{-4} , and Richardson extrapolation leaves a residual of 4.0×10^{-8} against the closed form, which is what a pure second-order error leaves behind; the call behaves the same way (Table 15 in Appendix B.9).

5.2 The jump operator: the Kou European put against COS

The parameters of Table 1 are those of the test case of [Toivanen \(2008\)](#): $\sigma = 0.15$, $\lambda = 0.1$, $p = 0.3445$, $\eta_1 = 3.0465$, $\eta_2 = 3.0775$, $r = 0.05$, $K = 100$, $T = 0.25$, so $\lambda T = 0.025$ and $\kappa = 0.007575914$. The reference values use $N = 2^{14}$ cosine terms and the truncation width $L_c = 20$; raising N to 2^{16} and L_c to 30 changes them by at most 3×10^{-11} . A narrower interval is not enough for this kernel: at $L_c = 12$ the cumulant interval cuts off the rare large jumps (the mean jump sizes $1/\eta_1$ and $1/\eta_2$ are about a third) and the value is off by up to 5.9×10^{-7} . The default setting $N = 2^{12}$, $L_c = 10$ differs from the reference by 7.0×10^{-6} , 5.1×10^{-6} and 3.8×10^{-6} at $S_0 = 90$, 100 and 110, which is the error of the default setting. The European

Table 1: Kou European put against the COS benchmark, Toivanen parameters ($\sigma = 0.15$, $\lambda = 0.1$, $p = 0.3445$, $\eta_1 = 3.0465$, $\eta_2 = 3.0775$, $r = 0.05$, $q = 0$, $K = 100$, $T = 0.25$).

M	N	Δx	$\lambda\Delta\tau$	$V_{\text{FD}}(S_0)$	error	order
$S_0 = 90$, COS reference ($N = 2^{14}$, $L_c = 20$) = 9.43045738						
100	50	0.03512	5.0×10^{-4}	9.42551036	4.95×10^{-3}	–
200	100	0.01756	2.5×10^{-4}	9.42884741	1.61×10^{-3}	1.62
400	200	0.00878	1.3×10^{-4}	9.43003141	4.26×10^{-4}	1.92
800	400	0.00439	6.3×10^{-5}	9.43034939	1.08×10^{-4}	1.98
1600	800	0.00220	3.1×10^{-5}	9.43043024	2.71×10^{-5}	1.99
3200	1600	0.00110	1.6×10^{-5}	9.43045054	6.84×10^{-6}	1.99
6400	3200	0.00055	7.8×10^{-6}	9.43045561	1.77×10^{-6}	1.96
$S_0 = 100$, COS reference ($N = 2^{14}$, $L_c = 20$) = 2.73125890						
100	50	0.03526	5.0×10^{-4}	2.64760107	8.37×10^{-2}	–
200	100	0.01763	2.5×10^{-4}	2.71112249	2.01×10^{-2}	2.05
400	200	0.00881	1.3×10^{-4}	2.72626547	4.99×10^{-3}	2.01
800	400	0.00441	6.3×10^{-5}	2.73001295	1.25×10^{-3}	2.00
1600	800	0.00220	3.1×10^{-5}	2.73094751	3.11×10^{-4}	2.00
3200	1600	0.00110	1.6×10^{-5}	2.73118101	7.79×10^{-5}	2.00
6400	3200	0.00055	7.8×10^{-6}	2.73123938	1.95×10^{-5}	2.00
$S_0 = 110$, COS reference ($N = 2^{14}$, $L_c = 20$) = 0.55236304						
100	50	0.03177	5.0×10^{-4}	0.53875872	1.36×10^{-2}	–
200	100	0.01906	2.5×10^{-4}	0.54716871	5.19×10^{-3}	1.88
400	200	0.00866	1.3×10^{-4}	0.55126380	1.10×10^{-3}	1.97
800	400	0.00433	6.3×10^{-5}	0.55208693	2.76×10^{-4}	1.99
1600	800	0.00222	3.1×10^{-5}	0.55229064	7.24×10^{-5}	2.00
3200	1600	0.00111	1.6×10^{-5}	0.55234488	1.82×10^{-5}	2.00
6400	3200	0.00055	7.8×10^{-6}	0.55235849	4.55×10^{-6}	1.98

Hat-function (P1) quadrature weights, grid centered at $\ln S_0$ with Δx snapped so that $\ln K$ is a node; dense jump operator for $M < 3200$ and the $O(M)$ recursion above. The COS reference changes by at most 2.9×10^{-11} when N is raised to 2^{16} and L_c to 30; the European values that [Salmi and Toivanen \(2011\)](#) list in their Table 1 for this test case of [Toivanen \(2008\)](#) (9.430457, 2.731259, 0.552363) agree with it to within 3.8×10^{-7} , the rounding of their six decimals. The operator self-checks, the error split and the domain check are reported in Section 5.2 and Appendix B.10.

values that [Salmi and Toivanen \(2011\)](#) list for this test case agree with the reference to within 4×10^{-7} , the rounding of their six decimals. Calls are priced through put–call parity.

Convergence. At $S_0 = 100$ the error falls from 8.37×10^{-2} on the (100, 50) grid to 1.95×10^{-5} on the (6400, 3200) grid with observed orders 2.05, 2.01, 2.00, 2.00, 2.00 and 2.00 (Figure 1). At $S_0 = 110$ and $S_0 = 90$ the non-nested coarse grids distort the start of the ladders; from the second refinement on the orders lie between 1.9 and 2.0 (Appendix B.10). The explicit treatment of the integral is far from its stability limit throughout: $\lambda\Delta\tau$ runs from 5.0×10^{-4} to 7.8×10^{-6} .

Where the error comes from. On the (800, 400) grid the error at $S_0 = 100$ is 1.25×10^{-3} . Refining time alone to (800, 3200) leaves it at 1.24×10^{-3} ; refining space alone to (6400, 400) reduces it to 2.13×10^{-5} . The error of this European ladder is therefore spatial, which is why its joint order coincides with the spatial one; its temporal order, measured on its own in Section 5.5, is two. Doubling the half-width at fixed Δx changes $V(100)$ by 6.1×10^{-8} , so the localization of Remark 3.1 is not a visible source of error at $L \approx 1.76$. The Rannacher start matters for the derivatives (Appendix B.16) and not for this price: with 0, 2 and 4 implicit-Euler half steps the (400, 200) error is 4.984×10^{-3} , 4.989×10^{-3} and 4.993×10^{-3} .




Figure 1: Convergence of the price at $S = K = 100$ along the ladders $N = M/2$ (both axes logarithmic; the dotted line has slope -2). Every series is an error against an independent reference: the Black-Scholes closed form (Table 15 in Appendix B.9), the COS value of the Kou put at $S_0 = 100$ (Table 1), the product-characteristic-function COS value with the announcement at $t_E = T/2$ (Table 20 in Appendix B.17), the binomial reference of case 1 in Panel A of Table 2 for the Black-Scholes American put, and the time-converged exact-LCP value of Section 5.5 for the Kou American put at $S_0 = 100$. All five series run parallel to the reference slope over the grids shown; the constants differ between problems, the slopes do not.

Operator identities and stability. The hat weights conserve mass to 3.4×10^{-15} on the (400, 200) grid, as Lemma 4.2(i) says (the trapezoid weights to 2.2×10^{-4}); the dense, FFT and recursive evaluations of J agree to 1.2×10^{-17} and 8.7×10^{-14} ; the orientation check of Section 4.2 would expose a reflected kernel; and with the discrete symbols of D and J on the (400, 200) grid the largest amplification modulus stays below one up to $\lambda\Delta\tau = 8$ (0.9999 there), far beyond the generic bound of Lemma 4.3 (Appendix B.10).

5.3 The complementarity problem: American puts

Without jumps (Panel A). For the at-the-money American put with $\sigma = 0.3$, $T = 0.5$ (case 1) and $\sigma = 0.2$, $T = 1$ (case 2), the reference is the binomial Black-Scholes price with Richardson extrapolation at 40 000 steps, 7.3940416 and 6.0903712; the Richardson value of the even-step binomial tree agrees with it to 5×10^{-7} , and the reference changes by at most 1×10^{-6} from 20 000 to 40 000 steps. Against it the splitting converges in case 1 with observed orders 1.97, 1.98, 1.99 and 2.04 (ratios of successive differences 3.89, 3.95 and 3.93), that is, at order two. In case 2 the orders rise toward the fine end, 2.06, 2.10, 2.26 and 2.66 (ratios 4.13, 4.17 and 4.50), because at $N = M/2$ the temporal error of the splitting is positive and cancels a growing part of the negative spatial error. Two runs on the same grids show this: with eight times the time steps the splitting's ratios are 3.96, 3.96 and 3.97 and its Richardson value lies -7×10^{-8} from the reference, while the exact discrete LCP (Brennan-Schwartz) at $N = M/2$ errs in the other direction, with ratios 3.72, 3.63 and 3.35. The Richardson values of the $N = M/2$ ladders, 7.394044 and 6.090393, lie $+3 \times 10^{-6}$ and $+2 \times 10^{-5}$ from the reference. On the (200, 100) grid Brennan-Schwartz and PSOR agree to 2.8×10^{-11} and 3.8×10^{-11} , the stopping tolerance, while

Table 2: American puts: $\lambda = 0$ against a converged binomial reference (Panel A), the Kou test case of [Toivanen \(2008\)](#) at $S_0 = 90, 100, 110$ (Panel B), and the Ikonen–Toivanen splitting against the exact discrete LCP (Panel C).

<i>Panel A: $\lambda = 0$ American put, $S_0 = K = 100$, $r = 0.05$; case 1: $\sigma = 0.3$, $T = 0.5$; case 2: $\sigma = 0.2$, $T = 1$</i>										
M	N	V_{IT} (1)	$V - V_{\text{ref}}$	ratio	order	V_{IT} (2)	$V - V_{\text{ref}}$	ratio	order	
100	50	7.376343	-1.77×10^{-2}	–	–	6.075032	-1.53×10^{-2}	–	–	
200	100	7.389512	-4.53×10^{-3}	–	1.97	6.086692	-3.68×10^{-3}	–	2.06	
400	200	7.392896	-1.15×10^{-3}	3.89	1.98	6.089515	-8.56×10^{-4}	4.13	2.10	
800	400	7.393753	-2.88×10^{-4}	3.95	1.99	6.090192	-1.79×10^{-4}	4.17	2.26	
1600	800	7.393971	-7.01×10^{-5}	3.93	2.04	6.090343	-2.84×10^{-5}	4.50	2.66	
V_{ref} (BBSR, $n = 40000$): 7.3940416 / 6.0903712; Richardson of the two finest grids 7.394044 / 6.090393 ($+2.5 \times 10^{-6}$ / $+2.2 \times 10^{-5}$ from V_{ref}) Ratios with Brennan–Schwartz: 3.73, 3.55, 3.42 / 3.72, 3.63, 3.35; splitting with $N = 4M$: 3.94, 3.94, 3.96 / 3.96, 3.96, 3.97, Richardson -8.5×10^{-7} / -7.0×10^{-8} from V_{ref}										
<i>Panel B: Kou American put, Toivanen parameters, $K = 100$, $T = 0.25$ (grid centered at $\ln S_0$)</i>										
		$S_0 = 90$			$S_0 = 100$			$S_0 = 110$		
M	N	V_{IT}	Δ	ratio	V_{IT}	Δ	ratio	V_{IT}	Δ	ratio
100	50	10.00000000	–	–	2.71663364	–	–	0.54775217	–	–
200	100	10.00000000	0	–	2.78521354	6.86×10^{-2}	–	0.55635068	8.60×10^{-3}	–
400	200	10.00273248	2.73×10^{-3}	–	2.80212595	1.69×10^{-2}	4.06	0.56065213	4.30×10^{-3}	2.00
800	400	10.00488567	2.15×10^{-3}	1.27	2.80641644	4.29×10^{-3}	3.94	0.56155562	9.03×10^{-4}	4.76
1600	800	10.00496520	7.95×10^{-5}	27.08	2.80750873	1.09×10^{-3}	3.93	0.56178912	2.34×10^{-4}	3.87
3200	1600	10.00505782	9.26×10^{-5}	0.86	2.80778614	2.77×10^{-4}	3.94	0.56185377	6.47×10^{-5}	3.61
6400	3200	10.00506787	1.00×10^{-5}	9.22	2.80785596	6.98×10^{-5}	3.97	0.56187065	1.69×10^{-5}	3.83
12800	6400	10.00507090	3.04×10^{-6}	3.31						
Richardson (error estimate) at $S_0 = 90, 100, 110$: 10.005072 (1.0×10^{-6}), 2.807879 (2.3×10^{-5}), 0.561876 (5.6×10^{-6}); literature values 10.005071, 2.807879, 0.561876; Richardson – literature $+9.2 \times 10^{-7}$, $+2.3 \times 10^{-7}$, $+2.7 \times 10^{-7}$ Brennan–Schwartz on the same grids: ratios at $S_0 = 100$ 4.05, 3.93, 3.89, 3.88, 3.85; Richardson 10.0050693, 2.8078765, 0.5618756, minus literature -1.7×10^{-6} , -2.5×10^{-6} , -3.9×10^{-7} Time-converged exact-LCP reference (Brennan–Schwartz with $N = 25600$ on $M = 3200$ and 6400, Richardson in space): $S_0 = 100$: 2.8078787 (-3.1×10^{-7} from the literature), $S_0 = 110$: 0.5618765 ($+5.2 \times 10^{-7}$ from the literature)										
<i>Panel C: Ikonen–Toivanen splitting (IT) against the exact discrete LCP (Brennan–Schwartz, BS) and PSOR</i>										
M	N	IT–BS (90)	IT–BS (100)	IT–BS (110)	$\max_i \text{IT} - \text{BS} $	factor	$\max_i \text{BS} - \text{PSOR} $	sweeps		
200	100	0	$+6.62 \times 10^{-5}$	$+8.14 \times 10^{-6}$	1.18×10^{-4}	–	4.5×10^{-10}	16.3/19		
400	200	$+3.71 \times 10^{-5}$	$+4.36 \times 10^{-5}$	$+8.60 \times 10^{-6}$	6.01×10^{-5}	1.96	7.6×10^{-10}	14.3/18		
800	400	$+2.90 \times 10^{-5}$	$+2.93 \times 10^{-5}$	$+6.42 \times 10^{-6}$	3.57×10^{-5}	1.68	1.1×10^{-9}	11.4/18		
1600	800	$+1.40 \times 10^{-5}$	$+1.61 \times 10^{-5}$	$+3.63 \times 10^{-6}$	1.90×10^{-5}	1.88				
3200	1600	$+6.51 \times 10^{-6}$	$+8.41 \times 10^{-6}$	$+1.95 \times 10^{-6}$	9.79×10^{-6}	1.94				
6400	3200	$+3.09 \times 10^{-6}$	$+4.12 \times 10^{-6}$	$+9.81 \times 10^{-7}$	4.77×10^{-6}	2.05				

Δ is the difference between successive grid values and “ratio” the ratio of successive Δ (4 for second order); “order” is $\log_2$ of the ratio of successive $|V - V_{ref}|$. Richardson uses the two finest grids with order two. V_{ref} in Panel A is the binomial Black–Scholes price with Richardson extrapolation (BBSR) of [Broadie and Detemple \(1996\)](#). The literature values of Panel B are the reference prices that [Salmi and Toivanen \(2011\)](#) list in their Table 1 for this test case of [Toivanen \(2008\)](#); Toivanen’s own table was not consulted. On the (6400, 3200) grid the exercise boundary at $\tau = T$ is $S^*(T) = 89.517$ ($\Delta x = 0.00055$, i.e. ± 0.05 in S); $S_0 = 90$ lies $+0.48$ above it, which is why its convergence is irregular. In Panel C, IT–BS is taken at S_0 on the grid of Panel B and $\max_i$ over the grid centered at $\ln K$, “factor” is the ratio of successive maxima, and PSOR ($\omega_{SOR} = 1.3$, max-norm tolerance 10^{-10}) runs on the three coarsest grids, “sweeps” being its mean/maximum number of sweeps per time step. PSOR and BS solve the LCP of each step to the stopping tolerance (their difference is the PSOR stopping error), whereas IT carries a splitting error that vanishes as $\Delta\tau \rightarrow 0$ on a fixed grid (Section 4.5).




Figure 2: Exercise boundaries of American puts. (a) $S^*(\tau)$ for the Kou put with the parameters of Table 1 ($\lambda = 0.1$) and for the Black-Scholes put with the same $\sigma = 0.15$, $K = 100$, $T = 0.25$, on the (6400, 3200) grid; at $\tau = T$ the Kou boundary is 89.517, which is why $S_0 = 90$ converges irregularly in Panel B of Table 2. (b) The boundary in calendar time for the put of Table 3 ($\sigma = 0.25$, $r = 0.04$, $T = 9/252$) with and without the announcement at $t_E = 4.5$ trading days, grid (1600, 800). With the announcement the boundary today is 69.945 against 90.671 without; as $t \uparrow t_E$ it falls to the grid floor 45.68 and the exercise region empties, the finite-grid form of Corollary 3.4; at t_E^+ it rejoins the no-event boundary at 92.918.

the splitting differs from them by at most 9.71×10^{-4} and 1.26×10^{-3} , its own $O(\Delta\tau)$ error.

With jumps (Panel B). With the Toivanen parameters the at-the-money ladder has ratios 4.06, 3.94, 3.93, 3.94 and 3.97 from (200, 100) to (6400, 3200); the Richardson limit 2.807879 differs from the finest-grid value by 2.3×10^{-5} and from the published value 2.807879 by $+2.3 \times 10^{-7}$. That agreement is closer than the extrapolation can resolve. Richardson extrapolation at order two removes neither the first-order splitting error, which leaves the splitting's Richardson value $+2.7 \times 10^{-6}$ above that of the exact discrete LCP on the same grids, nor the time error of the exact projected Crank-Nicolson step, whose observed temporal order is 1.2 to 1.5 (Section 5.5) and which leaves the exact-LCP Richardson value -2.5×10^{-6} from the published one. The two residuals have opposite signs and nearly cancel, so the agreement should be read as agreement within a few 10^{-6} . A time-converged exact-LCP reference, Brennan-Schwartz with $N = 25600$ on $M = 3200$ and 6400 extrapolated in space, is 2.8078787, -3×10^{-7} from the published value. At $S_0 = 110$ the ratios are 2.00, 4.76, 3.87, 3.61 and 3.83, the first two reflecting the non-nested coarse grids, and the limit 0.561876 differs from the published value by $+2.7 \times 10^{-7}$ with a finest-grid error estimate of 5.6×10^{-6} ; the time-converged reference there is 0.5618765 ($+5 \times 10^{-7}$ from the published value). The spot $S_0 = 90$ lies 0.48 above the boundary 89.517; on the two coarsest grids it sits inside the discrete exercise region, and from (400, 200) on the ratios are irregular because the distance between S_0 and the nearest node of the discrete boundary changes from grid to grid. The Richardson value there, 10.005072 against the published 10.005071, differs from it by an amount of the same size as elsewhere, but the agreement rests on the finest grid rather than on the extrapolation (Appendix B.11). The three published values are the reference prices that Salmi and Toivanen (2011) list in their Table 1 for this test case of Toivanen (2008) (see the note to Table 2). Figure 2(a) shows the exercise boundary that the solution carries: the Kou boundary lies below the Black-Scholes boundary of the same σ at every τ .

Three solvers (Panel C). Projected SOR and the Brennan–Schwartz elimination solve the discrete complementarity problem exactly, up to the stopping tolerance, and so measure the error of the production splitting of Ikonen and Toivanen (2004). On the grids of the ladders from (200, 100) to (6400, 3200) the largest difference over the grid between the splitting and Brennan–Schwartz solutions falls from 1.2×10^{-4} to 4.8×10^{-6} , by factors of 1.96, 1.68, 1.88, 1.94 and 2.05 when the step is halved: first order along this ladder, on which $\Delta\tau$ is proportional to Δx , and smaller than the change between successive grids on every grid of the ladder; at $S_0 = 100$ the difference is $+6.6 \times 10^{-5}$ on the coarsest of these grids and $+4.1 \times 10^{-6}$ on the finest. Brennan–Schwartz and PSOR agree to 4.5×10^{-10} , 7.6×10^{-10} and 1.1×10^{-9} on the three coarsest grids, which confirms the interval structure of the contact set that the elimination assumes (Section 4.5). PSOR with $\omega_{\text{SOR}} = 1.3$ needs 16.3, 14.3 and 11.4 sweeps per time step on average with the warm start from the previous level, at most 19.

American calls with a dividend yield. The American model of the empirical study prices calls with an ex-dividend date in their life, and the puts above do not touch the code that such calls use: the upper far field, the $e^{-q\tau}$ factors in the boundary data and the tails, and the exercise of calls. Table 17 (Appendix B.13) checks them in two ways. Without jumps, the American call with $q > 0$ converges to the binomial reference at observed orders 1.84 to 2.00 in three cases, two of them at maturities of the empirical sample, with an error of at most 8×10^{-5} on the (1600, 800) grid, and its early-exercise premium (0.021 to 0.180) is reproduced to 2×10^{-5} . And through put–call symmetry, $C(S, K, r, q) = P(K, S, q, r)$, which holds with an announcement as well if the put carries the law of $-Z$ under the share measure: the cosine prices of the two problems agree to 5×10^{-12} , and for an asymmetric announcement law three trading days ahead and $q = 0.08$ the finite-difference American call and the American put of the symmetric problem, each solved on its own grid, approach each other from 3×10^{-4} on the (400, 200) grid to 2×10^{-5} on the (1600, 800) grid.

5.4 The scheduled jump

The announcement step against the product characteristic function. The base case is the one of the empirical study in miniature: $\sigma = 0.25$, $r = 0.04$, $T = 9/252$, no background jumps, and one compensated jump $Z \sim \text{DE}(0.5, 12, 12)$ at $t_E = T/2$, so $\mathbb{E}|Z| = 0.0833$, $\text{sd}(Z) = 0.1179$ and $\kappa_E = 0.00696867$ (the linear compensator would be 0.00699301, larger by 0.35 percent, which is the distinction drawn in Section 2). The variance of $\ln S_T$ is 0.016121, of which the diffusion contributes 0.002232 and the announcement 0.013889, or 86 percent. The cosine expansion at $N = 2^{12}$ and 2^{14} returns the same put value 4.61125402 to all digits shown, the mixture quadrature returns 4.61125402 with an error estimate of 8×10^{-14} , and the two agree to 2.5×10^{-13} ; the Black–Scholes put without the event is 1.81278775, so the event premium is 2.798466, the same for the call, as parity requires when the forward is unchanged. Against this benchmark the two-leg solver converges at order two: the put error falls from 6.04×10^{-4} on (200, 100) to 2.36×10^{-6} on (3200, 1600) with observed order 2.00 at every level, the call error from 9.18×10^{-5} to 3.55×10^{-7} with orders 2.01, 2.00, 2.00, 2.00, and on the (1600, 800) grid $C - P$ exceeds $S_0 - Ke^{-rT}$ by 8.0×10^{-6} (Figure 1; Table 20 in Appendix B.17).

The announcement operator (36) passes its identities on its own, without time stepping (Appendix B.17): it returns one for the constant function to 6.9×10^{-14} or better and e^x for e^x up to a second-order interpolation error. Dropping the compensator breaks put–call parity by $+0.6993$, which equals $S_0(\mathbb{E}[e^Z] - 1) = +0.6993$, the failure that (6) forbids; moving t_E/T from 0.10 to 0.90 changes the European put by at most 8.0×10^{-11} , Proposition 2.3(v) on the grid; and with an asymmetric law, $\text{DE}(0.3, 20, 10)$, the finite-difference smile follows the cosine smile of the correct law and not that of the mirrored law $-Z$ (Figure 6 in Appendix B.17).

Table 3: American put with the scheduled announcement: early-exercise premium with and without the event, exercise boundary around t_E , dependence on the announcement time, and convergence of the two-leg solver.

Panel A: early-exercise premium $V_{\text{am}} - V_{\text{eu}}$ by spot, grid (1600, 800), announcement at $t_E = T/2$								
S	base case $r = 0.04$				$r = 0.10$			
	V_{eu}	V_{am}	EEP	EEP, no ev.	V_{eu}	V_{am}	EEP	EEP, no ev.
85	15.557228	15.618556	0.061328	0.142391	15.360172	15.515326	0.155154	0.356072
90	11.246699	11.298386	0.051687	0.120855	11.066030	11.198935	0.132905	0.331637
95	7.502765	7.539979	0.037213	0.039834	7.350742	7.448746	0.098003	0.129255
100	4.611245	4.634190	0.022945	0.008559	4.499573	4.560910	0.061337	0.026324
105	2.689500	2.702717	0.013217	0.001103	2.618047	2.653505	0.035458	0.003224
Panel B: exercise boundary S^* (grid resolution $\pm S \Delta x$) and the continuation value at the announcement								
case	S_0^*	S_0^* no ev.	$S^*(t_E^-)$	$S^*(t_E^+)$	$\min(V_c - \varphi)$	$\Delta V(S_0, t_E)$		
$r = 0.04$	69.945	90.671	–	92.918	$+3.67 \times 10^{-4}$	+3.135454		
$r = 0.10$	75.424	92.464	–	94.109	$+3.65 \times 10^{-4}$	+3.161020		
Panel C: dependence on the announcement time t_E (trading days from today), ATM, grid (1600, 800)								
t_E (d)	in life	Eur. put	ref.	Eur. call	Am. put	EEP	IV	premium
-1.0	no	1.812774	1.812788	1.955529	1.821341	0.008567	0.25000	-0.000013
0.5	yes	4.611245	4.611254	4.754008	4.647174	0.035929	0.62185	+2.798457
1.0	yes	4.611245	4.611254	4.754008	4.645765	0.034520	0.62185	+2.798457
2.0	yes	4.611245	4.611254	4.754008	4.642785	0.031541	0.62185	+2.798457
4.5	yes	4.611245	4.611254	4.754008	4.634190	0.022945	0.62185	+2.798457
7.0	yes	4.611245	4.611254	4.754008	4.623249	0.012004	0.62185	+2.798457
8.5	yes	4.611245	4.611254	4.754008	4.614762	0.003517	0.62185	+2.798457
10.0	no	1.812774	1.812788	1.955529	1.821341	0.008567	0.25000	-0.000013
12.0	no	1.812774	1.812788	1.955529	1.821341	0.008567	0.25000	-0.000013
Panel D: American put convergence with the Rannacher restart after the condition (base case, ATM)								
M	N	V_{am}	Δ	ratio				
200	100	4.63328400	–	–				
400	200	4.63395961	6.76×10^{-4}	–				
800	400	4.63414120	1.82×10^{-4}	3.72				
1600	800	4.63418981	4.86×10^{-5}	3.74				
3200	1600	4.63420256	1.28×10^{-5}	3.81				
Richardson		4.634207						
(400, 200): IT / PSOR / BS		4.633960	4.633953	4.633953				

$\sigma = 0.25$, $T = 9/252$, $K = 100$, $Z \sim \text{DE}(0.5, 12, 12)$ compensated. Both premiums are finite-difference American minus finite-difference European values on the same grid: “EEP” with the announcement, “EEP, no ev.” without it at the same σ and half-width. Panel B: S_0^* is the boundary today ($\tau = T$), with and without the event; $\min(V_c - \varphi)$ is the smallest value over the grid of the continuation value minus the payoff at t_E ; $\Delta V(S_0, t_E) = V(S_0, t_E^-) - V(S_0, t_E^+)$ is the value of the announcement optionality at the money. t_E^- is the instant just before the announcement (a dash means the exercise region is empty on the grid, grid floor $S_{\min} = 45.68$) and t_E^+ just after it (no-event boundary at the same time: 92.918 and 94.109). The continuation value at t_E never falls below the payoff (Jensen: $V(\cdot, t_E^+)$ is convex and the compensated jump has $\mathbb{E}[S_{t_E} | S_{t_E-}] = S_{t_E-}$), so the projection at the announcement instant is never active for a vanilla; exercise before the announcement happens only where the forgone interest beats the event optionality. Panel C: the European price is flat in t_E inside $(0, T)$ and equals Black–Scholes outside; “ref.” is the product-CF COS value inside the life and the Black–Scholes value 1.812788 outside; “IV” is the implied volatility of the European put and “premium” the European put price minus its no-event Black–Scholes value.

The American put with the announcement. Table 3 is the American side of the base case on the (1600, 800) grid, and three things in it matter. First, the announcement raises the early-exercise premium near the money and lowers it deep in the money: at the money it multiplies the premium by 2.7, at $S = 105$ by about twelve, while at $S = 85$ the premium without the event is the larger one. The announcement optionality is worth most near the money, the forgone interest on the strike deep in it; raising r to 0.10 raises all premiums without changing the pattern. Second, the properties of Section 3 hold on the grid (Panel B): the continuation value at the announcement never falls below the payoff, so the projection at the announcement instant is never active (Theorem 3.3; smallest margin $+3.67 \times 10^{-4}$), the exercise region just before the announcement is empty on the grid (Corollary 3.4 in its finite-grid form), and the boundary rejoins the no-event boundary immediately after (Figure 2(b)). Third, the boundary today, 69.945 against 90.671 without the event, is the number to carry into the empirical study: with an announcement four and a half trading days away, a put with $\sigma = 0.25$ is not exercised unless it is about thirty percent in the money.

Panel C moves the announcement date t_E from one trading day before today to twelve trading days from today, which is three trading days after the option expires at $T = 9/252$. Outside the option's life the European put is 1.812774 against the Black–Scholes value 1.812788 (the discretization error of this grid) and the premium over Black–Scholes is zero to five digits; inside the life the European put is 4.611245 at every placement from 0.5 to 8.5 days, while the American put falls monotonically from 4.647174 at $t_E = 0.5$ days to 4.614762 at 8.5 days, the early-exercise premium going from 0.035929 to 0.003517. This is Proposition 2.3(v) and Proposition 3.5 on the grid: the announcement date enters the price only through the early-exercise premium, and an earlier announcement is worth more to the holder. Panel D is the self-convergence of the American solver through the announcement step and the restart: the successive differences 6.76×10^{-4} , 1.82×10^{-4} , 4.86×10^{-5} and 1.28×10^{-5} have ratios 3.72, 3.74 and 3.81, a joint order of about 1.9 in the tested range, slightly below the Kou ladders of Table 2; on the (400, 200) grid the splitting error at the money is 7×10^{-6} against a discretization difference of 6.8×10^{-4} at that level.

Background jumps with the announcement: the published benchmark. Every test so far has either background jumps or an announcement. Table 21 (Appendix B.18) compares the solver with Table 4 of the arXiv version of Leung and Santoli (2014), the one published set of American prices for Kou dynamics with a double-exponential announcement jump: two parameter sets with jump intensities of 252 and 200 per year, nine strikes and three announcement dates, 54 American and 18 European puts printed to two decimals. The product-characteristic-function cosine prices match all 18 published European values within the rounding (largest difference 0.0046), which confirms the mapping of the parameters and the shared compensator convention, and the finite-difference European price with both jump terms differs from the cosine price by at most 3×10^{-4} on the (3200, 1600) grid. Richardson extrapolation of the (1600, 800) and (3200, 1600) American values lies within half a unit of the last published digit in 54 of the 54 cells (largest deviation 0.0049), and the finer grid alone in 53 (largest deviation 0.0052). The comparison has the resolution of the published digits and no more, since the source reports no error estimate for its transform method. Both the published and the computed values fall as the announcement moves later, as Proposition 3.5 requires. The day count matters at this resolution: taking a day as $1/365$ instead of $1/252$ of a year moves the computed values by up to 9×10^{-3} , and three of the 36 cells whose announcement date depends on it would then fall outside the rounding, so the published values are matched with trading days and not with calendar days.

Same total variance, three models. Table 4 compares three ways of putting the same variance, 0.016121 of $\ln S_T$ over nine trading days, into an option price: (a) the scheduled jump of the base case; (b) Kou dynamics with the same jump law and $\lambda = 28$ per year, so that one

Table 4: Same-variance comparison of the scheduled jump (a), a Poisson–Kou model with one expected jump (b), and Black–Scholes with inflated volatility (c): prices, implied volatilities, kurtosis and tail probabilities.

<i>Panel A: European prices (COS / closed form) and Black–Scholes implied volatilities, $S_0 = 100$</i>									
K	type	(a) price	(a) IV	(b) price	(b) IV	(c) price	(c) IV	FD–COS(a)	FD–COS(b)
90	put	1.240840	0.65113	1.183683	0.63938	1.343460	0.67185	-9.5×10^{-8}	-2.0×10^{-6}
100	put	4.611254	0.62185	4.243668	0.57298	4.987244	0.67185	-9.4×10^{-6}	-2.9×10^{-5}
100	call	4.754009	0.62185	4.386424	0.57298	5.129999	0.67185	-1.4×10^{-6}	-2.5×10^{-5}
110	call	1.723886	0.66321	1.718643	0.66233	1.775557	0.67185	$+3.8 \times 10^{-6}$	-2.3×10^{-6}
<i>Panel B: American put, FD (1600, 800) with Ikonen–Toivanen splitting, and early-exercise premium</i>									
K	type	(a) V_{am}	(a) EEP	(b) V_{am}	(b) EEP	(c) V_{am}	(c) EEP		
90	put	1.246907	0.0061	1.183863	0.0002	1.344860	0.0014		
100	put	4.634190	0.0229	4.244225	0.0006	4.994650	0.0074		
<i>Panel C: moments and terminal-density probabilities of $\ln(S_T/S_0)$</i>									
		(a)	(b)	(c)					
variance c_2		0.016121	0.016121	0.016121					
excess kurtosis c_4/c_2^2		2.227	4.453	0.000					
density at 0		4.007	5.215	3.138					
$\mathbb{P}(\ln S_T/S_0 > 0.05)$		0.6198	0.5331	0.6912					
$\mathbb{P}(\ln S_T/S_0 > 0.10)$		0.3511	0.2843	0.4292					
$\mathbb{P}(\ln S_T/S_0 > 0.20)$		0.1056	0.1045	0.1139					
$\mathbb{P}(\ln S_T/S_0 < 0.02)$		0.1509	0.1945	0.1188					

All three models have the same variance of $\ln S_T$ over $T = 9/252$ ($\sigma = 0.25$, $r = 0.04$, $q = 0$): (a) one compensated jump $Z \sim \text{DE}(0.5, 12, 12)$ at $t_E = T/2$; (b) Kou with $\lambda = 1/T$ ($= 28.000$ per year, so $\lambda T = 1$; $\mathbb{P}(0 \text{ jumps}) = 0.368$, $\mathbb{P}(\geq 2) = 0.264$) and the same jump law; (c) Black–Scholes with $\sigma = 0.671855$. European smile (COS), (a) / (b) / (c): $K = 80$: 0.7304 / 0.7650 / 0.6719; $K = 90$: 0.6511 / 0.6394 / 0.6719; $K = 100$: 0.6219 / 0.5730 / 0.6719; $K = 110$: 0.6632 / 0.6623 / 0.6719; $K = 120$: 0.7277 / 0.7667 / 0.6719. Terminal densities by Fourier inversion of the characteristic function on $[-1, 1]$ (masses 0.99999, 0.99996, 1.00000). Reading: (b) spends e^{-1} of its mass on “no earnings move” (narrow diffusion core) and $1 - 2e^{-1}$ on two or more moves (fat shoulders); (a) has exactly one move; (c) spreads the same variance as a Gaussian, with the highest at-the-money value and the thinnest tails.

jump is expected in the option’s life, with probability 0.368 of none and 0.264 of two or more; and (c) Black–Scholes with the volatility inflated to 0.671855, which is the Dubinsky–Johannes representation of Remark 2.2(iii) with a Gaussian in place of the double exponential. The finite-difference prices agree with the cosine prices to 9.4×10^{-6} for (a) and 2.9×10^{-5} for (b), so the comparison is between models and not between numerical methods; the terminal laws differ in shape, with excess kurtosis 2.227, 4.453 and 0 (Appendix B.19 and Figure 7 give the densities, tails and smiles). At the money the put is 4.611254 under the scheduled jump, 4.243668 under the Poisson model and 4.987244 under the Gaussian, ratios of 0.92 and 1.08, with implied volatilities 0.62185, 0.57298 and 0.67185; the Poisson model prices the straddle about eight percent below the scheduled jump and the Gaussian about eight percent above it. Prediction P7 of the pre-registered plan (Section 8; `output/preregistration.md`) was written from this table, and its at-the-money clause is what these numbers say. Its clause about the wings needs the crossing points of the smiles to be read correctly. At ten percent out of the money the Poisson model is below the scheduled jump (0.63938 against 0.65113 at $K = 90$, 0.66233 against 0.66321 at $K = 110$) and the Gaussian is above both (0.67185); at twenty percent out of the money the order reverses, the Poisson model at 0.7650 and 0.7667 and the Gaussian at 0.6719 against 0.7304 and 0.7277 for the scheduled jump. So the statement that both alternatives under-price “the wings” holds for the Poisson model at ten percent out of the money and for the Gaussian at twenty percent out, and fails in the other two cells; I return to this when P7 is evaluated. The early-exercise premium at the money is 0.0229 under the scheduled jump, 0.0006 under the

Poisson model and 0.0074 under the Gaussian: in (a) the uncertainty is resolved at t_E and the put is then a plain diffusion put, exercised when the announcement has moved it deep into the money, whereas in (b) another jump can arrive at any time and keeps the continuation value above the payoff, and in (c) the inflated diffusion plays the same role more weakly. Models that agree on the variance thus disagree on the exercise decision by a factor of forty, so the premium cannot be borrowed from a model of equal variance, although at these parameters both premia are a fraction of a percent of the price; Section 8.7 compares the premium in the data with the half-spread of the quotes.

5.5 The time step and the early-exercise premium

Table 16 (Appendix B.12) refines time alone at fixed Δx : the Kou put on $M = 3200$ and the scheduled-jump base case on $M = 1600$, each with N from 25 to 3200 for the European solver, the splitting and the exact discrete LCP (Brennan–Schwartz), and deeper Brennan–Schwartz time ladders, to $N = 25600$, from which time-converged exact-LCP references are built by extrapolation in space. For the European put the ratios of successive time differences are 3.99 to 4.01: Crank–Nicolson with the Rannacher start is second order in time. For the exact discrete LCP they are 2.36 to 2.61 for the Kou put at the money, an observed temporal order of 1.2 to 1.4 that rises to 1.4 to 1.5 at the finer steps of the deeper ladders (ratios 2.57 to 2.81), and 2.53 to 2.90 with the announcement, an order of 1.3 to 1.5: the projected Crank–Nicolson step loses order at the free boundary, by about as much as the plan of this paper anticipated (an order between 1.3 and 1.5). The production splitting has irregular temporal ratios, 3.15 to 4.95 for the Kou put at the money, 2.18 to 3.02 at $S_0 = 110$ and 2.67 to 3.37 with the announcement, and a smaller temporal error than the exact LCP, because its first-order splitting error, of the opposite sign, offsets part of the time error of the projected step.

What this means for the ladders depends on the problem. For the Kou put at the money the temporal part of the error on the (3200, 1600) grid is 1 percent with the splitting and 9 percent with the exact LCP; on the (800, 400) grid refining time alone moves the splitting’s error from -1.46×10^{-3} to -1.45×10^{-3} , while refining space alone takes it to -3.23×10^{-5} . With the announcement the temporal part is larger: on the (1600, 800) grid it is 10 percent of the splitting’s error, 22 percent of the exact LCP’s and 1 percent of the European error. The joint order of about two found in Sections 5.3 and 5.4 is therefore a statement about the tested range, $N = M/2$ with M up to 6400, where the spatial error dominates. Within it the joint ratios of the exact LCP fall slowly, from 3.93 to 3.85 for the Kou put and from 3.69 to 3.59 with the announcement, while those of the splitting do not (3.94 to 3.97 and 3.72 to 3.81).

The model Kou-SEJ-Q-Am of the empirical study adds to the cosine European price the premium $V_{\text{FD}}^{\text{am}} - V_{\text{FD}}^{\text{eu}}$ on one grid, so this difference is the only finite-difference quantity the study uses. It converges more slowly than the American value: the spatial errors of the two values partly cancel in the difference, the temporal error of the American value does not. On the scheduled-jump ladder the premium’s ratios of successive differences are 3.25, 3.37 and 3.58 with the splitting (3.19, 3.22 and 3.17 with the exact LCP), a joint order of 1.7 to 1.8, and on the (1600, 800) grid its error is 22 percent temporal with the splitting and 41 percent with the exact LCP.

5.6 Accuracy of the production configuration

Kou-SEJ-Q-Am computes the premium on the (400, 200) grid, centered at $\ln K$ with $\ln K$ a node, and reads it at the pre-event spot by the cubic spline of Appendix B.1, a configuration that no ladder above uses. The pricing script checks this configuration on development quotes (output/results_price_checks.json): 40 puts of 12 development events, with maturities of 10 to 32 trading days and K/S from 0.93 to 1.17, priced with their own Kou-SEJ-Q parameters and rates on the production grid and on grids up to (3200, 1600). What it measures is self-

convergence of the premium toward the finest grid, not an error against an exact value. On the production grid the premium differs from the (3200, 1600) value by at most 6×10^{-5} (median 8×10^{-6}), that is, by at most 6×10^{-5} of the American price and 4×10^{-4} of the quote's half-spread, and the largest difference falls by factors of 3.6 and 4.3 on the next two grids. The finite-difference European value on the same grid errs by up to 6×10^{-4} against the cosine price, about ten times more; that error cancels in the difference, which is why the model uses the premium and not the American value itself. The development rates were close to zero (median premium 3×10^{-4}), so the study repeats the comparison with the same quotes at a rate of 5.5 percent, where the premium is larger (median 6×10^{-2}): the production grid is then off by at most 4×10^{-4} , 2×10^{-3} of the half-spread. One solve takes about 0.04 seconds. The grid study was computed on the corrected inputs of deviation (viii) in Section 8.1. The study covers puts only, in the range just given; the calls with an ex-dividend date, on which the American model's gain of Section 8.7 rests, and the maturities and moneyness of the sample outside that range (Section 7) are covered only by the tests of Section 5.3 and Appendix B.

5.7 Further checks, and what the suite does not test

Appendix B reports the checks that support the tables above without being benchmarks of their own. In a sweep of the background volatility from 0.3 to 1.5 with background jumps on a single (800, 200) grid, the European put differs from the cosine value by -1.64×10^{-4} to -4.28×10^{-4} and the American put never falls below its payoff (Appendix B.14); on a coarsely resolved kernel with a large intensity the trapezoid weights price an option worth 13.5 at 6114 on the coarsest grid, while the hat weights converge at order two when space and time are refined together (Figure 4); a march with $\lambda\Delta\tau$ up to four times the generic bound of Lemma 4.3 stays stable, as the kernel-specific scan predicts; each wrong sign or time direction of the Crank–Nicolson step shows itself on the coarsest grid (Appendix B.15); the Rannacher start removes the gamma oscillation of a plain Crank–Nicolson start at a large parabolic ratio; and delta, gamma and theta with the announcement are checked against their cosine references (Appendix B.16).

What the suite does not test matters as much as the tables.

- (1) *The time step beyond the tested range.* The fixed- Δx ladders stop at $N = 3200$ (25600 for the exact LCP); whether the joint ratios of the splitting begin to fall, as those of the exact LCP do, on grids finer than (6400, 3200) is not tested.
- (2) *Dividends.* A dividend yield is tested for American calls without jumps and, through put–call symmetry, with an announcement; it is not tested together with background jumps.
- (3) *Solver paths.* The empirical runs use the splitting; the mirrored Brennan–Schwartz sweep for calls is compared only with PSOR, in the test suite.
- (4) *Background jumps with the announcement.* The combination is tested for European options against the product characteristic function and for American puts against the published values of Leung and Santoli (2014), at the resolution of their two decimals. On this point the empirical models stay inside the tested configurations, since they carry no background jumps.
- (5) *Code paths.* The FFT forms of the announcement and jump operators inside a full time march are not exercised.
- (6) *Derivatives.* Delta, gamma and theta are checked against the cosine reference for the Kou European put and for the European put with the announcement; for American options they are checked only by self-convergence.

- (7) *Range.* The announcement tests use one symmetric law for the American put, one asymmetric law for the calls, maturities of $T = 9/252$ to $20/252$ for the base cases and three months for the published benchmark, and spots within fifteen percent of the strike (twenty for the published benchmark), the Kou tests spots within ten percent, and the stability scan stops at $\lambda\Delta\tau = 8$. The American model Kou-SEJ-Q-Am prices puts and calls with an ex-dividend date in their life, with asymmetric announcement laws and at the maturities and moneyness of the sample (Section 7); the tests above cover these configurations only as listed.
- (8) *References.* The binomial reference carries about 1×10^{-6} (its change from 20 000 to 40 000 steps), the published Kou values six decimals and the published values with the announcement two, and the Richardson limits assume an expansion that the $S_0 = 90$ ladder shows is not always present and that the American ladders satisfy in time only approximately.
- (9) *Timing.* The run times of Appendix B.20 are wall clock on one machine that was running other work.

The pytest suite in `code/tests` carries the reference numbers of Tables 1, 2 and 3 and of Tables 15 and 20 in Appendix B as literals, so that the configurations that are tested stay tested (Appendix B.22).

6 Estimation

The models of Section 8 need three groups of inputs: the law of the announcement return under the physical measure $\mathbb{P}$, which DJ-hist and Kou-SEJ-P take from the firm's past announcements; the background parameters $(\sigma, \lambda, p, \eta_1, \eta_2)$ of (1); and the announcement variance under $\mathbb{Q}$, which DJ-implied and Kou-SEJ-Q take from the option chain before the announcement. This section states the estimator of each input, summarizes its accuracy at the sample sizes of the data, and reports the estimates on the 2020–2025 primary sample of Section 7.4, the events with at least one primary expiration, development and holdout together. In-sample and out-of-sample refer here only to whether a law was fitted to the move on which it is evaluated. The two announcement laws are different objects: the historical law describes the moves that occurred, the option-implied law the moves that option prices paid for, and the gap between them is a difference between two laws, not an error of the pricing model (Section 6.7). Appendix C collects the derivations, the Monte Carlo studies, the predictive comparison in detail and the firm-level laws.

6.1 Announcement returns

For announcement i of a firm, let $x_i = \ln((S_{\text{post},i} + D_i)/S_{\text{pre},i})$ be the split- and dividend-adjusted log move between the parity-implied spots of the pre-event and post-event snapshots (Section 7.3). Besides the announcement, the interval between the two snapshots contains d_i ordinary trading days (one day for 82.4% of the events of the primary sample), so I write

$$x_i = Z_i + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, v_i), \quad v_i = \sigma_{\text{HV},i}^2 d_i/252, \quad (41)$$

with Z_i the announcement return and $\sigma_{\text{HV},i}$ the trailing background volatility of Section 6.5. The diffusion-corrected squared move and the historical announcement variance are

$$z_i^2 = x_i^2 - v_i, \quad \hat{\sigma}_J^2 = \frac{1}{n} \sum_{i=1}^n z_i^2. \quad (42)$$

Table 5: Announcement moves, 2020–2025: pooled distribution and per-firm double-exponential laws.

	Development 2020–2021	Holdout 2022–2025	All
<i>A. Pooled announcement moves (primary-sample events)</i>			
events / firms	8,407 / 1,577	13,197 / 1,702	21,604 / 1,918
share of positive moves (%)	50.2	51.0	50.7
mean up move / mean down move (%)	5.26 / 5.48	5.76 / 6.38	5.57 / 6.03
mean $ x $ / RMS of x (%)	5.37 / 7.71	6.07 / 8.57	5.79 / 8.25
diffusion-corrected $\sigma_J = (\overline{z^2})^{1/2}$ (%)	6.96	8.26	7.78
diffusion share of $\overline{x^2}$ (%)	18.8	7.1	11.0
<i>B. Standardized moves $u = x/\text{RMS}_{f,-i}$ (firms with at least 8 events)</i>			
events / firms	2,144 / 268	10,192 / 978	19,397 / 1,308
99th percentile of $ u $ (Gaussian 2.58)	4.31	3.51	3.46
share $ u > 3$ / $ u > 5$ (%; Gauss. 0.27 / 0.00)	3.26 / 0.37	1.84 / 0.22	1.78 / 0.21
log-lik. per event minus Gaussian: Laplace	+0.072	+0.016	+0.027
Student $t_{\hat{\nu}}$ ($\hat{\nu}$)	+0.075 (3.7)	+0.034 (5.9)	+0.040 (5.3)
DE, sign split	+0.072	+0.017	+0.027
DE with diffusion, MLE	+0.107	+0.035	+0.045
mean up / mean down move of u	0.948	0.981	0.965
<i>C. Per-firm laws</i>			
firms (minimum number of events)	264 (8)	296 (12)	1,017 (12)
AIC picks DE / Gaussian / t_4 (%)	12 / 71 / 17	4 / 77 / 19	11 / 66 / 23
DE picks if truth DE / Gaussian (%) [†]	13 / 7	18 / 6	17 / 6
median $\hat{\eta}_{1,E}$ / $\hat{\eta}_{2,E}$, pricing estimator	30.9 / 29.1	23.7 / 24.7	24.3 / 23.6
median $\hat{\eta}_{1,E}$ / $\hat{\eta}_{2,E}$, MLE with diffusion	31.5 / 24.6	21.6 / 21.7	22.2 / 21.5
MLE on a parameter bound (%)	12.9	0.3	0.8
sign share: dispersion ratio (p -value)	1.18 (0.02)	1.04 (0.29)	1.06 (0.10)

x is the realized announcement log move between the pre- and post-event snapshots (parity-implied spots, split and dividend adjusted); $z^2 = x^2 - \sigma_{\text{HV}}^2 d/252$ removes the diffusion of the d non-announcement trading days between the snapshots, with σ_{HV} the trailing background volatility. Missing σ_{HV} : in A the diffusion-corrected σ_J and the diffusion share use only the moves with σ_{HV} (129 / 166 / 295 moves lack it in the three columns); in C a missing σ_{HV} (64 / 121 / 237 moves) takes the median σ_{HV} of the firm's other events in the same sample, which is admissible for these descriptive full-sample fits (the predictive comparison of Table 24 uses the current event's σ_{HV} instead). B: each move divided by the root mean square of the other moves of its firm in the same sample; log-likelihood differences against a Gaussian with free mean and variance; the Student t and Laplace laws have location zero. C: The pricing estimator takes the sign split of the raw moves ($\hat{p}_E$, mean up and mean down move) and rescales both rates so that the variance of the law equals z^2 , with $\eta_{1,E} \geq 1.5$; the MLE maximizes the exact likelihood of $Z + \mathcal{N}(0, \sigma_{\text{HV}}^2 d/252)$, $Z \sim \text{DE}(p_E, \eta_{1,E}, \eta_{2,E})$. Per-firm MLE and AIC for firms with at least 12 events (8 in the development sample, which has at most eight per firm); the MLE keeps $p_E \in [0.01, 0.99]$ and each mean jump below 5 times the RMS of the firm's moves, and both median-rate rows use the firms whose MLE is interior; AIC = $2k - 2\ell$ with $k = 3$ (DE, closed-form MLE), 2 (Gaussian), 2 (t_4 , location and scale), as in Table 23. [†]Post-freeze: share of DE picks if every firm had the law of Table 23 or a Gaussian law with its standard deviation, given its own n and diffusion share (100 samples per firm). The dispersion ratio is the cross-firm variance of the share of positive moves over its binomial sampling variance under a common share, with the p -value of the χ^2 homogeneity test. Written by `code/scripts/estimate_firms.py`; numbers in `output/estimation_firms.json`.

If ε_i has mean zero and is independent of Z_i , then $\mathbb{E}[z_i^2] = \mathbb{E}[Z_i^2]$. The notation is that of a square, but z_i^2 is negative whenever $x_i^2 < v_i$; such values are averaged as they are and never floored. Means of z_i^2 over events are the definition of realized announcement variance that every section of the paper uses.

Table 5 describes the 21,604 moves of the primary sample, by 1,918 firms. In Panel A, 50.7% of the moves were positive, the mean up move was 5.57% and the mean down move 6.03%, and the diffusion-corrected announcement volatility $\hat{\sigma}_J$ was 6.96% in the development sample and 8.26% in the holdout. The diffusion term v_i accounted for 18.8% of the mean squared move in the development sample, where the background volatility was higher (Section 7.5), and for 7.1% in the holdout; without the correction, an estimator of the announcement variance would

have attributed that share to the announcement. Panel B pools the shape across firms: the moves standardized by the root mean square of the other moves of the same firm had fat tails (1.78% beyond three in absolute value, a fitted Student t law with 5.3 degrees of freedom), more than the noise of the scale alone produces, and against a Gaussian the Laplace law gained 0.027 nats (natural-log units of the likelihood) per move (Appendix C.16). Panel C fits one law per firm. The Akaike criterion picked the Gaussian for 66% of the 1,017 firms with at least twelve announcements in the primary sample and the double exponential for 11%; Section 6.3 compares these picks with those that each law would produce at the firms' own sample sizes; and the share of positive moves varied across firms by little more than binomial sampling noise around a common share of 0.512 (p -value 0.10 for homogeneity), so a firm's own share of up moves is mostly noise, which matters for the shape estimator below.

The moves come from spots that passed the spot-consistency screen of Section 7.3. The screen removed one move, that of SR on 4 February 2021, whose post-event spot was computed from another underlying's chain and gave a log move of 1.836; the move is neither an event nor a prior move of the firm's later announcements anywhere in this section or in Section 8. As a second check, I flagged any remaining move larger than 20 times the root mean square of the same firm's other moves in the primary sample, for the firms with at least 8 of them; none of the remaining moves exceeded that bound.

6.2 The announcement law under $\mathbb{P}$

The announcement jump at a known date is that of Dubinsky and Johannes (2004) and Dubinsky et al. (2019), and its double-exponential law on Kou dynamics that of Leung and Santoli (2014) (Section 2). The law $Z \sim \text{DE}(p_E, \eta_{1,E}, \eta_{2,E})$ of (4) has

$$\mathbb{E}Z = \frac{p_E}{\eta_{1,E}} - \frac{1-p_E}{\eta_{2,E}}, \quad \mathbb{E}Z^2 = \frac{2p_E}{\eta_{1,E}^2} + \frac{2(1-p_E)}{\eta_{2,E}^2}, \quad \mathbb{E}|Z| = \frac{p_E}{\eta_{1,E}} + \frac{1-p_E}{\eta_{2,E}} \quad (43)$$

(Appendix C.1): p_E is the probability of an up move, and $1/\eta_{1,E}$ and $1/\eta_{2,E}$ are the mean sizes of the up and the down moves. Given n announcement returns $z_1, \dots, z_n$ observed without noise, n_+ of them positive and $n_- = n - n_+$ negative, the maximum-likelihood estimator is in closed form,

$$\hat{p}_E = \frac{n_+}{n}, \quad \hat{\eta}_{1,E} = \frac{n_+}{\sum_{z_i > 0} z_i}, \quad \hat{\eta}_{2,E} = \frac{n_-}{\sum_{z_i < 0} |z_i|}, \quad (44)$$

the sign split: the share of up moves and the reciprocals of the mean up and mean down moves.

The estimator used in pricing. Observed moves contain the diffusion of (41), and (44) applied to them attributes that diffusion to the announcement. Kou-SEJ-P therefore keeps the shape of the sign split and replaces its scale. For an event with $n \geq 4$ prior announcements of the same firm, with moves $x_1, \dots, x_n$ and diffusion-corrected second moment $s_z^2 = n^{-1} \sum z_i^2 > 0$, let $\hat{\mu}_+$ and $\hat{\mu}_-$ be the mean up move and the mean absolute down move. The announcement law of Kou-SEJ-P is

$$\begin{aligned} \hat{p}_E &= \frac{n_+}{n}, \quad \hat{\eta}_{1,E} = \max \left\{ \frac{1}{c \hat{\mu}_+}, 1.5 \right\}, \quad \hat{\eta}_{2,E} = \max \left\{ \frac{1}{c \hat{\mu}_-}, 10^{-3} \right\}, \\ c^2 &= \frac{s_z^2}{2\hat{p}_E \hat{\mu}_+^2 + 2(1-\hat{p}_E) \hat{\mu}_-^2 - (\hat{p}_E \hat{\mu}_+ - (1-\hat{p}_E) \hat{\mu}_-)^2}, \end{aligned} \quad (45)$$

so that the variance of the fitted law equals s_z^2 (Appendix C.4); when all prior moves have one sign it falls back to the symmetric law $\hat{p}_E = \frac{1}{2}$, $\hat{\mu}_+ = \hat{\mu}_- = n^{-1} \sum |x_i|$. The variance, rather than the second moment, is matched because option prices depend on the law of Z only up to a shift: adding a constant to Z adds the same constant to κ_E in (5), so $Z - \kappa_E$ in (4) does not

change (Lemma C.1). The mean of the fitted law is a by-product of the shape. DJ-hist uses the announcement variance of the pre-registered design, the mean squared prior move $n^{-1} \sum x_i^2$, without the diffusion correction; on the events of the main study, s_z^2 had a median of 0.82 times the DJ-hist variance. When $s_z^2 \leq 0$, which happened for 196 of the 1,644 development events with at least 4 prior announcements and 145 of the 5,776 holdout events, Kou-SEJ-P is not priced, while DJ-hist is (Appendix C.4 describes the fitted laws).

The estimator (45) is a deviation from the pre-registered design, which specified a maximum-likelihood fit of the double-exponential law. The shape is the closed-form maximum-likelihood shape; the scale is a moment estimator corrected for the diffusion. The estimator was written for the study of recent announcements, which was priced before the freeze (Section 8.10), and frozen with the specification of the main study on the development sample, before any holdout quote was priced, for two reasons: it keeps the two studies on one estimator, and the development sample offers at most 7 prior announcements per event (median 5), too few for the per-firm likelihood with diffusion, which ended on a parameter bound for 12.9% of the development firms even with eight moves.

6.3 Accuracy of the announcement-law estimators

How much can one firm’s history reveal about its announcement law? The primary-sample events with at least 4 prior announcements had a median of 5 of them in the development sample and 11 in the holdout (at most 21). Appendix C.3 simulates samples of $n = 8, 12, 20$ and 40 announcement returns from a fixed double-exponential law whose standard deviation, 8.00%, is close to the pooled $\hat{\sigma}_J$ of Table 5 (Table 23, Figure 8). Three findings carry over to the pricing models. First, the announcement variance, which is the input of DJ-hist and the scale of Kou-SEJ-P, is unbiased but imprecise: the standard deviation of $\hat{\sigma}_J^2$ was 0.81 times the truth at $n = 8$ and 0.50 times at $n = 20$, and the diffusion correction removes the upward bias that the ordinary days between the snapshots add to the mean squared move (+11.9% at $n = 8$ without it, -0.3% with it). Second, the rates of (44) have no finite mean and no finite variance at any n (Appendix C.2) and are imprecise (bias +44.7% conditional on at least two up moves and standard deviation 1.04 times the truth conditional on at least three, at $n = 8$), while the standard deviation implied by the fitted law is nearly unbiased (-0.2%) and exact intervals for the rates have close to nominal coverage. Third, the shape is not identified for a single firm: when the truth was double exponential, the Akaike criterion picked it in 22.3% of the samples at $n = 8$ and 31.7% at $n = 20$, and when the truth was Gaussian it still picked the double exponential in 7.2% at $n = 8$. Across the hundreds of firms of Table 5, Panel C, the picks do carry information. A post-freeze simulation that gives every firm of the panel its own number of announcements and its own diffusion share (the row of Panel C marked †) shows what each law would produce: in the development sample, where every firm has eight announcements, the criterion would pick the double exponential for 12.7% of the firms under a double-exponential law and 7.2% under a Gaussian one, against 11.7% observed, close to the double-exponential value; in the holdout (median 12 announcements) the two laws would give 18.0% and 6.3%, against 4.1% observed, below even the Gaussian value. The development picks are thus compatible with a double-exponential law, while the longer holdout histories resemble a Gaussian law more than a double-exponential one, as does the predictive comparison below.

6.4 Alternative laws and out-of-sample prediction (exploratory)

A direct test of an announcement law is how well it predicts the next announcement of the same firm. This comparison is a post-freeze exploratory analysis: the pre-registration names no predictive comparison of announcement laws, the comparison uses the holdout, and its results were computed after the specification freeze; no priced model depends on it. For every primary-sample event with at least 4 prior announcements, I fitted each law on the firm’s prior

moves only, with information available at d_{pre} , and evaluated its density $\hat{f}_{L,i}$ at the realized move:

$$\text{LS}_L = \frac{1}{N} \sum_{i=1}^N \ln \hat{f}_{L,i}(x_i), \quad \Delta_{L,i} = \ln \hat{f}_{L,i}(x_i) - \ln \hat{f}_{G,i}(x_i), \quad (46)$$

the mean log predictive score in nats per event and its difference against the Gaussian law G with the DJ-hist variance (positive values favor L ; Table 24 in Appendix C.7). Fat tails helped: the Student t_4 law improved on the Gaussian by +0.070 nats per event in the development sample and +0.059 in the holdout, with intervals that exclude zero. The firm’s own double-exponential shape did not: the announcement law of Kou-SEJ-P scored -0.041 $[-0.087, +0.004]$ and -0.007 $[-0.041, +0.030]$ nats against a Gaussian with the same corrected variance, a difference indistinguishable from zero, and the sign-split double exponential with the DJ-hist variance did worse than the Gaussian. With a shape pooled across earlier events the same law improved on its firm-shape version by +0.128 and +0.121 nats: the double-exponential shape helps when it is estimated from the cross-section rather than from one firm. The pricing study keeps the frozen firm-shape estimator, since these comparisons include the holdout. The option-implied Gaussian improved on the corrected historical Gaussian by +0.139 and +0.129 nats, with intervals that exclude zero, but the corrected Gaussian itself scored below the uncorrected one. Against the uncorrected Gaussian of DJ-hist the option-implied Gaussian gained +0.080 $[+0.044, +0.120]$ nats in the development sample and only +0.038 $[-0.009, +0.075]$ in the holdout, and on the same holdout events it scored -0.019 $[-0.063, +0.015]$ nats against the historical Student t_4 law. The evidence that the options carry information about the coming announcement beyond the firm’s past announcements, as Dubinsky et al. (2019) find, therefore comes mainly from the development sample. These comparisons concern the law of the move; the payoff clause of prediction P8 is evaluated in Section 8.

6.5 Background parameters

The background volatility σ_{HV} of an event is the zero-mean, time-weighted volatility of the firm’s parity-implied spot returns over the trailing 120 trading days before d_{pre} , with every return that overlaps the window of a past announcement of the firm removed (Appendix D.4); the window held a median of 54 returns. If ordinary days contain jumps, σ_{HV}^2 estimates the total variance $\sigma^2 + \lambda \mathbb{E}[Y^2]$ of ordinary returns. BS-HV, CRR-HV, DJ-hist and Kou-SEJ-P use it as their background volatility; DJ-implied and Kou-SEJ-Q take the background level from option prices. All models set $\lambda = 0$ in the background.

Could the data identify a background jump intensity and its shape? The two estimators of Appendix C.8 say no. A threshold rule that calls a day a jump when its absolute return exceeds the sample standard deviation flagged a share of 0.318 of the days of Gaussian returns without jumps at $n = 2520$ and in the limit hands the pricer 2.48 times the true variance (Table 25): it measures the threshold, not the jumps. Maximum likelihood of the Kou law (Ramezani and Zeng, 2007) worked with ten years of daily returns but not at the sample sizes of a trailing window (at $n = 80$ the 5–95% range of $\hat{\lambda}$ was 47.5–1553.9 per year against a truth of 75.6; Table 26), and on returns without jumps its median $\hat{\lambda}$ was the upper bound of the search region while the total volatility was recovered: in these experiments the data identify the total variance of ordinary returns, not its split into a diffusion and many small jumps. The likelihood has a second, separate difficulty: as Honoré (1998) shows for jump-diffusions, it can be unbounded, which makes unrestricted maximum likelihood inconsistent; for the Kou law it diverges as $\sigma \rightarrow 0$, so the search region bounds σ from below (Appendix C.8). The parity-spot series is less informative still (until 30 August 2024 the snapshots came three times a week, and the window of σ_{HV} holds fewer returns than the smallest sample of Table 26). The models of the main study therefore set $\lambda = 0$ in the background, as the pre-registered design prescribed before these experiments were run, and let σ_{HV} carry the identified total variance. This does

not assert that ordinary days have no jumps; it keeps the jump content of the model where its timing is known and its law can be estimated: in the announcement.

6.6 The announcement variance under $\mathbb{Q}$

Before the announcement, the at-the-money implied variance of an expiration that contains it mixes the background variance, which grows with the time to expiration, and the announcement variance, which does not. With two such expirations $T_1 < T_2$ and at-the-money implied volatilities IV_1 and IV_2 , the decomposition of Dubinsky and Johannes (2004) solves $IV_j^2 T_j = \sigma_{bg}^2 T_j + \sigma_j^2$, $j = 1, 2$:

$$\hat{\sigma}_{bg}^2 = \frac{IV_2^2 T_2 - IV_1^2 T_1}{T_2 - T_1}, \quad \hat{\sigma}_j^{2,\mathbb{Q}} = T_1 IV_1^2 - \hat{\sigma}_{bg}^2 T_1 = \frac{T_1 T_2}{T_2 - T_1} (IV_1^2 - IV_2^2). \quad (47)$$

The relation is exact in the model of Dubinsky and Johannes (2004), Black–Scholes dynamics with a Gaussian announcement jump, because the compensated log price is then Gaussian with variance $\sigma_{bg}^2 T + \sigma_j^2$ and every strike has implied variance $\sigma_{bg}^2 + \sigma_j^2/T$. In the data the two expirations are the two shortest that contain the announcement and have an at-the-money implied volatility at d_{pre} , each implied volatility is the mean of the call and the put at the strike nearest the parity spot, and T_j is measured in calendar time (Appendix D.4; Section 6.7 repeats the decomposition on the trading clock). The decomposition was available for 99.8% of the kept events. A negative $\hat{\sigma}_j^{2,\mathbb{Q}}$, which arises when the shorter expiration’s at-the-money implied volatility is below the longer one’s, occurred for 7.4% of the 21,235 primary-sample events with both $\hat{\sigma}_j^{2,\mathbb{Q}}$ and z^2 (9.7% on the trading clock); such values stay in the means of Section 6.7, while DJ-implied is priced only when both parts of (47) are positive. Quote noise leaves the decomposition unbiased but makes it noisy for small announcements and wide spreads: its standard error ranged from 11% of the truth for the largest announcement volatility and the tightest spread to 134% for the smallest announcement volatility and the widest spread in the Monte Carlo study of Appendix C.9 (Table 27).

Kou-SEJ-Q. DJ-implied prices every option with (47) and a Gaussian announcement law. A double-exponential law with the same variance gives a lower at-the-money implied volatility: in Table 4 the scheduled double-exponential jump has an at-the-money implied volatility of 0.622 where a Gaussian with the same total variance has 0.672. Putting the decomposition’s variance into a fat-tailed law would therefore under-price at-the-money options. Kou-SEJ-Q is instead calibrated to the at-the-money quotes of the same two expirations. The shape $(\hat{p}_E, \hat{\mu}_+/\hat{\mu}_-)$ comes from (45), the announcement variance V sets the scale as s_z^2 does there, and with $IV_j^{\text{mod}}(\bar{\sigma}, V)$ the mean of the Black–Scholes implied volatilities of the model’s call and put at the at-the-money strike of expiration j and IV_j^{mkt} the same quantity for the market mids,

$$(\hat{\sigma}, \hat{V}) = \arg \min_{\bar{\sigma} > 0, V > 0} \sum_{j=1}^2 \left(IV_j^{\text{mod}}(\bar{\sigma}, V) - IV_j^{\text{mkt}} \right)^2, \quad (48)$$

by least squares in $(\ln \bar{\sigma}, \ln V)$, started from the decomposition’s values where they are positive, with at most 200 evaluations; I call a calibration exact when both residuals are below 10^{-4} in implied volatility. The background $\bar{\sigma}$ is a calendar-clock volatility, converted to the trading clock of expiration j as $\bar{\sigma}((\text{DTE}_j/365)/T_j)^{1/2}$, so that DJ-implied, Kou-SEJ-Q and its pre-registered variant share the clock of the decomposition; this departs from the trading-time convention of the plan and is recorded as a deviation with the frozen specification. When the firm’s law is undefined, because the event has fewer than 4 prior announcements or $s_z^2 \leq 0$, the shape is the pooled shape of the event’s calendar quarter, estimated from at least 250 standardized moves of events that settled before the quarter began (Appendix C.6); before 2021Q2, when no such shape exists, it is the symmetric law $p_E = \frac{1}{2}$, $\eta_{1,E} = \eta_{2,E}$.

The calibration was exact for 94.9% of the 3,845 development events with a decomposition pair and for 97.7% of the 6,132 holdout events (99th percentile of the larger residual 1.62 and 0.72 volatility points). Because firm histories were short, the shape of the development sample came from the firm for only 38% of the events (8% pooled, 54% symmetric), against 92% in the holdout. The calibrated announcement variance exceeded the decomposition’s by a median factor of 1.16 (10th and 90th percentiles 1.05 and 1.30), the at-the-money effect just described, while the background volatility barely moved (median ratio 0.986). *The calibration (48) is a deviation from the pre-registered design*, which took the announcement variance from the decomposition and the asymmetry from the pre-event skew; it is the specification of the recent-announcement study, frozen with the main study’s specification on the development sample, and the literal variant is priced alongside.

The pre-registered variant. Kou-SEJ-Q-prereg takes $V = \hat{\sigma}_J^{2,\mathbb{Q}}$ and $\hat{\sigma}_{\text{bg}}$ from (47) and a double-exponential law with mean zero, whose one free parameter p_E matches the slope of the pre-event smile of the nearest expiration (equation (66) of Appendix C.12). No admissible p_E matched that slope for 47% of the development events and 41% of the holdout events, mostly because the market’s skew was steeper than that of the closest member of the family (Appendix C.12): the pre-event skew contains more than the asymmetry of a zero-mean announcement jump on a flat diffusion.

6.7 The earnings variance premium

The ratio

$$R = \frac{N^{-1} \sum_{i=1}^N \hat{\sigma}_{J,i}^{2,\mathbb{Q}}}{N^{-1} \sum_{i=1}^N z_i^2}, \quad (49)$$

over the N primary-sample events with both quantities defined, compares the announcement variance that the options paid for before the announcement with the variance realized across it. If the two laws coincided and the background implied variance were the same at both expirations, as (47) assumes, R would be one up to sampling error. Under that assumption a ratio above one is an earnings variance risk premium, the compensation that sellers of announcement risk receive, as found by Dubinsky and Johannes (2004) and Barth and So (2014) for earnings and by Kelly et al. (2016) for political events. Gao et al. (2018), by contrast, find positive average returns to at-the-money straddles bought three days before announcements, which would correspond to a ratio below one, and read them as underestimated uncertainty; Leung and Santoli (2014) compare the risk-neutral and the historical double-exponential announcement laws in the same spirit. Either way the ratio is a statement about two laws, not about option prices being wrong: a model that uses the historical law and is evaluated against market quotes measures the difference between the laws plus its own error. Both means are pooled, so events with large announcements weigh more, and negative values of $\hat{\sigma}_{J,i}^{2,\mathbb{Q}}$ and z_i^2 are kept, which keeps both means unbiased. The intervals are the 5th and 95th percentiles of 1,000 event-bootstrap draws (seed 20261014), each subsample bootstrapped on its own.

The ratio was 1.70 [1.60, 1.82] in the development sample (8,226 events), 1.14 [1.09, 1.19] in the holdout (13,009) and 1.31 [1.27, 1.36] over all events (Table 6, Panel A); in standard-deviation units the option-implied announcement volatility exceeded the realized one by factors of 1.31, 1.07 and 1.15. The premium rose with the VIX in the development sample, from 1.14 in the lowest tercile to 2.16 [1.96, 2.39] in the highest, but hardly in the holdout (1.13, 1.13 and 1.19, with overlapping intervals; Panel B). Appendix C.10 gives the premium by decile of the implied variance, by calendar year and by tercile (Figure 9); in the holdout it fell with the firm’s past move, from 1.39 in the lowest tercile to 1.04 in the highest. These patterns bear on the first two clauses of prediction P8, whose verdict under the pre-registered failure rule is given in Section 8.

Table 6: The earnings variance premium: option-implied over realized announcement variance, 2020–2025.

	Development 2020–2021		Holdout 2022–2025		All	
	<i>n</i>	ratio [90% CI]	<i>n</i>	ratio [90% CI]	<i>n</i>	ratio [90% CI]
<i>A. Ratio $\sigma_J^{2,Q}/z^2$, all events of the sample</i>						
all events	8,226	1.70 [1.60, 1.82]	13,009	1.14 [1.09, 1.19]	21,235	1.31 [1.27, 1.36]
trading clock [†]	8,226	1.52 [1.42, 1.63]	13,009	1.06 [1.01, 1.11]	21,235	1.20 [1.15, 1.24]
<i>B. By VIX tercile at the pre-event snapshot</i>						
low	2,747	1.14 [1.01, 1.31]	5,347	1.13 [1.05, 1.20]	8,094	1.13 [1.07, 1.20]
middle	2,748	1.79 [1.63, 1.96]	5,107	1.13 [1.06, 1.22]	7,855	1.29 [1.22, 1.36]
high	2,731	2.16 [1.96, 2.39]	2,555	1.19 [1.08, 1.32]	5,286	1.61 [1.50, 1.74]
<i>C. By tercile of the firm's mean absolute past move</i>						
low	2,291	2.10 [1.81, 2.40]	3,653	1.39 [1.27, 1.54]	5,944	1.65 [1.52, 1.80]
middle	2,301	1.73 [1.48, 2.02]	5,565	1.25 [1.18, 1.34]	7,866	1.35 [1.27, 1.44]
high	2,301	1.78 [1.64, 1.93]	3,581	1.04 [0.97, 1.10]	5,882	1.22 [1.15, 1.28]
<i>D. By tercile of the share price</i>						
low	2,706	1.71 [1.54, 1.90]	3,605	1.21 [1.11, 1.32]	6,311	1.38 [1.30, 1.48]
middle	2,747	2.02 [1.85, 2.21]	4,369	1.22 [1.14, 1.31]	7,116	1.44 [1.36, 1.53]
high	2,773	1.36 [1.22, 1.51]	5,035	0.98 [0.93, 1.05]	7,808	1.08 [1.02, 1.13]
<i>E. By calendar year</i>						
2020	4,154	1.88 [1.72, 2.06]			4,154	1.88 [1.72, 2.06]
2021	4,072	1.47 [1.32, 1.61]			4,072	1.47 [1.32, 1.61]
2022			4,125	1.14 [1.05, 1.24]	4,125	1.14 [1.05, 1.24]
2023			4,207	1.05 [0.97, 1.13]	4,207	1.05 [0.97, 1.13]
2024			3,492	1.22 [1.14, 1.32]	3,492	1.22 [1.14, 1.32]
2025 (to 22 Apr.)			1,185	1.23 [1.09, 1.40]	1,185	1.23 [1.09, 1.40]
<i>F. Option-implied versus historical (events with $n_{\text{past}} \geq 4$)</i>						
$\sigma_J^{2,Q}/\sigma_J^{2,\text{hist}}$	2,997	0.90 [0.86, 0.95]	11,933	1.30 [1.27, 1.33]	14,930	1.22 [1.19, 1.24]
$\sigma_J^{2,\text{hist}}/z^2$	2,997	1.50 [1.34, 1.67]	11,933	0.88 [0.85, 0.92]	14,930	0.97 [0.93, 1.00]
$\sigma_J^{2,Q}/z^2$	2,997	1.35 [1.21, 1.51]	11,933	1.15 [1.10, 1.20]	14,930	1.18 [1.13, 1.22]
<i>G. Expected absolute announcement move</i>						
Median (%)	Development 2020–2021		Holdout 2022–2025		All	
events	2,603		11,483		14,086	
realized $ x $	3.35		4.40		4.21	
ATM straddle / spot, nearest expiration	7.80		8.01		7.97	
option-implied $\mathbb{E} Z $, Gaussian: $\sigma_J^Q \sqrt{2/\pi}$	4.86		5.54		5.40	
option-implied $\mathbb{E} Z $, firm DE shape	4.54		5.12		5.02	
historical mean $ x $ of prior moves	5.24		4.97		5.01	
historical $\mathbb{E} Z $, Kou-SEJ-P law	3.74		3.94		3.91	
median of implied / historical (ratio)	0.88		1.07		1.04	
share with implied > historical (%)	38.9		57.2		53.8	

Ratio of the mean option-implied announcement variance $\sigma_J^{2,Q}$ (two-expiry decomposition at the pre-event snapshot, negative values kept) to the mean diffusion-corrected realized second moment $z^2 = x^2 - \sigma_{\text{HV}}^2 d/252$, over primary-sample events with both defined; under a flat background term structure a ratio above one is an earnings variance risk premium (a statement about the risk-neutral and physical announcement laws, not about mispricing). The maturities of the decomposition are in calendar time, as in the priced models; [†]the trading-clock row counts them in NYSE trading days over 252 with the same total variances $\text{IV}_j^2 \text{DTE}_j/365$ (post-freeze, exploratory; by VIX tercile, with week-clustered intervals and with two placebo decompositions in Table 28). 90% intervals: 5th and 95th percentiles of 1,000 event-bootstrap draws (seed 20261014), each subsample bootstrapped on its own. Terciles use the cut points of the development sample. Data screens: the primary sample leaves out the events whose parity-implied spot fails the spot-consistency screen of Section 7.3; of the remaining moves of firms with at least 8 primary-sample moves, none is larger than 20 times the root mean square of the same firm's other primary-sample moves. F: $\sigma_J^{2,\text{hist}}$ is the mean squared prior announcement move (DJ-hist), over every prior move as in the pricing models. G: medians over events where every measure is defined; the straddle measure includes the diffusion to expiration. Written by `code/scripts/estimate_firms.py`; numbers in `output/estimation_firms.json`.

Clock and term structure (post-freeze, exploratory). Two assumptions of (47) bear on R , and I checked both after the specification freeze; neither check was pre-registered (Table 28 and Appendix C.11). The first is the clock. The decomposition counts maturities in calendar time, while the diffusion term of z^2 is on the trading clock. With the maturities counted in NYSE trading days and the same total implied variances, the decomposition gave the smaller announcement variance for 87.2% of the events, and R fell to 1.52 [1.42, 1.63] in the development sample and 1.06 [1.01, 1.11] in the holdout (Table 6, Panel A), while the share of negative decompositions rose from 9.4% to 12.8% and from 6.1% to 7.8%. These data do not show which clock option prices follow, so the two values bracket R rather than correct it. Announcements of the same week share market news; with intervals that resample announcement weeks, the holdout interval was [1.08, 1.22] on the calendar clock and [0.99, 1.13] on the trading clock.

The second is the term structure: (47) counts as announcement variance any difference between the background implied variances of the two expirations. Two placebos apply it where no announcement lies inside the pair: to the same two expirations at the post-event snapshot (placebo 1), and to two expirations of matched maturity 20 trading days later that end before the firm’s next announcement (placebo 2). Under a flat background both would average zero. Relative to the mean z^2 of the same events, placebo 1 averaged 0.48 in the development sample and 0.37 in the holdout, where it was positive for 73.4% of the events, and placebo 2 averaged 0.41 and 0.29. Placebo 1 also contains the elevated volatility that follows an announcement, and neither placebo measures the background slope before it, so I do not subtract them from R ; they show that the background term structure is not flat, by amounts of the order of the premium. On the events of each placebo and on both clocks, $R - 1$ exceeded the placebo in the development sample and fell short of it in the holdout. Placebo 1 also rose with the VIX in the development sample, by 0.51 from the lowest to the highest tercile against 1.02 for R , so part of the rise of the premium with the VIX is a term-structure effect. R is therefore an estimate that assumes a flat background term structure. The development premium survives both checks; the holdout ratio is within clock and term-structure uncertainty of one.

Panel F compares both measures with the historical announcement variance of DJ-hist on the events with at least 4 prior announcements, all of them 2021 announcements in the development sample. In the holdout the options priced more announcement variance than the history, a ratio of 1.30, and the history under-predicted the realized variance (0.88). In the development sample the order was reversed: the history, made mostly of 2020 moves, exceeded both the option-implied variance (ratio 0.90) and the realized one (1.50). The historical announcement variance therefore errs in both directions, by amounts of the order of the premium itself. Panel G compares expected absolute moves: the median realized absolute move was 4.21%, the option-implied Gaussian expected absolute move 5.40%, the historical mean absolute move 5.01% and the expected absolute move of the Kou-SEJ-P law 3.91%, lower because the corrected variance excludes the diffusion and because at a given standard deviation the symmetric double exponential expects a smaller absolute move than the Gaussian (Appendix C.10).

Implications for the pricing study. The Kou-SEJ-P law is a noisy input whose double-exponential shape adds nothing to a Gaussian with the same variance in prediction; the background is a diffusion whose variance σ_{HV}^2 is the identified total variance of ordinary returns; and Kou-SEJ-Q matches the at-the-money level of the two decomposition expirations exactly in most events, so that its comparison with DJ-implied measures the shape of the smile away from the money. A model with the physical announcement law that is evaluated against market quotes inherits the difference between the two laws as a bias, which is not mispricing, whether it is a premium or partly the slope of the background term structure, and with the historical law it also inherits the error of the history, which Panel F shows can have either sign.

7 Data

The empirical part rests on four public sources: end-of-day option quotes, an earnings calendar that records whether each report came before the open or after the close, Treasury yields and the VIX. This section describes the sources, the features of the option data that shaped the design, the construction of the announcement sample and the quote filters. Appendix D gives the cleaning rules and every intermediate count. All downloads are frozen with their retrieval dates, and every derived file is written by a script from the frozen inputs (see the data and code availability statement).

7.1 Option quotes

The quotes come from the `options` database that the DoltHub user `post-no-preference` publishes ([post-no-preference, 2025](#)). I used it in its state of 22 April 2025, frozen as a single file. It holds 79,614,858 end-of-day quotes on 2,183 underlying symbols, taken on 927 snapshot dates from 9 February 2019 to 22 April 2025. A row is one listed contract on one date, with its closing bid and ask and the vendor’s Black–Scholes implied volatility and Greeks. The database has no trade prices, volumes, open interest or prices of the underlying, so I inferred the underlying price from put–call parity (Section 7.3), and I recomputed implied volatilities and Greeks from the quotes (Section 7.4). Over the whole database, 16.9% of the quotes had a zero bid, and the median relative spread $(\text{ask} - \text{bid})/\text{mid}$ of the quotes with a positive bid was 20.0%.

Three features of the database shaped the design (Appendix D.1). First, its cadence changed twice: 48 weekly snapshots in 2019, 717 snapshots on Mondays, Wednesdays and Fridays from 4 January 2020 to 30 August 2024, which covered 59% of the trading days of that period, and 162 daily snapshot labels from 2 September 2024 to 22 April 2025. Under the three-day cadence the last close before an announcement often had no snapshot, which is why the event matching allows a lag of up to 3 trading days. Second, the vendor quoted only a few expirations per symbol and date, 2.35 on average from 2020 on, the nearest with a median of 14–15 days to expiration, and weekly expirations dropped out of the feed in their last days: only 2.8% of the weekly series were ever quoted with fewer than 11 days to expiration, against 75.6% of the monthly series. The weekly expiration right after an announcement was therefore usually quoted before the announcement but not after it; such weeklies form a secondary sample that is evaluated on the settlement payoff. Third, a snapshot dated on a weekend or a holiday carried the closing quotes of the previous trading day; I mapped every snapshot to its effective trading date and dropped the 32 snapshots whose effective date already had one.

7.2 Earnings announcements and event dates

Announcement dates come from the DoltHub database `earnings` ([post-no-preference, 2026a](#)), in its table `earnings_calendar` as retrieved on 13 October 2026: 95,112 rows for 6,875 symbols, each with a date and a timing flag that reads ‘Before market open’, ‘After market close’ or is missing (24.9% of the rows). The table starts on 22 January 2020, which fixes the start of the sample. The calendar also keeps superseded projections that rolled forward week by week before a firm confirmed its date, so I treated rows of one symbol less than 30 days apart as one announcement and kept the last date, preferring a row with a timing flag. This left 36,543 announcements by 2,085 symbols between 22 January 2020 and 22 April 2025 (Table 7, Panel A; Table 32 in Appendix D.2).

The timing flag places the announcement jump between two closes. For a report after the close of day E , the last close before it, T_{pre} , is the close of the last trading day on or before E and the first close after it, T_{post} , that of the first trading day after E ; for a report before the open of E , T_{pre} is the close of the last trading day before E and T_{post} the close of the first trading day on or after E . When the flag was missing (2,590 announcements, 7.1%), I treated

the report as made after the close, as the design prescribes, and flagged it. The pre-event snapshot d_{pre} was the latest snapshot on or before T_{pre} in which the symbol was quoted, and the post-event snapshot d_{post} the earliest on or after T_{post} , each at most 3 trading days away. 29,127 announcements matched on both sides, 11,282 in the development period and 17,845 in the holdout; of the 7,416 that did not, 6,323 were announcements of symbols that the vendor did not quote within 30 calendar days of the announcement. In the primary sample, d_{pre} fell on T_{pre} itself for 61.5% of the events and d_{post} on T_{post} for 50.2%, so the return from d_{pre} to d_{post} usually contains one ordinary trading day besides the announcement (82.4% of the events), whose diffusion the realized announcement variance of Section 6 removes. Ticker changes matter only for a firm’s history: I linked 38 changes in which the old holders kept their shares one for one and the surviving company was predominantly the old one (FB to META, for example) and split the histories of 8 tickers that a different company had used earlier (Table 33 in Appendix D.2).

A report without a timing flag that in fact came before the open has its pre-event snapshot after the release. Some of the 917 such announcements with a primary expiration (Section 7.4) look mis-timed: the at-the-money implied volatility of their nearest expiration fell from d_{pre} to d_{post} by a median of 3.8 volatility points, against 10.0 for the announcements with a flag, and it did not fall at all for 34% of them, against 12%. The design keeps them. *A check added after the freeze, and not pre-registered*, drops them: 117 of the 6,893 events of the common sample of Section 8.1. Over the nine models of the common sample, both samples and both metrics, the largest change in an event-average RMSE was 0.10 volatility points; no verdict of P1 or P5 changed, and the premium ratio (49) was 1.67 in the development sample and 1.13 in the holdout, against 1.70 and 1.14 with them (Appendix D.2).

7.3 Spot, interest rates and dividends

For every snapshot and symbol I computed a parity-implied spot from the nearest expiration with at least one day to expiration and at least 2 strikes at which both the call and the put had a positive bid below the ask. Among those strikes, the 3 nearest the money (smallest $|C_{\text{mid}} - P_{\text{mid}}|$) give $S_k = C_{\text{mid}} - P_{\text{mid}} + Ke^{-rT}$; the spot S is their median and the dispersion their range divided by the median. The relation ignores dividends and the early-exercise premia of the American quotes; the dispersion statistic and the check against closing prices below measure the consequence.

The design written before the sample was built accepted a spot when at least 2 pairs were valid and the dispersion was at most 1% of the spot. For low-priced stocks that bound is tighter than the quotes allow (in the lowest price tercile the median at-the-money half-spread was 1.23% of the spot, against 0.41% in the highest; Table 34 in Appendix D.3), so I replaced it by a tolerance rule: at least 2 valid pairs and a dispersion, in dollars, of at most the larger of 1% of the spot and twice the at-the-money half-spread. *This is a deviation from the pre-registered design*, made while building the sample, before any pricing error was computed. The plain rule is kept as a flag on every event; it would have removed 684 (7%) of the events of the final sample (Appendix D.3).

A spot can pass the tolerance rule and still be wrong: on a few dates the database files the chain of another underlying under a ticker, and the parity pairs then agree closely on the other underlying’s price. A spot-consistency screen therefore rejects a parity spot that lies more than a factor of 3 above both the previous and the next parity spot of the same ticker, or more than that factor below both, a jump that reverts at the next snapshot. Of the 1,243,993 spots with two neighbors, 3 failed, 2 of them the parity spot of SPY on the same day. They belong to 2 matched announcements, 1 of which had passed the tolerance rule: SR on 4 February 2021, whose post-event spot was 388.06 against 61.85 at the pre-event snapshot and 64.40 at the next one. An announcement with a failing spot leaves the sample, and its move enters no later announcement’s history; a failing spot also leaves the return series of the background volatility. *I added this screen after the main study had been priced*, when a check of the estimation inputs

turned up the move of SR. The threshold was set from the parity series alone. The screen corrects a data error and is recorded as an amendment of the specification freeze; both samples were then priced again on the corrected inputs with unchanged code (Appendix D.3). For the 60 most-quoted symbols the median absolute deviation of the parity spot from the daily close of the `ohlcv` table of the DoltHub database `stocks` (post-no-preference, 2026c) was 0.05% (Table 35 and Figure 13 in Appendix D.3).

The interest rate is the 1-month Treasury constant-maturity yield from the table `us_treasury` of the DoltHub database `rates` (post-no-preference, 2026b), completed with the FRED series DGS1MO (Board of Governors of the Federal Reserve System, 2026) (20 calendar days, counting days carried forward from a FRED value) and carried forward on the 11 trading days without a value in either source; I used the yield of d_{pre} (of d_{post} for post-event quotes) as a continuously compounded rate. It ranged from 0.00% to 6.02% between 2019 and mid-2025. The VIX close at d_{pre} , from the FRED series VIXCLS (Chicago Board Options Exchange, 2026), defines the VIX terciles. Dividends come from the `dividend` table of the DoltHub database `stocks` (29,429 cash dividends of 1,497 symbols). When an ex-dividend date fell in an option’s life, I set the dividend yield of that option to $q = D/(ST)$, with D the dividends going ex in the life and T the time to expiration in trading years, and otherwise $q = 0$; 23% of the kept quotes had an ex-dividend date in their life. Splits come from the `split` table, with duplicated rows removed; no sample event had a split between d_{pre} and d_{post} . The realized announcement move is $\ln((S_{\text{post}} + D)/S_{\text{pre}})$, with parity-implied spots at both snapshots and D the dividends going ex between them, adjusted for splits.

7.4 Expirations and quote filters

For each matched announcement whose spot passed the tolerance rule and the consistency screen on both dates, the primary expirations were those quoted at both d_{pre} and d_{post} that expire on or after T_{post} , so that the announcement lies in the option’s life, with 5 to 45 calendar days to expiration at d_{pre} . There were 30,544 such (event, expiration) pairs in 21,604 events. These events form the primary sample; Section 6 estimates the announcement laws and the variance premium on all of them, development and holdout together. Weekly expirations at most 7 days after the announcement that were quoted at d_{pre} but were not primary form the secondary sample (1,837 pairs).

For every quote I computed the mid, the relative spread, and the European Black–Scholes implied volatility of the mid, the bid and the ask by bisection, at the parity spot, the 1-month rate, the dividend yield of the option’s life and the time to expiration in trading years. The options are American; the European inversion is a quoting convention that expresses pricing errors in volatility units and supplies the delta for the moneyness filter and the vega that converts dollar errors, not a pricing model. On the kept quotes my implied volatilities had a correlation of 0.979 with the vendor’s (Appendix D.5).

I applied the filters of the pre-registered design in the following order: bid > 0; ask > bid; relative spread at most 20% (at most 10% for a core subsample); mid at least \$0.20; a mid strictly inside the European no-arbitrage bounds, so that an implied volatility exists; implied volatility in (0.05, 3); $0.15 \leq |\Delta_{\text{BS}}| \leq 0.85$ at the quote’s own implied volatility; and, for each (event, expiration) pair, at least 2 strikes passing every other filter on the call side and on the put side. Table 7 follows the quotes through the filters. Of 792,270 pre-event quotes, 151,969 (19.2%) were kept, in 9,996 events, 13,897 (event, expiration) pairs and 1,281 firms. The spread and delta filters removed most quotes; the strikes rule then removed whole pairs: 16,035 events had at least one quote that passed every quote-level filter, but only 9,996 had a pair with 2 such calls and 2 such puts. Applying the rule to each side separately would have kept 161,685 quotes in 12,663 events; the sample follows the rule as written, and the one-sided variant is stored as a flag. The 14,873 puts and 11,681 calls that passed the first four filters with a mid at or below the European lower bound were deep in the money, and for puts such a mid is what an

Table 7: Construction of the main sample (announcements 2020-01-22 to 2025-04-22). Panel A follows announcements from the cleaned earnings calendar to the primary sample; a primary expiration has $5 \leq \text{DTE} \leq 45$ at d_{pre} and is quoted at both d_{pre} and d_{post} . Panel B applies the quote filters cumulatively, in the order listed, to every pre-event quote of the primary expirations; each row gives the quotes and the events with at least one quote left after that filter; the strikes rule requires at least two strikes on each of the call and the put side of an (event, expiration) pair. Implied volatilities and deltas are European Black–Scholes values at the parity-implied spot, the 1-month Treasury rate and T in trading years. Panel C is the secondary sample of earnings-week weeklies quoted only before the announcement. Development: 2020–2021; holdout: 2022 to 22 April 2025.

	Quotes			Events		
	Dev.	Holdout	All	Dev.	Holdout	All
<i>A. Announcements</i>						
Cleaned announcements in the window				13,831	22,712	36,543
Matched to a pre- and a post-event snapshot				11,282	17,845	29,127
Parity spot on both dates				10,435	16,839	27,274
Parity spot within tolerance on both dates				9,335	14,517	23,852
and consistent with the neighboring snapshots				9,334	14,517	23,851
With ≥ 1 primary expiration				8,407	13,197	21,604
primary (event, expiration) pairs				11,874	18,670	30,544
With ≥ 1 quote in the database at d_{pre}				8,407	13,197	21,604
With ≥ 1 kept quote				3,861	6,135	9,996
(event, expiration) pairs with ≥ 1 kept quote				5,289	8,608	13,897
firms with ≥ 1 kept quote				1,003	1,052	1,281
<i>B. Pre-event quotes of the primary expirations (cumulative filters)</i>						
All strikes, calls and puts	299,894	492,376	792,270	8,407	13,197	21,604
bid > 0	256,632	429,701	686,333	8,407	13,197	21,604
ask $>$ bid	256,630	429,339	685,969	8,407	13,197	21,604
relative spread $\leq 20\%$	123,179	211,607	334,786	7,411	11,103	18,514
mid $\geq \$0.20$	122,723	210,552	333,275	7,409	11,103	18,512
mid inside European no-arbitrage bounds	111,252	195,469	306,721	7,254	10,962	18,216
implied volatility in $(0.05, 3)$	110,751	195,462	306,213	7,191	10,962	18,153
$0.15 \leq \Delta_{\text{BS}} \leq 0.85$ at own IV	62,927	107,543	170,470	6,468	9,567	16,035
≥ 2 strikes per side and expiration = kept	54,840	97,129	151,969	3,861	6,135	9,996
Memo: kept, relative spread $\leq 10\%$ (core)	25,139	53,595	78,734	3,393	5,615	9,008
Memo: kept, same contract quoted at d_{post}	42,817	77,464	120,281	3,822	6,091	9,913
Memo: kept, post-event quote passes filters	22,234	39,274	61,508	3,451	5,470	8,921
Memo: strikes rule on the own side only	59,085	102,600	161,685	5,008	7,655	12,663
Memo: mid $\leq$ European lower bound, puts	7,022	7,851	14,873			
Memo: mid $\leq$ European lower bound, calls	4,449	7,232	11,681			
<i>C. Earnings-week weeklies (pre-event quotes, settlement payoff)</i>						
(Event, weekly expiration) pairs				437	1,400	1,837
Quotes, all strikes	8,170	30,318	38,488	437	1,400	1,837
kept (same filters as Panel B)	390	1,887	2,277	63	275	338
kept, with a validated settlement spot	379	1,768	2,147	61	256	317

early-exercise premium allows (Table 37 in Appendix D.5).

The filters changed the composition of the sample. The spread filter is relative, so it removed most quotes of low-priced and illiquid names: by development-sample tercile of the spot, 29%, 42% and 64% of the events with a primary expiration kept at least one quote, and by tercile of the event’s median relative spread 90%, 49% and 6% (Appendix D.6 gives the shares by VIX and past-move tercile). The results of Section 8 therefore concern the liquid, higher-priced part of the optionable cross-section, and the pre-registered spread cut rests mainly on the first two terciles.

7.5 The sample

Table 38 and Figure 14 in Appendix D.5 describe the kept sample by year, quarter, moneyness and days to expiration. The median kept quote had 25 days to expiration and a relative spread of 9.7%, and 72% of the kept (event, expiration) pairs were standard monthly expirations. The median event had an at-the-money implied volatility of 44.6% and a straddle-implied move of 8.5% against a median absolute realized move of 4.3%; the straddle prices the option’s whole remaining life, several weeks of ordinary diffusion besides the announcement, so these two numbers do not measure the same object. The delta filter kept strikes close to the money: the 1st and 99th percentiles of $\ln(K/S_{\text{pre}})$ were -0.19 and 0.21 . Each event carries the inputs the models need at d_{pre} (Table 36 in Appendix D.4): the background volatility σ_{HV} of Section 6.5, available for 99.5% of the kept events; the two-expiry decomposition of Section 6.6, available for 99.8% and negative for 2.9% of the events where it was available; and the firm’s past announcement moves. Because the calendar starts on 22 January 2020, firm histories are short early in the sample: only 43% of the development events had the 4 past announcements that DJ-hist and Kou-SEJ-P require, against 94% of the holdout events.

The development sample consists of the announcements from 22 January 2020 to 31 December 2021, with 3,861 events and 54,840 kept quotes; the holdout consists of those from 1 January 2022 to 22 April 2025, with 6,135 events and 97,129 quotes. The filters, the tercile cut points (Table 39 in Appendix D.6) and the specification choices were fixed on the development sample, and the analysis plan with the predictions tested in Section 8 was written before any pricing error was computed. The two periods differed beyond the history length: the median 1-month rate rose from 0.07% to 4.67%, the median VIX at d_{pre} fell from 21.2 to 18.8, and the median background volatility from 36% to 30%. After the announcement, 40.5% of the kept quotes had a post-event quote of the same contract that passed the quote filters (79.1% had one at all), so every post-event comparison runs on a smaller set of quotes than the pre-event one; on the linked kept quotes the median change in implied volatility from d_{pre} to d_{post} was -6.7 volatility points, as expected when the announcement variance leaves the options’ remaining life.

7.6 Announcements after April 2025

The frozen database ends on 22 April 2025. For the announcements from 1 May 2025 to 10 October 2026 I queried the live table `option_chain` of the same DoltHub database on 13 October 2026. I started from a seed list of 40 large, actively traded firms with large past announcement moves; the 31 of them with at least one announcement in the window and quotes at the time of collection form the universe. Of their 176 announcements, 165 had both a pre-event and a post-event chain and 11 had no snapshot within the probed days. In this period the database quoted a handful of expirations per symbol and date at roughly constant tenors, the nearest about two weeks out, so no earnings-week weekly is available. The spot is parity-implied as above (median absolute deviation from the close 0.12%), and the rate is the 1-month Treasury yield of the snapshot date. Section 8.10 states the remaining differences in design and reports the results. The latest snapshot, of 12 October 2026, also gives model prices for all 24 announcements scheduled after it; Appendix E records them before the events.

8 Results

This section tests the models of Section 6 on the sample of Section 7 against the pre-registered plan. I wrote the analysis plan of this section on 13 October 2026, before any announcement had been matched to its option chains and before any model price or pricing error had been computed. The plan was not deposited with a registry or time-stamped by a third party. It is a file of the replication package (`output/preregistration.md`, SHA-256 beginning `c75ee59a8931df9f`), and the time it was written rests on my own record and on the file’s modification time, which a reader cannot check independently; the freeze record described below hashes the pricing code, not the plan. In this paper, pre-registered means written down in that file before any announcement was matched to its chains or priced. Before this study I had worked with the same options database in a master’s project of April 2025, which applied an earlier and faulty implementation of a Kou model with an announcement jump to an earnings-week trading exercise; it computed none of the statistics of this section, and none of its results are used here. The plan fixes the samples, the models, the primary metric, the tests, the cuts, eight predictions (P1–P8) and the rule that decides when a prediction fails, and it will be deposited with the replication archive (see the data and code availability statement, which also gives the address of the code repository). The plan splits the announcements into a development sample, from 22 January 2020 to 31 December 2021, on which the specification could be adjusted, and a holdout, from 1 January 2022 to 22 April 2025, to be priced once with the specification frozen. I froze the specification on 14 October 2026 at 22:05 UTC, after the development sample had been priced and before any holdout quote was; the freeze record `output/spec_freeze.json` holds the time, every setting and the SHA-256 hashes of the pricing code (and, in later amendments, of its input files), and no parameter or model changed afterwards; a data correction found after both samples had been priced was recorded as an amendment, and both samples were priced again with the frozen code (deviation (viii) below); later amendments, to non-numerical code and to the ex-ante table of Appendix E, changed no price of this section. Every table reports the two samples separately, and every pre-registered test is reported, including those that fail; Appendix F holds the secondary tables.

The main findings are three. First, Black–Scholes with the trailing historical volatility (BS-HV) was far too cheap before an announcement, and every event-jump model reduced its error; the models that use the event’s own at-the-money quotes came closest to the market, and most of their advantage is the at-the-money level, which they fit by construction. Second, the double-exponential announcement law, with the asymmetry and tail ratio of the firm’s past moves, did not improve on a Gaussian one with the same calibrated level: Kou-SEJ-Q was worse than DJ-implied away from the money, and its large gain over Kou-SEJ-P, which takes its whole law from the firm’s past moves, measures the option-implied announcement variance, not the shape. The pre-registered variant, which fits the asymmetry to the skew slope but keeps the decomposition’s variance, was also worse in the wings; a law whose shape is calibrated to the smile was not priced. Third, the option-implied announcement variance exceeded the realized one on average, a difference between the risk-neutral and the realized law and not a pricing error; in the holdout the ratio was within clock and term-structure uncertainty of one (Section 8.8). Of the eight predictions (stated in Table 12), P5 and P7 held in both samples, P3 and P8 in the holdout only, and P1, P2, P4 and P6 failed in both (Section 8.9).

8.1 Models, samples and metrics

Models. Every model was priced at the pre-event snapshot d_{pre} with information of that close, apart from the known dividends noted below: the parity-implied spot, the 1-month rate, the dividend yield of the option’s life (Section 7.3) and the event inputs of Section 6. Nine models are priced on every event, eight European and one American; a tenth, Kou-SEJ-Q-Am, is priced on a stratified tenth of the events. **BS-HV** is Black–Scholes with the background volatility σ_{HV} ,

and **CRR-HV**, the American one, the binomial price with the same volatility. **BS-IVlong** uses the at-the-money implied volatility of the longer expiration of the decomposition pair (47); **Mkt-ATM** uses that of the quote’s own expiration and is the market reference of P3 rather than a forecast. **DJ-hist** is the deterministic-time Gaussian announcement jump of Dubinsky and Johannes (2004) on Black–Scholes, with total variance $\sigma_{\text{HV}}^2 T + \sigma_J^{2,\text{hist}}$ and the firm’s mean squared past announcement move (at least 4 of them) as $\sigma_J^{2,\text{hist}}$; **DJ-implied** is the same with the option-implied background and announcement variance of (47) (Dubinsky and Johannes, 2004; Dubinsky et al., 2019), and inside the two decomposition expirations it coincides with Mkt-ATM. **Kou-SEJ-P** is Kou-SEJ with $\lambda = 0$, background σ_{HV} and the historical ($\mathbb{P}$) announcement law (45), the law of the comparison in Leung and Santoli (2014), whose Kou model with a double-exponential announcement jump Kou-SEJ adopts. **Kou-SEJ-Q**, the primary option-implied specification, takes the shape of that law and calibrates the background volatility and the announcement variance to the at-the-money implied volatilities of the two decomposition expirations (48); the fit was exact to 10^{-4} in all but 15 of the 6,893 events of the common sample. **Kou-SEJ-Q-prereg**, the literal pre-registered variant (66), puts the decomposition’s background and $\sigma_J^{2,Q}$ into a zero-mean double-exponential law whose sign probability matches the slope of the nearest expiration’s smile; in 2,800 of the 6,893 events (41%) no admissible law matched that slope, and the closest one was used. **Kou-SEJ-Q-Am** adds to Kou-SEJ-Q the early-exercise premium of the finite-difference solver of Section 4, on one tenth of the events stratified by year and past-move tercile.

Samples and metrics. I priced the 151,969 kept quotes of 9,996 events (Table 7). The comparisons use the common sample on which all nine models are defined: 19,373 quotes in 1,386 events of 636 firms in the development sample and 87,516 quotes in 5,507 events of 940 firms in the holdout. DJ-hist and Kou-SEJ-P need 4 prior announcements and the calendar starts on 22 January 2020, so the development part of the common sample consists of the announcements of 2021 only. Errors are model minus market, in volatility points, that is percentage points of implied volatility (‘vol. points’ or ‘IV points’ in the tables). The pre-registered primary metric is the price error divided by the Black–Scholes vega at the quote’s own implied volatility, and it decides every verdict. The descriptive tables use the exact implied-volatility difference, a deviation in presentation: it is the metric of the recent study (Section 8.10), and it avoids the first-order vega approximation, which misstates errors of several volatility points away from the money (BS-HV’s RMSE for $0.15 \leq |\Delta| < 0.35$ is 10.33 vega-scaled against 11.47 in implied volatility, development sample); since the ranking of option models can depend on the loss (Christoffersen and Jacobs, 2004), I report both. Each statistic is a per-event RMSE or mean averaged over events, the event being the unit of inference; 90% intervals are percentiles of 1,000 bootstrap draws that resample events with all their quotes. Mispricing captured, $\text{MC} = 1 - \sum |V_{\text{model}} - \text{mid}| / \sum |V_{\text{bench}} - \text{mid}|$, is a ratio of sums of dollar errors over quotes. The event bootstrap, which the plan names, treats announcements as independent, while announcements cluster in weeks with shared market news and within firms, and the tests of Section 8.5 cluster by week. As a post-freeze sensitivity check, not pre-registered, the tables also give intervals from a bootstrap that resamples announcement weeks instead (column ‘Week-cl.’; 48 development and 167 holdout weeks). Over the counted components of the predictions the median ratio of their width to that of the event intervals is 1.4 in the development sample and 1.4 in the holdout, and the largest ratio 3.4; for the vega-scaled RMSEs of Table 8 the median ratios are 1.6 and 2.1. Resampling firms widens the intervals less (median ratios 1.1 and 1.4). Every verdict of this section uses the event bootstrap; Section 8.9 lists the component verdicts that the week-clustered intervals would change.

What each comparison measures. BS-HV, CRR-HV, DJ-hist and Kou-SEJ-P use no option price and are out of sample everywhere; their distance from the market combines model error, the noise of estimating an announcement law from a few past moves, and the earnings

variance risk premium (Section 8.8). Mkt-ATM, DJ-implied and Kou-SEJ-Q reproduce the at-the-money implied volatilities of the two decomposition expirations by construction, and 99.5% (development) and 99.3% (holdout) of the common-sample quotes lie in those expirations. Their advantage over the first group is therefore largely the value of the market’s at-the-money level, not evidence about the announcement law, and their differences from one another measure the smile away from the at-the-money strike: flat for Mkt-ATM and DJ-implied, double-exponential for Kou-SEJ-Q. Their out-of-sample test is on those strikes and on expirations outside the pair, of which the main sample has few (624 holdout quotes in 50 events) and the recent sample of Section 8.10 one per chain.

Deviations from the pre-registered plan. Each is recorded in the freeze record or the analysis output.

- (i) The primary Kou-SEJ-Q is the calibration (48) of the recent study, chosen to keep the two studies on one specification; the literal pre-registered variant is reported in every table and verdict. The recent study of Section 8.10 was priced, and its results seen, on 13 October 2026, before the freeze, and deviations (i), (ii) and (iv) adopt its calibration, its estimator and its descriptive metric; it is therefore not an independent test of the frozen specification.
- (ii) Kou-SEJ-P uses the moment estimator (45) instead of maximum likelihood, because in the development period a firm had at most eight announcements, fewer than the twelve that the per-firm likelihood fit requires.
- (iii) Option-implied volatilities accrue in calendar time, the clock of the decomposition, and σ_{HV} in trading time; CRR-HV uses a binomial control variate.
- (iv) The descriptive tables use the implied-volatility difference. The first complete version of the analysis judged the predictions on it without labeling the change; a re-computation from the stored prices flagged this, and the verdicts were moved to the pre-registered metric, with the implied-volatility verdicts kept as a sensitivity check.
- (v) The pre-registered failure rule leaves open how components combine and which statistic represents each prediction. I fixed these rules after the headline RMSEs of both samples had been displayed: three labels (holds, inconclusive, fails); a prediction holds if every counted component holds and fails if any fails; components with fewer than 30 events are not judged; RMSE differences below 10^{-6} volatility points are ties; and each prediction has the statistic given in the notes of Table 13. No model or setting of the specification changed afterwards, and the one re-pricing on changed inputs, that of deviation (viii), applied the frozen code to corrected inputs; the order of these steps cannot be verified from the files.
- (vi) The earnings-week weeklies of the secondary sample were priced in the analysis step, from the frozen per-event parameters.
- (vii) The dividend yield uses the realized dividend table, a known-dividend assumption (declaration dates are not in the data), for 35,430 of the 151,969 kept quotes; the same yield enters the model and the market implied volatilities. The spot rule of Section 7.3 is a further deviation, labeled there.
- (viii) The spot-consistency screen of Section 7.3 was added after both samples had been priced. It removed 1 announcement, none of whose quotes had been kept, from the primary sample and its move from the histories of the firm’s later announcements, and it corrected σ_{HV} of 2 other announcements; the development tercile cut points, and with them the strata of the American subsample, were recomputed by the frozen rules. Both samples were priced

again with the frozen code, and every number of this section comes from that pricing. Before the correction the holdout common sample had 5,504 events and 87,502 quotes, against 5,507 and 87,516 after it: with the corrected σ_{HV} of the two SSD announcements, the diffusion-corrected second moment of the firm’s past moves, which had not been positive, became positive for its three holdout announcements with kept quotes, so that Kou-SEJ-P, and with it the common sample, was defined for them (14 quotes). Through the recomputed strata the American subsample of the holdout changed from 546 events and 9,303 quotes to 545 and 9,246. No event-average RMSE recorded in the analysis log moved by more than 0.01 points, and every verdict and the model confidence sets were unchanged (output/results_analysis_log.txt).

8.2 Pre-event pricing errors

Table 8 reports the errors at d_{pre} . BS-HV was far too cheap: its event-average RMSE was 12.84 [12.41, 13.33] volatility points in the development sample and 15.59 [15.25, 15.94] in the holdout, with mean errors of -10.25 and -12.79 ; a volatility estimated from ordinary days leaves out the announcement variance. The American binomial price changed little (CRR-HV 12.16 and 14.45). Adding the firm’s historical announcement variance removed most of the bias (DJ-hist 0.78 and -1.88) but much less of the dispersion: DJ-hist and Kou-SEJ-P stayed 8.99 and 9.00 points from the market in the development sample and 8.91 and 9.40 in the holdout. A few past moves are a noisy estimate of the variance the market prices for the next announcement, and the market prices the risk-neutral law. The flat at-the-money volatility of the longer expiration did better than either (BS-IVlong 6.98 and 8.81): one option-implied number carried more information than the firm’s announcement history.

The models that use the event’s own at-the-money quotes were closest: Mkt-ATM at 4.04 and 4.55, DJ-implied at 4.04 and 4.56, Kou-SEJ-Q at 4.20 [4.00, 4.43] and 4.65 [4.55, 4.76]. This mainly records that they fit the at-the-money level. Within the group the double-exponential law was slightly worse than the flat implied volatility in both samples and under both metrics (vega-scaled 3.81 against 3.58, and 4.24 against 4.04). The literal variant Kou-SEJ-Q-prereg was biased low (-2.94 and -3.90): it gives the double-exponential law the variance of the Gaussian decomposition without refitting the level, and at equal variance a fat-tailed announcement law has a lower at-the-money implied volatility (Table 4). This at-the-money effect of the equal-variance comparison, on which P7 builds, is why the primary specification calibrates at-the-money prices rather than the variance.

Figure 3 shows the errors across moneyness. Every model under-priced the outer bins of standardized moneyness, in part by selection: the delta filter keeps a strike far from the money only when its own implied volatility is high. The in-the-money bin ($0.65 < |\Delta| \leq 0.85$) had the largest RMSE for every model, as prices there are dominated by intrinsic value and most exposed to early exercise and dividends. In the broad sample that includes 2020 and on the core quotes with a relative spread of at most 10%, the option-implied models again came first, and on the few out-of-calibration quotes of the holdout Kou-SEJ-Q had a lower RMSE than DJ-implied while Mkt-ATM was lower still (Table 41 in Appendix F.1).

8.3 Mispricing captured

Table 9 reports the share of BS-HV’s dollar error that each model removed. Kou-SEJ-Q captured 0.79 [0.76, 0.82] in the development sample and 0.80 [0.79, 0.81] in the holdout, DJ-implied 0.83 and 0.82, and Kou-SEJ-P 0.28 [0.21, 0.34] and 0.42 [0.40, 0.44]; the RMSE reductions in the same table, the statistic of P1, group the models in the same way. For DJ-implied and Kou-SEJ-Q most of the captured mispricing is the at-the-money level they are fitted to, so their MC measures how far BS-HV is from the market’s level rather than the merit of an announcement law. For Kou-SEJ-P, which uses no option price, MC is an out-of-sample number, and what it leaves

Table 8: Pre-event pricing errors of the main study, common sample (kept quotes at d_{pre} for which every model of Section 8.1 is defined). Errors are model minus market in vol. points, per event and averaged over events, with 90% event-bootstrap intervals (1,000 draws). RMSE vega, the pre-registered primary metric (price error over the Black–Scholes vega at the quote’s own IV), decides the verdicts of Tables 12 and 13; the other columns use the implied-volatility difference, a deviation in presentation (see the note).

Model	Pre-registered		Implied-volatility difference			Price			
	RMSE vega [90% CI]	Week-cl.	RMSE [90% CI]	RMSE pooled	Bias	RMSE (\$)	MAE (\$)	MAPE (%)	Inside (%)
<i>A. Development 2021 (common sample): 1,386 events, 636 firms, 19,373 quotes</i>									
BS-HV	11.29 [10.93, 11.74]	[10.34, 12.51]	12.84 [12.41, 13.33]	19.36	−10.25	1.543	1.477	18.7	20.0
CRR-HV	10.96 [10.58, 11.40]	[10.02, 12.18]	12.16 [11.73, 12.64]	18.68	−9.78	1.515	1.454	18.4	20.7
BS-IVlong	6.23 [6.05, 6.43]	[5.85, 6.67]	6.98 [6.74, 7.24]	9.12	−5.05	0.737	0.646	8.1	47.6
Mkt-ATM	3.58 [3.44, 3.74]	[3.43, 3.76]	4.04 [3.85, 4.27]	6.00	−1.56	0.398	0.316	3.6	77.1
DJ-hist	8.62 [8.20, 9.05]	[7.96, 9.42]	8.99 [8.58, 9.42]	16.66	0.78	1.111	1.038	13.5	35.3
DJ-implied	3.58 [3.44, 3.74]	[3.43, 3.76]	4.04 [3.85, 4.27]	5.99	−1.56	0.398	0.316	3.6	77.1
Kou-SEJ-P	8.46 [8.08, 8.87]	[7.88, 9.17]	9.00 [8.60, 9.41]	16.35	−1.83	1.123	1.043	13.0	34.5
Kou-SEJ-Q	3.81 [3.67, 3.97]	[3.64, 4.00]	4.20 [4.00, 4.43]	6.14	−1.18	0.447	0.365	4.0	73.7
Kou-SEJ-Q-prereg	4.00 [3.86, 4.16]	[3.79, 4.22]	4.44 [4.25, 4.66]	6.32	−2.94	0.465	0.398	4.9	66.0
<i>B. Holdout 2022–2025 (common sample): 5,507 events, 940 firms, 87,516 quotes</i>									
BS-HV	13.71 [13.39, 14.05]	[13.01, 14.47]	15.59 [15.25, 15.94]	26.39	−12.79	1.691	1.616	22.3	16.0
CRR-HV	13.19 [12.87, 13.53]	[12.50, 13.96]	14.45 [14.13, 14.79]	25.44	−11.85	1.640	1.558	21.9	17.9
BS-IVlong	7.88 [7.76, 8.01]	[7.51, 8.31]	8.81 [8.66, 8.98]	11.01	−6.50	0.864	0.756	10.3	38.7
Mkt-ATM	4.04 [3.96, 4.12]	[3.88, 4.20]	4.55 [4.45, 4.66]	6.30	−1.64	0.419	0.333	4.1	69.1
DJ-hist	8.37 [8.09, 8.69]	[7.95, 8.83]	8.91 [8.63, 9.23]	20.70	−1.88	0.968	0.889	12.8	33.8
DJ-implied	4.04 [3.97, 4.12]	[3.89, 4.20]	4.56 [4.45, 4.67]	6.30	−1.64	0.419	0.333	4.2	69.1
Kou-SEJ-P	8.76 [8.48, 9.09]	[8.32, 9.24]	9.40 [9.12, 9.73]	21.06	−3.94	1.027	0.945	13.3	31.6
Kou-SEJ-Q	4.24 [4.16, 4.32]	[4.09, 4.40]	4.65 [4.55, 4.76]	6.31	−1.11	0.454	0.370	4.6	66.6
Kou-SEJ-Q-prereg	4.92 [4.84, 5.00]	[4.74, 5.08]	5.47 [5.36, 5.58]	7.30	−3.90	0.541	0.470	6.0	50.4
<i>C. All 2021–2025 (common sample): 6,893 events, 1,016 firms, 106,889 quotes</i>									
BS-HV	13.22 [12.95, 13.51]	[12.60, 13.91]	15.04 [14.75, 15.33]	25.26	−12.28	1.661	1.588	21.6	16.8
CRR-HV	12.74 [12.47, 13.02]	[12.11, 13.42]	13.99 [13.71, 14.27]	24.36	−11.43	1.615	1.537	21.2	18.5
BS-IVlong	7.55 [7.44, 7.66]	[7.21, 7.91]	8.45 [8.31, 8.58]	10.69	−6.21	0.839	0.734	9.9	40.5
Mkt-ATM	3.95 [3.87, 4.01]	[3.81, 4.09]	4.45 [4.35, 4.54]	6.24	−1.62	0.415	0.329	4.0	70.7
DJ-hist	8.42 [8.19, 8.69]	[8.06, 8.84]	8.93 [8.70, 9.19]	20.03	−1.34	0.997	0.919	12.9	34.1
DJ-implied	3.95 [3.88, 4.02]	[3.82, 4.10]	4.45 [4.36, 4.55]	6.25	−1.63	0.415	0.330	4.0	70.7
Kou-SEJ-P	8.70 [8.46, 8.96]	[8.33, 9.12]	9.32 [9.09, 9.59]	20.29	−3.51	1.046	0.964	13.3	32.1
Kou-SEJ-Q	4.15 [4.08, 4.22]	[4.02, 4.30]	4.56 [4.47, 4.65]	6.28	−1.13	0.452	0.369	4.5	68.0
Kou-SEJ-Q-prereg	4.73 [4.65, 4.81]	[4.58, 4.90]	5.26 [5.16, 5.36]	7.14	−3.70	0.525	0.455	5.8	53.6

99.5% (development) and 99.3% (holdout) of the quotes belong to the two expirations whose ATM IVs Mkt-ATM, DJ-implied and Kou-SEJ-Q reproduce by construction (Kou-SEJ-Q to within 10^{-4} in IV in all but 15 of the 6,893 events), so their differences from BS-HV, DJ-hist and Kou-SEJ-P largely measure the value of the ATM implied level and their differences from each other measure the smile away from the money; out-of-calibration results are in Table 41 in Appendix F.1. The development common sample contains announcements of 2021 only: DJ-hist and Kou-SEJ-P need four prior announcements of the firm and the sample starts in January 2020, so no 2020 event enters; the broad sample of Table 41 covers 2020. Week-cl.: 90% interval from a bootstrap that resamples announcement weeks (Monday-dated week of the announcement, the clusters of the Diebold–Mariano tests; 48 development and 167 holdout weeks) instead of events, a post-freeze sensitivity check, not pre-registered. Implied-volatility difference: both IVs inverted at the quote’s own S , r , q , T ; it is the metric of the recent-events study and of the other tables, and the exact inversion avoids the first-order vega approximation, which misstates errors of several vol. points away from the money, where the price is convex in volatility (BS-HV’s RMSE for $|\Delta|$ 0.15–0.35 is 10.33 vega-scaled against 11.47 (development) and 12.85 vega-scaled against 14.50 (holdout)). RMSE pooled is over quotes; bias is the event average of the mean signed IV error; RMSE and MAE in dollars, MAPE (mids of at least \$1) and the share of model prices inside the bid–ask spread are event averages. The other tables of this section use the same common sample and conventions. Models as defined in Section 8.1; the primary Kou-SEJ-Q is a deviation from the pre-registration recorded in `output/spec_freeze.json`, and Kou-SEJ-Q-prereg is the literal pre-registered variant. In 2,800 of the 6,893 common-sample events (41%) its zero-mean double-exponential family cannot match the market’s skew slope, and p is set to the value whose slope is closest. The dividend yield q of a quote sums the dividends with an ex-date in $(d_{\text{pre}}, \text{expiration}]$ in the realized dividend table, a known-dividend assumption (declaration dates are not in the data; 35,430 of the 151,969 kept quotes); the same q enters the model and the market IVs.




Figure 3: Pre-event errors across moneyness, common sample (development: announcements of 2021; holdout: 2022–2025). (a, b) Mean signed implied-volatility error, model minus market in volatility points, over the quotes in each bin of standardized moneyness $\ln(K/F)/(\sigma_{ATM}\sqrt{\tau})$, with 90% event-bootstrap intervals; the bins were fixed before any error by moneyness was computed. (c, d) Event-average RMSE by $|\Delta|$ bin with 90% intervals. Mkt-ATM is not drawn; inside the calibration expirations it coincides with DJ-implied. The outer moneyness bins are affected by the delta filter (see the text).

uncaptured contains the variance risk premium as well as model error. Relative to BS-IVlong the historical-variance models captured negative fractions (DJ-hist -0.90 and -0.34 ; Kou-SEJ-P -0.89 and -0.42), against 0.45 and 0.50 for Kou-SEJ-Q.

8.4 Where the models differ

Appendix F.2 reports the RMSE and the mispricing captured by every pre-registered cut (Tables 42–44, Figure 16); three patterns carry the argument. First, moneyness. The option-implied models fitted to the at-the-money level (Mkt-ATM, DJ-implied and Kou-SEJ-Q) had their smallest errors for $0.35 \leq |\Delta| \leq 0.65$, as they should; Kou-SEJ-Q-prereg, which is not refitted to that level, had its smallest in the wing bin (Appendix F.2). For $0.15 \leq |\Delta| < 0.35$ the double-exponential smile of Kou-SEJ-Q was worse than the flat volatility of DJ-implied (RMSE 2.42 against 2.01 in the development sample, 2.97 against 2.50 in the holdout), although its mean

Table 9: Mispricing captured, $MC = 1 - \sum |V_{\text{model}} - \text{mid}| / \sum |V_{\text{bench}} - \text{mid}|$ (ratio of sums over quotes), the reduction of the event-average RMSE relative to BS-HV, $1 - \overline{\text{RMSE}}_{\text{model}} / \overline{\text{RMSE}}_{\text{BS-HV}}$, with the implied-volatility difference and with the pre-registered vega-scaled error (the statistic of P1 in Table 12), and the share of quotes on which the model is closer to the mid than BS-HV. Brackets: 90% event-bootstrap intervals (1,000 draws); Week-cl.: the interval of MC relative to BS-HV from the bootstrap of announcement weeks of Table 8 (sensitivity check). Common sample and conventions as in Table 8.

Model	Mispricing captured relative to				RMSE reduction vs BS-HV (%)		Closer than
	BS-HV	Week-cl.	CRR-HV	BS-IVlong	IV diff. [90% CI]	vega	BS-HV (%)
<i>A. Development 2021 (common sample)</i>							
CRR-HV	0.01 [0.01, 0.01]	[0.01, 0.01]	–	–1.60 [–1.97, –1.32]	5.3 [4.3, 6.4]	3.0	46.1
BS-IVlong	0.62 [0.57, 0.67]	[0.57, 0.68]	0.62 [0.57, 0.66]	–	45.6 [44.1, 47.3]	44.8	85.7
Mkt-ATM	0.83 [0.80, 0.85]	[0.80, 0.86]	0.83 [0.80, 0.85]	0.55 [0.52, 0.57]	68.5 [66.9, 70.0]	68.3	91.8
DJ-hist	0.27 [0.20, 0.35]	[0.16, 0.37]	0.27 [0.20, 0.34]	–0.90 [–1.35, –0.61]	30.0 [27.7, 32.5]	23.7	70.0
DJ-implied	0.83 [0.80, 0.85]	[0.80, 0.86]	0.83 [0.80, 0.85]	0.55 [0.52, 0.57]	68.5 [67.0, 70.0]	68.3	91.8
Kou-SEJ-P	0.28 [0.21, 0.34]	[0.19, 0.35]	0.27 [0.21, 0.34]	–0.89 [–1.31, –0.60]	29.9 [27.8, 32.1]	25.1	75.2
Kou-SEJ-Q	0.79 [0.76, 0.82]	[0.76, 0.82]	0.79 [0.76, 0.81]	0.45 [0.42, 0.48]	67.3 [65.7, 68.8]	66.3	90.5
Kou-SEJ-Q-prereg	0.78 [0.75, 0.81]	[0.75, 0.82]	0.78 [0.75, 0.81]	0.42 [0.40, 0.45]	65.4 [63.9, 66.9]	64.6	91.2
<i>B. Holdout 2022–2025 (common sample)</i>							
CRR-HV	0.03 [0.03, 0.04]	[0.03, 0.04]	–	–1.36 [–1.46, –1.26]	7.3 [7.0, 7.7]	3.8	51.2
BS-IVlong	0.59 [0.57, 0.61]	[0.57, 0.61]	0.58 [0.56, 0.59]	–	43.5 [42.4, 44.6]	42.5	87.2
Mkt-ATM	0.82 [0.81, 0.83]	[0.81, 0.84]	0.82 [0.81, 0.83]	0.57 [0.56, 0.58]	70.8 [70.0, 71.6]	70.6	91.7
DJ-hist	0.45 [0.43, 0.47]	[0.41, 0.49]	0.43 [0.41, 0.45]	–0.34 [–0.44, –0.26]	42.9 [41.8, 43.9]	38.9	79.9
DJ-implied	0.82 [0.81, 0.83]	[0.81, 0.84]	0.82 [0.81, 0.83]	0.57 [0.56, 0.58]	70.8 [69.9, 71.6]	70.5	91.7
Kou-SEJ-P	0.42 [0.40, 0.44]	[0.38, 0.45]	0.40 [0.38, 0.42]	–0.42 [–0.52, –0.34]	39.7 [38.7, 40.7]	36.1	83.0
Kou-SEJ-Q	0.80 [0.79, 0.81]	[0.78, 0.81]	0.79 [0.78, 0.80]	0.50 [0.49, 0.52]	70.2 [69.3, 70.9]	69.1	90.7
Kou-SEJ-Q-prereg	0.74 [0.73, 0.75]	[0.72, 0.76]	0.73 [0.72, 0.75]	0.37 [0.35, 0.39]	64.9 [64.1, 65.8]	64.1	90.0
<i>C. All 2021–2025 (common sample)</i>							
CRR-HV	0.03 [0.03, 0.03]	[0.03, 0.03]	–	–1.39 [–1.50, –1.31]	7.0 [6.6, 7.4]	3.7	50.3
BS-IVlong	0.59 [0.58, 0.61]	[0.58, 0.61]	0.58 [0.57, 0.60]	–	43.8 [43.0, 44.8]	42.9	86.9
Mkt-ATM	0.82 [0.82, 0.83]	[0.81, 0.84]	0.82 [0.81, 0.83]	0.57 [0.55, 0.58]	70.4 [69.7, 71.1]	70.2	91.7
DJ-hist	0.42 [0.40, 0.44]	[0.38, 0.46]	0.40 [0.38, 0.43]	–0.43 [–0.53, –0.34]	40.7 [39.7, 41.5]	36.3	78.1
DJ-implied	0.82 [0.82, 0.83]	[0.81, 0.84]	0.82 [0.81, 0.83]	0.57 [0.55, 0.58]	70.4 [69.7, 71.1]	70.1	91.7
Kou-SEJ-P	0.40 [0.37, 0.42]	[0.36, 0.43]	0.38 [0.36, 0.40]	–0.49 [–0.59, –0.40]	38.0 [37.1, 38.8]	34.2	81.6
Kou-SEJ-Q	0.80 [0.79, 0.80]	[0.78, 0.81]	0.79 [0.78, 0.80]	0.50 [0.48, 0.51]	69.7 [69.0, 70.3]	68.6	90.7
Kou-SEJ-Q-prereg	0.75 [0.74, 0.76]	[0.73, 0.76]	0.74 [0.73, 0.75]	0.38 [0.36, 0.39]	65.0 [64.3, 65.7]	64.2	90.2

error there was smaller. The market smile was the familiar put skew of equity options, before and after the announcement (P7 in Section 8.9), and the firm-specific shape that Kou-SEJ-Q takes from a few past moves added more dispersion than it removed bias; in two of the four typical events of Figure 15 in Appendix F.2 (ALGT 2021-07-28 and MPWR 2023-10-30) the asymmetry taken from the firm’s history tilts the Kou-SEJ-Q smile against the market’s. Second, the variance that BS-HV omits was largest where little time remained and where past moves had been large: BS-HV’s error was largest for 5–15 days to expiration (27.48 and 28.45) and rose from the bottom to the top tercile of the firm’s past announcement moves, and so did the share of it captured by Kou-SEJ-Q. Third, the errors of Kou-SEJ-Q were largest in the top VIX tercile (in the development sample a small tercile, because its cut points include the high-VIX months of 2020), and in the holdout the lowest price tercile, the pre-registered size proxy, had the largest errors of every model.

8.5 Tests of equal accuracy

Table 45 in Appendix F.3 reports Diebold–Mariano tests (Diebold and Mariano, 1995) on per-event loss differentials, with standard errors clustered by announcement week, the conditional test of Giacomini and White (2006) and Mincer–Zarnowitz regressions. Under the implied-volatility loss every model beat BS-HV with $|t| > 3$, the hurdle of Harvey et al. (2016), in both samples; under the vega-scaled loss DJ-hist missed it in the development sample ($t = -2.17$). The option-implied announcement variance beat the historical one (Kou-SEJ-Q against Kou-SEJ-P: $t = -4.89$ and -2.80 ; DJ-implied against DJ-hist: -4.70 and -2.61); the holdout statistics fell

short of the hurdle, consistent with the heavy tail of the historical models' per-event losses, whose pooled RMSE is more than twice the event average (Table 8). Whether the double-exponential shape lost to the Gaussian depended on the loss. With the implied-volatility difference Kou-SEJ-Q and DJ-implied were indistinguishable over all quotes ($t = 0.99$ and -0.97), while on the quotes with $|\Delta| < 0.35$ Kou-SEJ-Q had the larger loss ($t = 4.47$ and 8.68). With the pre-registered vega-scaled loss it had the larger loss over all quotes as well ($t = 4.07$ and 6.08): where the price is convex in volatility, dividing a price error by vega enlarges over-pricing errors and shrinks under-pricing ones, and the steeper smile of Kou-SEJ-Q over-priced part of the wings. No prediction rests on this test. On the out-of-calibration quotes Kou-SEJ-Q had the smaller loss, not significantly ($t = -1.14$, 50 events). The Mincer–Zarnowitz slopes of the option-implied models are close to one, largely by construction, since they are fitted to the at-the-money level of the same expirations; those of BS-HV and the historical-variance models are well below. The 90% Model Confidence Set (10% level) of Hansen et al. (2011) held only models that use the at-the-money quotes, always including Mkt-ATM, and with the vega-scaled loss Kou-SEJ-Q left it (Table 46 in Appendix F.3). Regressions of the loss differentials on the pre-registered characteristics locate the disadvantage of Kou-SEJ-Q relative to DJ-implied in the top tercile of the firm's past moves, where the shape taken from history hurt most, under every loss and every kind of standard error (holdout: two-way clustered $t = 4.97$ with the reported loss, wild cluster bootstrap p -value < 0.001 ; Tables 47 and 48 in Appendix F.3). A top-VIX-tercile coefficient appeared only with the squared implied-volatility loss and not with the per-event RMSE or the pre-registered vega-scaled loss, so I do not read the VIX terciles as a location of the disadvantage.

8.6 After the announcement

At d_{post} every model was priced with its background diffusion only, on the contracts quoted at d_{pre} whose post-event quote passed the filters (932 development and 3,986 holdout events; Table 49 in Appendix F.4). P1 predicted that the event-jump models would then lose their advantage over BS-HV, which presumes a common background, and the models do not share one. DJ-hist and Kou-SEJ-P use σ_{HV} , so their post-event prices are BS-HV's by construction. DJ-implied, Kou-SEJ-Q and Kou-SEJ-Q-prereg use option-implied backgrounds, which were closer to the post-event market: in implied-volatility differences the RMSE of DJ-implied was 33% and 22% below BS-HV's, that of Kou-SEJ-Q 29% and 17% (vega-scaled, the statistic of P1: 34.1% and 25.3%, and 30.6% and 20.6%). The post-event comparison therefore measures the choice of background volatility, not the announcement model. The post-event sample is also selected on the move: after a large move the strikes quoted before the announcement leave the delta window of the filters, so the events with a post-event quote (67% of the development and 72% of the holdout common sample) had smaller moves relative to their option-implied announcement s.d. (median $|x|/\sigma_J^Q$ of 0.60 against 0.72 for the other events, holdout). Without the post-event filters, on every post-linked contract with a post-event implied volatility (a sensitivity check, not pre-registered), the vega-scaled reductions were 30.2% [13.3, 43.0] and 21.0% [12.6, 30.1] for Kou-SEJ-Q, still far outside the predicted $\pm 5\%$. Across the announcement, DJ-implied and Kou-SEJ-Q predicted larger falls of the at-the-money implied volatility than the market's, because their background was below the post-event market level (Table 50 in Appendix F.4), with or without the post-event filters.

8.7 The early-exercise premium

Table 10 reports the early-exercise premium of the American Kou-SEJ-Q on its subsample (142 development and 545 holdout events) and of CRR-HV on all quotes. By Theorem 3.3 a put is not exercised at the instant before the announcement; its premium comes from exercise after the jump has moved it into the money and grows with the interest rate. In the development period the median 1-month rate was 0.07% and the put premium negligible (median \$0.000). In the

Table 10: Early exercise. Panel A: distribution of the early-exercise premium of the American Kou-SEJ-Q (finite differences, stratified 10% subsample of events) and of CRR-HV (binomial control variate, all quotes), for puts by $|\Delta|$ bin and for calls whose life holds an ex-dividend date; median (%) is relative to the European price; >HS: share of quotes whose premium exceeds the half-spread (ask – bid)/2 of the quote, and Med./HS the median ratio of the premium to the half-spread (post-freeze additions, descriptive). Panel B: event-average RMSE in implied-volatility points under the European and the American model (second and third column of each block) and the reduction from using the American price, with 90% event-bootstrap intervals; Kou-SEJ-Q-Am on the subsample with its own price, CRR-HV on the common sample. Market IVs are European inversions of American quotes; Kou-SEJ-Q is calibrated to such IVs at the money.

Method	Quotes	Development 2021						Holdout 2022–2025					
		n	Median (\$)	P90 (\$)	Median (%)	>HS (%)	Med./HS	n	Median (\$)	P90 (\$)	Median (%)	>HS (%)	Med./HS
<i>A. Early-exercise premium (American minus European price)</i>													
FD Kou-SEJ-Q	puts	1,012	0.000	0.001	0.00	0.0	0.00	4,664	0.026	0.342	0.71	2.7	0.14
	puts, $ \Delta $ 0.15–0.35	178	0.000	0.000	0.00	0.0	0.00	1,044	0.008	0.054	0.56	0.7	0.08
CRR-HV	puts, $ \Delta $ 0.35–0.65	426	0.000	0.001	0.00	0.0	0.00	1,956	0.027	0.214	0.68	2.8	0.14
	puts, $ \Delta $ 0.65–0.85	408	0.000	0.002	0.00	0.0	0.00	1,664	0.074	0.696	0.87	3.7	0.18
	calls with a dividend	149	0.049	0.383	1.11	8.1	0.17	1,278	0.027	0.304	0.70	12.4	0.14
	puts	9,751	0.000	0.001	0.00	0.0	0.00	44,106	0.013	0.260	0.46	1.1	0.07
	puts, $ \Delta $ 0.15–0.35	1,803	0.000	0.000	0.00	0.0	0.00	9,794	0.002	0.016	0.30	0.0	0.02
	puts, $ \Delta $ 0.35–0.65	4,058	0.000	0.000	0.00	0.0	0.00	18,093	0.014	0.112	0.44	0.4	0.07
	puts, $ \Delta $ 0.65–0.85	3,890	0.000	0.002	0.00	0.0	0.00	16,219	0.067	0.601	0.72	2.6	0.15
	calls with a dividend	2,122	0.049	0.492	1.15	12.7	0.18	11,747	0.011	0.295	0.41	8.4	0.06
<i>B. Errors of the European and the American model: events, RMSE in IV points (event average) European, American, and the reduction (%) with 90% interval</i>													
Kou-SEJ-Q-Am vs Kou-SEJ-Q	all quotes	142	3.84	3.52		8.4 [4.1, 12.9]		545	4.77	4.22		11.7 [8.3, 15.1]	
	calls with a dividend		6.64	4.19		36.9 [28.0, 44.4]			8.34	5.76		30.9 [25.8, 35.5]	
CRR-HV vs BS-HV	all puts		3.11	3.10		0.0 [0.0, 0.1]			3.56	3.65		−2.5 [−4.8, 0.1]	
	puts, $ \Delta $ 0.15–0.35		2.13	2.13		0.0 [0.0, 0.0]			2.86	2.79		2.5 [2.0, 3.1]	
	puts, $ \Delta $ 0.35–0.65		1.39	1.39		0.0 [0.0, 0.0]			1.74	1.91		−9.7 [−11.7, −8.0]	
	puts, $ \Delta $ 0.65–0.85		3.98	3.98		0.0 [0.0, 0.1]			4.91	5.10		−3.8 [−6.9, −0.6]	
	all quotes	1,386	12.84	12.16		5.3 [4.3, 6.4]		5,507	15.59	14.45		7.3 [7.0, 7.7]	
	calls with a dividend		15.38	11.37		26.1 [23.0, 29.2]			16.47	13.43		18.4 [17.3, 19.5]	
	all puts		11.27	11.25		0.2 [0.1, 0.3]			13.94	12.97		7.0 [6.7, 7.3]	
	puts, $ \Delta $ 0.15–0.35		11.80	11.80		0.0 [0.0, 0.0]			15.36	15.32		0.2 [0.2, 0.3]	
	puts, $ \Delta $ 0.35–0.65		10.27	10.26		0.0 [0.0, 0.0]			12.96	12.78		1.4 [1.3, 1.4]	
	puts, $ \Delta $ 0.65–0.85		11.34	11.30		0.3 [0.2, 0.5]			13.87	11.58		16.5 [15.9, 17.1]	

holdout, at a median rate of 4.67%, the median put premium was \$0.026 (0.71% of the European price), its 90th percentile \$0.342, and the median for in-the-money puts \$0.074. Against the cost of trading the option the premium was small: it exceeded the quote’s half-spread for none of the development puts, for 2.7% of the holdout puts (3.7% of those in the money; median premium 0.14 of the half-spread) and for 8.1% and 12.4% of the calls with a dividend.

The American price helped for calls whose life holds an ex-dividend date, cutting the RMSE of Kou-SEJ-Q on them by 36.9% [28.0, 44.4] and 30.9% [25.8, 35.5]. The binomial CRR-HV gained on the same kind of calls over BS-HV (26.1% and 18.4%), and on the American subsample the finite-difference and binomial premia of these calls had a correlation of 0.96 (development) and 0.96 (holdout). The gain therefore comes from American exercise under the continuous dividend yield $q = D/(ST)$ of Section 7.3, which both methods share, rather than from the finite-difference solver, and that yield is an approximation to a discrete payment (Appendix F.5). For puts it did not help: in the holdout the American Kou-SEJ-Q had the higher RMSE on puts (–2.5% [–4.8, 0.1]), with the largest relative increase in the at-the-money bin, and its mean put error rose from 0.35 to 1.15 points. The European Kou-SEJ-Q is calibrated to European-inverted implied volatilities of American at-the-money quotes, which already contain the premium, so adding the finite-difference premium counts it twice at the money. This is a property of the frozen specification, which I report rather than change. A post-freeze exploratory analysis, not pre-registered (Appendix F.5, Table 51), recalibrates the American price to the American quotes. It confirms the double count at the calibration strikes of expirations without a dividend, where the mean put error of the frozen American model, 0.30, falls to 0.09, but it removes only 25% [23, 28] of the rise of the mean error over all puts, a rise that is largest in the in-the-money bin; calibrated consistently, the American price still did not help for puts, while it still helped for calls with a dividend. The double count at the money therefore explains only part of the put result. CRR-HV’s gain for in-the-money puts in the holdout (16.5%) reflected mainly that

BS-HV under-priced everything. On the American subsample, with Kou-SEJ-Q-Am among the candidate models, the 90% MCS was {Kou-SEJ-Q-Am} in the holdout (545 events), on the strength of the dividend-paying calls.

8.8 The earnings variance premium, and the historical against the option-implied law

Panel A of Table 11 reports the ratio (49) of the mean option-implied announcement variance to the mean diffusion-corrected squared realized move over all primary-sample announcements with a decomposition, the numbers of Table 6: 1.70 [1.60, 1.82] in the development period (2020 included) and 1.14 [1.09, 1.19] in the holdout. In the development period it rose with the VIX, from 1.14 in the bottom tercile to 2.16 in the top (Table 6; difference 1.02 [0.77, 1.29]); in the holdout the difference was 0.07 [−0.06, 0.22], an interval that contains zero. These ratios count maturities in calendar time and assume a flat background term structure. On the trading clock they were 1.52 and 1.06 (Panel A; holdout interval with announcement-week clusters [0.99, 1.13]), and a placebo decomposition over the same expirations after the announcement averaged 0.48 and 0.37 of the realized variance and rose with the VIX (Section 6.7, Appendix C.11): the holdout ratio is within clock and term-structure uncertainty of one, and part of the development VIX gradient is a term-structure effect. Under a flat background term structure, a ratio above one means that the risk-neutral announcement law had more variance than the realized moves, a variance risk premium in the sense of Dubinsky and Johannes (2004): a statement about two laws, not about mispricing. On the priced events of the holdout the ratio was only 1.05 [0.99, 1.11], an interval that contains one, which is closer to the finding of Gao et al. (2018) that at-the-money straddles bought shortly before announcements earned positive returns on average.

Panel B compares the two laws on the objects each should be good at. On pre-event mids Kou-SEJ-Q reduced the RMSE of Kou-SEJ-P by 53.3% and 50.6%, as a law fitted to the same chain should. On realized values the historical law was at best slightly better. On the settlement payoffs of the earnings-week weeklies, the pre-registered secondary sample, Kou-SEJ-P had the lower RMSE in the holdout (4.59 against 4.67 percent of the spot over 216 events, a reduction of 1.8% [0.7, 3.0]; the development sample has 19 such events), but the test had little power: payoffs were dominated by the realized move, and BS-HV, which under-priced on average, had the lowest payoff RMSE of all (4.57). On the post-event mid of the same contract, a secondary comparison that I added at the analysis stage, with the verdict rules and after the headline RMSEs had been displayed, and that is not counted because P8 names payoffs, the historical law was worse in both samples: Kou-SEJ-Q was closer (−2.9% and −0.7% from using Kou-SEJ-P instead). The two laws' expected absolute moves were about equally far from the realized absolute move (mean absolute errors 3.90 and 3.87 percentage points in the holdout; means 4.52%, 5.96% and 6.17% for the historical law, the option-implied law and the realized move, which also contains the ordinary trading day between the snapshots).

8.9 The pre-registered predictions

Tables 12 and 13 evaluate P1–P8 with the pre-registered rule: a component fails if its point estimate has the opposite sign or if its 90% interval excludes the predicted magnitude by more than its own width; it holds if the point estimate satisfies the prediction and is inconclusive otherwise. Verdicts use the pre-registered metric and the pre-registered event bootstrap; the column 'IV diff.' repeats them with the implied-volatility difference, and the column 'Week-cl.' gives the interval with announcement weeks resampled and, where it differs, the verdict the rule would give with it. With the implied-volatility difference every overall verdict is the same, and 4 development and 2 holdout components change (in italics in Table 12): the development P1 and P2 past-move components of DJ-hist, and the P2 wing components of DJ-hist and Kou-SEJ-P in both samples. With week-clustered intervals one development and one holdout

Table 11: The earnings variance premium and the comparison of the risk-neutral and historical announcement laws. Panel A: ratio of the mean option-implied announcement variance (two-expiry decomposition) to the mean diffusion-corrected squared realized move $z^2 = x^2 - \sigma_{HV}^2 d/252$ over primary-sample events (the definition of Table 6, which also gives the ratio by VIX tercile); under a flat background term structure a ratio above one is a variance risk premium, a statement about the two laws, not about mispricing. [†]Trading clock: the same decomposition with maturities counted in NYSE trading days (Appendix C.11), a post-freeze exploratory addition. Panel B: fit to pre-event mids (common sample); RMSE of the model price against the realized value of the contract, the discounted settlement payoff of earnings-week weeklies (secondary sample, priced from the frozen event parameters) or the post-event mid of the same contract; mean absolute error of the model’s expected absolute announcement move against the realized absolute move. Panel A covers the primary-sample announcements of each sample period; Panel B is built on the common sample of Table 8, so in the development column Panel A covers 2020–2021 and Panel B 2021 (weeklies: 2021). The pre-event RMSEs of Panel B are in IV-difference points; the P8 verdict uses the vega-scaled error (Table 13). Brackets: 90% event-bootstrap intervals; week-clustered interval: the same statistic with announcement weeks resampled instead of events (sensitivity check, not pre-registered).

	Development			Holdout			All		
	<i>n</i>	Estimate	[90% CI]	<i>n</i>	Estimate	[90% CI]	<i>n</i>	Estimate	[90% CI]
<i>A. Earnings variance premium $\sigma_J^{2,Q}/z^2$, primary-sample announcements of 2020–2021 / 2022–2025 / 2020–2025</i>									
All primary-sample events	8,226	1.70	[1.60, 1.82]	13,009	1.14	[1.09, 1.19]	21,235	1.31	[1.27, 1.36]
week-clustered interval			[1.50, 1.94]			[1.08, 1.22]			[1.23, 1.42]
High minus low VIX tercile		1.02	[0.77, 1.29]		0.07	[−0.06, 0.22]		0.48	[0.35, 0.62]
week-clustered interval			[0.55, 1.45]			[−0.14, 0.34]			[0.22, 0.79]
Events of the priced sample	3,823	1.56	[1.44, 1.70]	6,103	1.05	[0.99, 1.11]	9,926	1.19	[1.14, 1.25]
Trading clock [†] , all primary-sample events	8,226	1.52	[1.42, 1.63]	13,009	1.06	[1.01, 1.11]	21,235	1.20	[1.15, 1.24]
high minus low VIX tercile		0.88	[0.64, 1.13]		0.02	[−0.10, 0.17]		0.39	[0.27, 0.52]
events of the priced sample	3,823	1.39	[1.28, 1.52]	6,103	0.97	[0.92, 1.03]	9,926	1.09	[1.04, 1.14]
<i>B. Risk-neutral (Q) against historical (P) announcement law, common sample (2021 / 2022–2025 / 2021–2025) and its subsets</i>									
Pre-event RMSE, Kou-SEJ-P (IV pts)	1,386	9.00	[8.60, 9.41]	5,507	9.40	[9.12, 9.73]	6,893	9.32	[9.09, 9.59]
Pre-event RMSE, Kou-SEJ-Q (IV pts)		4.20	[4.00, 4.43]		4.65	[4.55, 4.76]		4.56	[4.47, 4.65]
reduction Q vs P (%)		53.3	[50.9, 55.5]		50.6	[49.0, 52.1]		51.1	[49.8, 52.4]
Weekly payoff RMSE, Kou-SEJ-P (% of S)	19	3.80	[2.84, 4.85]	216	4.59	[4.20, 4.98]	235	4.52	[4.16, 4.89]
Weekly payoff RMSE, Kou-SEJ-Q (% of S)		3.91	[2.99, 5.01]		4.67	[4.31, 5.06]		4.61	[4.27, 4.96]
reduction P vs Q (%)		3.0	[−0.8, 7.5]		1.8	[0.7, 3.0]		1.9	[0.9, 3.1]
Post-event mid RMSE, Kou-SEJ-P (% of S)	1,370	3.43	[3.31, 3.57]	5,472	4.01	[3.94, 4.09]	6,842	3.90	[3.83, 3.96]
Post-event mid RMSE, Kou-SEJ-Q (% of S)		3.34	[3.21, 3.47]		3.99	[3.92, 4.06]		3.86	[3.79, 3.91]
reduction P vs Q (%)		−2.9	[−3.7, −2.1]		−0.7	[−1.3, −0.2]		−1.1	[−1.7, −0.7]
MAE of $\mathbb{E}[Z]$ vs $ x $, Kou-SEJ-P (%)	1,386	3.41	[3.26, 3.56]	5,507	3.90	[3.80, 3.99]	6,893	3.80	[3.72, 3.88]
MAE of $\mathbb{E}[Z]$ vs $ x $, Kou-SEJ-Q (%)		3.33	[3.20, 3.47]		3.87	[3.79, 3.96]		3.76	[3.69, 3.84]

component change (P1 and P6 below), and no prediction’s verdict changes; with firms resampled, no component verdict changes.

P1: fails (development), fails (holdout). The pre-event half held for the option-implied models (Kou-SEJ-Q cut BS-HV’s RMSE by 66.26% and 69.08%) and for the historical ones in the holdout (DJ-hist 38.94%, Kou-SEJ-P 36.13%). In the development sample the historical models fell short of 30% (DJ-hist 23.69% [20.91, 26.52], Kou-SEJ-P 25.08% [22.77, 27.49]) by less than the width of their intervals, so these components are inconclusive; with the implied-volatility difference the DJ-hist component held by 0.018 percentage points. The post-event half failed for DJ-implied and Kou-SEJ-Q (Kou-SEJ-Q 30.63% and 20.57%, against $\pm 5\%$ predicted) for the reason of Section 8.6, and held by construction for DJ-hist and Kou-SEJ-P. With week-clustered intervals the holdout post-event component of Kou-SEJ-Q would be inconclusive ([15.10, 26.62] against [16.37, 24.49] with events resampled), while that of DJ-implied would still fail ([20.33, 30.90]), so the verdict would not change.

Table 12: Pre-registered predictions P1–P4 (P5–P8 in Table 13; `output/preregistration.md`, written before any pricing error was computed), development and holdout separately. Estimates and verdicts use the pre-registered primary metric, the price error divided by the Black–Scholes vega at the quote’s own IV (post-event: at d_{post}); IV diff. repeats the point estimate and the verdict with the exact implied-volatility difference of the other tables, and Week-cl. the interval with announcement weeks resampled, two sensitivity checks (in italics where the verdict differs). Bold: overall verdict of the prediction. Pre-registered failure rule: a component fails if the point estimate has the opposite sign or the 90% event-bootstrap interval excludes the predicted magnitude by more than its own width. *Literal pre-registered variant Kou-SEJ-Q-prereg in place of the primary Kou-SEJ-Q (overall verdict in parentheses when it differs); †reported, not counted; inconcl.: inconclusive; too few: fewer than 30 events, not judged. Notes below Table 13.

Prediction and component	Pred.	Development 2021				Holdout 2022–2025			
		Estimate [90% CI]	Week-cl.	Verdict	IV diff.	Estimate [90% CI]	Week-cl.	Verdict	IV diff.
P1 Event-jump models cut the pre-event RMSE vs BS-HV by $\geq 30\%$; at d_{post} the reduction is within $\pm 5\%$			fails	fails	fails		fails	fails	fails
pre: DJ-hist	≥ 30	23.69 [20.91, 26.52]	[14.22, 31.08]	inconcl.	<i>30.02 holds</i>	38.94 [37.83, 40.14]	[35.34, 42.09]	holds	42.85 holds
pre: DJ-impl.	≥ 30	68.28 [66.81, 69.69]	[66.06, 70.48]	holds	68.51 holds	70.52 [69.69, 71.37]	[68.54, 72.31]	holds	70.78 holds
pre: Kou-SEJ-P	≥ 30	25.08 [22.77, 27.49]	[18.51, 30.03]	inconcl.	29.94 inconcl.	36.13 [35.07, 37.15]	[33.39, 38.61]	holds	39.69 holds
pre: Kou-SEJ-Q	≥ 30	66.26 [64.80, 67.82]	[64.00, 68.48]	holds	67.27 holds	69.08 [68.20, 69.93]	[67.14, 70.80]	holds	70.18 holds
pre: Q-prereg*	≥ 30	64.61 [63.19, 66.12]	[62.46, 66.93]	holds	65.39 holds	64.15 [63.25, 65.07]	[62.20, 65.99]	holds	64.92 holds
post: DJ-hist	± 5	0.00 [0.00, 0.00]	[0.00, 0.00]	holds	0.00 holds	0.00 [0.00, 0.00]	[0.00, 0.00]	holds	0.00 holds
post: DJ-impl.	± 5	34.08 [30.44, 37.76]	[27.64, 40.84]	fails	32.86 fails	25.35 [21.40, 29.07]	[20.33, 30.90]	fails	22.42 fails
post: Kou-SEJ-P	± 5	0.00 [0.00, 0.00]	[0.00, 0.00]	holds	0.00 holds	0.00 [0.00, 0.00]	[0.00, 0.00]	holds	0.00 holds
post: Kou-SEJ-Q	± 5	30.63 [26.72, 34.41]	[23.55, 37.96]	fails	29.04 fails	20.57 [16.37, 24.49]	[15.10, 26.62]	<i>inconcl.</i>	17.28 fails
post: Q-prereg*	± 5	34.08 [30.44, 37.76]	[27.64, 40.84]	fails	32.86 fails	25.35 [21.40, 29.07]	[20.33, 30.90]	fails	22.42 fails
P2 The reduction is larger for $ \Delta $ 0.15–0.35 than at the money, and in the top past-move tercile			fails	fails	fails		fails	fails	fails
wing – ATM: DJ-hist	> 0	–4.4 [–6.9, –2.1]	[–7.7, –1.0]	fails	<i>2.8 holds</i>	–4.5 [–5.2, –3.8]	[–5.2, –3.7]	fails	<i>0.9 holds</i>
wing – ATM: DJ-impl.	> 0	–6.3 [–7.2, –5.3]	[–7.3, –5.3]	fails	–4.9 fails	–7.1 [–7.6, –6.7]	[–8.0, –6.3]	fails	–5.4 fails
wing – ATM: Kou-SEJ-P	> 0	–4.6 [–6.6, –2.6]	[–7.3, –1.8]	fails	<i>1.0 holds</i>	–4.3 [–4.9, –3.6]	[–4.9, –3.6]	fails	<i>0.6 holds</i>
wing – ATM: Kou-SEJ-Q	> 0	–9.2 [–10.4, –8.0]	[–10.5, –8.0]	fails	–7.2 fails	–9.8 [–10.4, –9.2]	[–10.5, –9.1]	fails	–7.4 fails
wing – ATM: Q-prereg*	> 0	0.8 [–0.1, 1.6]	[–0.3, 1.8]	holds	2.3 holds	2.4 [1.9, 2.8]	[1.7, 3.0]	holds	4.3 holds
past-move T3 – T1: DJ-hist	> 0	–3.2 [–9.7, 3.6]	[–15.7, 6.7]	fails	<i>4.2 holds</i>	24.6 [22.3, 27.0]	[21.9, 27.3]	holds	27.8 holds
past-move T3 – T1: DJ-impl.	> 0	16.9 [13.3, 20.8]	[13.1, 21.3]	holds	18.4 holds	25.5 [23.4, 27.5]	[22.7, 28.3]	holds	28.1 holds
past-move T3 – T1: Kou-SEJ-P	> 0	12.2 [7.0, 17.6]	[1.6, 21.0]	holds	18.4 holds	28.7 [26.8, 30.6]	[26.5, 30.8]	holds	31.9 holds
past-move T3 – T1: Kou-SEJ-Q	> 0	16.4 [12.6, 20.4]	[13.0, 20.6]	holds	17.6 holds	23.0 [20.9, 25.0]	[20.3, 25.7]	holds	25.4 holds
past-move T3 – T1: Q-prereg*	> 0	14.3 [10.6, 18.1]	[10.3, 18.8]	holds	15.7 holds	20.9 [18.7, 23.1]	[18.3, 23.6]	holds	23.8 holds
P3 No model has a lower ATM RMSE than the same-expiry market IV (Mkt-ATM)			fails	fails	fails		holds	holds	holds
min over models: ATM RMSE – Mkt-ATM (lowest: DJ-impl.)	≥ 0	–0.0003 [–0.0021, 0.0014]	[–0.0028, 0.0015]	fails	–0.0003 fails	0.0088 [0.0050, 0.0136]	[0.0033, 0.0156]	holds	0.0088 holds
the same, models other than DJ-impl. [†] (lowest: Kou-SEJ-Q)	≥ 0	0.1365 [0.1136, 0.1608]	[0.1180, 0.1539]	holds	0.1364 holds	0.1737 [0.1587, 0.1890]	[0.1531, 0.1973]	holds	0.1726 holds
P4 Residual errors are largest in the top VIX and top relative-spread terciles			fails	fails (lit. holds)	fails (lit. holds)		fails	fails (lit. holds)	fails (lit. holds)
Kou-SEJ-Q: VIX T3 – max(T1, T2)	> 0	1.80 [0.64, 2.94]	[1.21, 3.44]	holds	1.84 holds	0.81 [0.57, 1.05]	[0.49, 1.15]	holds	0.75 holds
Kou-SEJ-Q: spread T3 – max(T1, T2)	> 0	–0.53 [–1.65, 0.58]	[–1.64, 0.76]	fails	–0.33 fails	–0.002 [–0.422, 0.421]	[–0.435, 0.470]	fails	–0.14 fails
Q-prereg: VIX T3 – max(T1, T2)*	> 0	1.57 [0.68, 2.48]	[0.92, 3.94]	holds	1.74 holds	0.99 [0.73, 1.24]	[0.61, 1.40]	holds	0.90 holds
Q-prereg: spread T3 – max(T1, T2)*	> 0	0.03 [–1.09, 1.13]	[–0.97, 1.26]	holds	0.15 holds	0.30 [–0.10, 0.72]	[–0.11, 0.73]	holds	0.12 holds

P2: fails, fails. DJ-hist, DJ-implied, Kou-SEJ-P and Kou-SEJ-Q reduced the error by less in the wings than at the money (wing-minus-ATM difference of the reductions for Kou-SEJ-Q: -9.2 and -9.8 percentage points). The wing component held only for the literal variant Kou-SEJ-Q-prereg (0.8 $[-0.1, 1.6]$ and 2.4 $[1.9, 2.8]$), and as in P6 this came from its worse at-the-money fit, not from better wings: its wing RMSE was close to that of Kou-SEJ-Q (2.37 against 2.42 in the development sample, 3.12 against 2.97 in the holdout), its at-the-money RMSE much larger (2.47 against 1.51 , and 3.46 against 1.76). With the implied-volatility difference DJ-hist and Kou-SEJ-P had positive differences (DJ-hist 2.8 and 0.9), because under that metric BS-HV’s wing errors are larger relative to its at-the-money errors (ratio 1.06 against 0.97 vega-scaled, development). The past-move half held in the holdout for all four models and in the development sample for all but DJ-hist (-3.2 $[-9.7, 3.6]$).

P3: fails, holds. The minimum over models of the at-the-money RMSE minus that of Mkt-ATM was attained by DJ-implied, which coincides with Mkt-ATM inside the calibration expirations, so the difference came from the out-of-calibration quotes alone: -0.0003 $[-0.0021,$

Table 13: Pre-registered predictions P5–P8, continued from Table 12 (same conventions; same: the component uses no pricing error, so both metrics give it).

Prediction and component	Pred.	Development 2021				Holdout 2022–2025			
		Estimate [90% CI]	Week-cl.	Verdict	IV diff.	Estimate [90% CI]	Week-cl.	Verdict	IV diff.
P5 Option-implied announcement variance beats historical (DJ-implied, Kou-SEJ-Q vs DJ-hist, Kou-SEJ-P)			holds	holds	holds		holds	holds	holds
DJ-impl. vs DJ-hist	> 0	58.4 [56.2, 60.6]	[54.7, 62.5]	holds	55.0 holds	51.7 [50.0, 53.6]	[49.7, 53.7]	holds	48.9 holds
Kou-SEJ-Q vs Kou-SEJ-P	> 0	55.0 [52.7, 57.2]	[52.1, 58.2]	holds	53.3 holds	51.6 [49.9, 53.3]	[49.6, 53.5]	holds	50.6 holds
Q-prereg vs Kou-SEJ-P*	> 0	52.8 [50.4, 55.1]	[49.7, 56.5]	holds	50.6 holds	43.9 [41.9, 45.8]	[41.7, 46.1]	holds	41.8 holds
P6 The DE shape (Kou-SEJ-Q vs DJ-implied) helps more in the wings; <5% at the money gain vs DJ-impl., wing – ATM: Kou-SEJ-Q			fails	fails	fails		fails	fails	fails
gain vs DJ-impl., wing – ATM: Kou-SEJ-Q	> 0	–11.8 [–16.7, –7.3]	[–16.4, –8.0]	fails	–10.5 fails	–10.1 [–11.8, –8.4]	[–13.3, –7.2]	fails	–8.6 fails
gain vs DJ-impl. at the money: Kou-SEJ-Q	±5	–10.0 [–11.8, –8.2]	[–11.1, –8.7] <i>fails</i>	inconcl.	–10.0 inconcl.	–10.4 [–11.4, –9.4]	[–11.8, –9.0]	fails	–10.3 fails
gain vs DJ-impl., wing – ATM: Q-prereg*	> 0	63.3 [56.1, 69.9]	[55.0, 70.3]	holds	63.0 holds	92.8 [88.8, 96.8]	[86.8, 98.3]	holds	92.5 holds
gain vs DJ-impl. at the money: Q-prereg*	±5	–80.2 [–87.3, –73.4]	[–89.5, –69.5]	fails	–80.6 fails	–116.7 [–121.4, –112.4]	[–122.8, –109.9]	fails	–117.3 fails
P7 BS-HV under-prices in every $ \Delta $ bin; DJ-implied under-prices the wings relative to the money			holds	holds	holds		holds	holds	holds
BS-HV bias, wing bin	< 0	–8.23 [–8.86, –7.63]	[–9.75, –6.83]	holds	–9.44 holds	–10.58 [–11.02, –10.10]	[–11.65, –9.42]	holds	–12.34 holds
BS-HV bias, ATM bin	< 0	–8.54 [–9.05, –8.05]	[–10.17, –7.05]	holds	–8.69 holds	–10.98 [–11.34, –10.62]	[–12.16, –9.73]	holds	–11.18 holds
BS-HV bias, deep bin	< 0	–9.41 [–9.87, –8.95]	[–10.78, –8.17]	holds	–11.70 holds	–11.61 [–11.96, –11.24]	[–12.61, –10.54]	holds	–14.57 holds
DJ-impl. bias, wing – ATM	< 0	–0.84 [–0.94, –0.73]	[–0.96, –0.73]	holds	–0.87 holds	–0.83 [–0.88, –0.78]	[–0.94, –0.72]	holds	–0.87 holds
P8 Implied announcement variance exceeds realized, more so in high VIX; Q fits mids, P forecasts payoffs			inconcl.	inconcl.	inconcl.		holds	holds	holds
premium ratio – 1	> 0	0.70 [0.60, 0.82]	[0.50, 0.94]	holds	same	0.14 [0.09, 0.19]	[0.08, 0.22]	holds	same
premium ratio, VIX T3 –	> 0	1.02 [0.77, 1.29]	[0.55, 1.45]	holds	same	0.07 [–0.06, 0.22]	[–0.14, 0.34]	holds	same
T1 premium ratio – 1, trading clock [†]	> 0	0.52 [0.42, 0.63]	[0.34, 0.72]	holds	same	0.06 [0.01, 0.11]	[–0.01, 0.13]	holds	same
premium ratio, VIX T3 –	> 0	0.88 [0.64, 1.13]	[0.46, 1.26]	holds	same	0.02 [–0.10, 0.17]	[–0.16, 0.28]	holds	same
T1 , trading clock [†]									
mids: Kou-SEJ-Q vs	> 0	55.0 [52.7, 57.2]	[52.1, 58.2]	holds	53.3 holds	51.6 [49.9, 53.3]	[49.6, 53.5]	holds	50.6 holds
Kou-SEJ-P									
mids: Q-prereg vs	> 0	52.8 [50.4, 55.1]	[49.7, 56.5]	holds	50.6 holds	43.9 [41.9, 45.8]	[41.7, 46.1]	holds	41.8 holds
Kou-SEJ-P*									
weekly payoffs: Kou-SEJ-P vs Kou-SEJ-Q	> 0	3.0 [–0.8, 7.5]	[1.1, 8.5]	too few	same	1.8 [0.7, 3.0]	[0.7, 3.2]	holds	same
weekly payoffs: Kou-SEJ-P vs Q-prereg	> 0	–0.5 [–4.0, 2.6]	[–3.7, 2.8]	too few	same	0.3 [–0.7, 1.2]	[–0.8, 1.3]	holds	same
post-event mids: Kou-SEJ-P vs Kou-SEJ-Q [†]	> 0	–2.9 [–3.7, –2.1]	[–4.3, –1.7]	fails	same	–0.7 [–1.3, –0.2]	[–1.6, 0.0]	fails	same

Notes to Tables 12 and 13. Reductions and gains in percent (pp: percentage points), errors and their differences in vol. points, premium ratios unit-free. Operationalization of the failure rule: a component holds if the point estimate satisfies the prediction and is otherwise inconclusive; a prediction holds if every counted component holds and fails if any counted component fails; components with fewer than 30 events are not judged. These rules and the statistic of each prediction were coded after the headline RMSEs of both samples had been displayed; afterwards no model or setting of the specification changed, the re-pricing after the data correction (deviation (viii) of Section 8.1) applied the frozen code to corrected inputs, and the re-pricings after later edits of non-numerical code reproduced every frozen price. The two P8 rows on the trading clock use the announcement variance of the same decomposition with maturities counted in trading days (Appendix C.11), a post-freeze exploratory addition that is reported and not counted; no verdict depends on them. Week-cl.: 90% interval from a bootstrap that resamples announcement weeks (Monday-dated week of the announcement, the clusters of the Diebold–Mariano tests; 48 development and 167 holdout weeks) instead of events, a post-freeze sensitivity check, not pre-registered. Its column gives the interval and, in italics, the verdict the failure rule gives with it where that differs (in the row of a prediction, the overall verdict); every reported verdict uses the pre-registered event bootstrap. P3: DJ-implied equals Mkt-ATM inside the calibration expirations, so its difference rests on the out-of-calibration at-the-money quotes of 4 (development) and 48 (holdout) events; the minimum of 30 events counts the events of the component's sample, every event with at-the-money quotes. The primary Kou-SEJ-Q (shape from the firm's history, ATM calibration) is a deviation from the pre-registration recorded in `output/spec_freeze.json`. In 2,800 of the 6,893 common-sample events (41%; development 545 of 1,386, holdout 2,255 of 5,507) the zero-mean double-exponential family of Kou-SEJ-Q-prereg cannot match the market's skew slope, and p is set to the value whose slope is closest. P4 is judged on the residuals of the primary Kou-SEJ-Q (Kou-SEJ-Q-prereg for the literal variant); the verdict depends on that choice: the spread component is positive in both samples for BS-HV, CRR-HV, DJ-hist, Kou-SEJ-P, Kou-SEJ-Q-prereg, negative in both for Kou-SEJ-Q and of mixed sign for BS-IVlong, Mkt-ATM, DJ-implied (all nine models in `output/results_summary.json`). Common sample and conventions as in Table 8.

0.0014] volatility points in the development sample (4 events with such at-the-money quotes) and 0.0088 [0.0050, 0.0136] in the holdout (48 events). Under the rule a negative point estimate, however small, is the opposite sign, so P3 fails in the development sample. The statistic is an average over all development events with at-the-money quotes, so the minimum of 30 events of deviation (v) does not exclude it, although only these 4 events can move it (2 of their differences are negative and 2 positive); I report the verdict the rule gives. Every other model was worse than Mkt-ATM at the money (the closest, Kou-SEJ-Q, by 0.1365 and 0.1737). For the models that use no option price P3 is an efficiency statement; for DJ-implied it is close to an identity.

P4: fails, fails (literal variant: holds, holds). Kou-SEJ-Q’s residuals were largest in the top VIX tercile (1.80 and 0.81 points above the larger of the other two) but not in the small top relative-spread tercile (-0.53 [$-1.65, 0.58$] and -0.002 [$-0.422, 0.421$]); the holdout failure rested on the sign of a number close to zero. The pre-registration speaks of ‘residual errors’ without naming a model, and the spread component was positive in both samples for BS-HV, CRR-HV, DJ-hist, Kou-SEJ-P and Kou-SEJ-Q-prereg and negative in both only for Kou-SEJ-Q, so the verdict depends on whose residuals are tested.

P5: holds, holds. DJ-implied reduced the RMSE of DJ-hist by 58.4% and 51.7%, Kou-SEJ-Q that of Kou-SEJ-P by 55.0% and 51.6%. Much of this was the value of the at-the-money level, and part of the remainder the variance risk premium that the historical law omits.

P6: fails, fails. The double-exponential shape did not help in the wings: the wing-minus-ATM difference of the gain of Kou-SEJ-Q over DJ-implied was -11.8 and -10.1 percentage points, and at the money the gain was -10.0% and -10.4% , outside the $\pm 5\%$ band (inconclusive under the rule in the development sample). The literal variant’s large positive wing-minus-ATM difference came from its much worse at-the-money fit (-116.7% in the holdout), not from better wings. With two other bootstrap seeds one component verdict changed, the development at-the-money component of P6 (inconclusive to fails); no prediction’s verdict changed. The same component fails with week-clustered intervals ($[-11.1, -8.7]$ against $[-11.8, -8.2]$); it sits at the boundary of the rule, and the prediction fails through its wing component either way.

P7: holds, holds. BS-HV under-priced in every $|\Delta|$ bin (mean errors -8.23 , -8.54 and -9.41 from the wing to the in-the-money bin in the development sample, -10.58 , -10.98 and -11.61 in the holdout), and DJ-implied under-priced the wings relative to the money (-0.84 and -0.83). The empirical clause compares each model’s wing error with its own at-the-money error, so the crossings of the equal-variance smiles noted in Section 5.4 drop out: relative to its own at-the-money volatility, the smile of a fat-tailed announcement law rises away from the money while the Gaussian smile is flat (Table 4; Figure 7 in Appendix B.19). The clause held as registered, but an exploratory breakdown that was not pre-registered shows that it measures the background equity smile rather than the tails of the announcement law. It nets an under-pricing of out-of-the-money puts against an over-pricing of out-of-the-money calls: against the at-the-money error of the same side, -2.00 [$-2.11, -1.89$] for puts and 0.62 [$0.49, 0.74$] for calls in the development sample, -2.64 and 1.80 in the holdout, with puts 54% and 57% of the wing quotes. The market smile was thus the familiar put skew with some curvature, which on average over the two wings rises away from the money. The flat post-event at-the-money model shows the same sign after the announcement, when no announcement remains in the chain (-0.92 [$-1.05, -0.79$] and -0.59 [$-0.66, -0.52$], on bins of the post-event $|\Delta|$), and on the same contracts the pre-event value was not more negative than the post-event one (pre minus post 0.37 [$0.20, 0.54$] and 0.10 [$0.00, 0.20$]). The first clause is the omitted-variance reading of the verification study.

P8: inconclusive, holds. The premium components held in both samples (ratio minus one 0.70 and 0.14; high minus low VIX tercile 1.02 and 0.07), and Kou-SEJ-Q fitted the mids better. The payoff component held in the holdout and was not judged in the development sample, whose 19 weekly events are fewer than 30; hence the inconclusive verdict. As Section 8.8 showed, the payoff comparison had little power. Section 9 reads the failures.

8.10 Recent announcements

The frozen database of Section 7 ends on 22 April 2025. To test the models on recent announcements of large, heavily traded firms, I collected end-of-day chains from the live version of the same source for 165 announcements by 31 large firms with large past announcement moves between 1 May 2025 and 10 October 2026, each with a snapshot at the last close before the announcement and at the first close after it (Section 7.6; Appendix F.6 lists the firms). This sample is not part of the pre-registered design, and it is not a test of a frozen specification on unused data: I priced it on 13 October 2026, before the freeze of the main specification on 14 October 2026, and its results informed that specification, whose calibration of Kou-SEJ-Q, estimator of Kou-SEJ-P and descriptive metric were adopted from it (deviations (i), (ii) and (iv) of Section 8.1). It is an out-of-calibration check: every model is evaluated on an expiration that no calibrated model uses.

Design. The models are those of Section 8.1, with three differences: BS-HV uses the vendor’s 30-day historical volatility as the background volatility, because no daily spot series is stored for these dates; BS-IVlong uses the at-the-money implied volatility of the longest quoted expiration; and DJ-hist and Kou-SEJ-P take the announcement law from the firm’s previous announcements (median 20 prior events), corrected for the diffusion between the two snapshots. Kou-SEJ-Q is calibrated as in (48) to the two shortest expirations (the calibration converged in 162 events), and Kou-SEJ-Q-Am adds the finite-difference early-exercise premium to the puts. The vendor quotes three or four expirations at roughly constant tenors, the nearest about two weeks out, so I evaluate all models on the *third* expiration, which neither DJ-implied nor Kou-SEJ-Q uses: 3,184 quotes in 141 events of 28 firms after the filters of Section 7. This is the out-of-calibration test that the main sample can supply for only a few events. Mkt-ATM uses the at-the-money volatility of the third expiration and so remains a reference fitted to the evaluated expiration. In the 82 events whose chain has only three expirations, the longest is the third, so BS-IVlong then uses the at-the-money volatility of the evaluation expiration too and coincides with Mkt-ATM; only in the other 59 events is it a forecast from a longer expiration (Table 56 in Appendix F.6). The fitted expirations (6,729 quotes) are in Panel B of Table 14. Errors are implied-volatility differences.

Results. Black–Scholes with historical volatility missed the market by 11.93 volatility points on the out-of-sample expiration, with a bias of -7.86 points: as in the main study, a diffusion calibrated to ordinary days was far too cheap before an announcement. Adding the historical announcement law removed about a third of that error (32% for DJ-hist and 31% for Kou-SEJ-P) but left the models 8.08 and 8.18 points from the market. The option-implied models predicted the third expiration to within 2.97 points (DJ-implied) and 3.00 points (Kou-SEJ-Q), 2.83 points with the early-exercise premium, against 2.40 points for the reference Mkt-ATM. Relative to BS-HV they captured 0.86 [0.83, 0.88] of BS-HV’s absolute dollar error (90% bootstrap interval; Table 53 in Appendix F.6). Relative to BS-IVlong the gain was 0.11 $[-0.07, 0.27]$, an interval that contains zero, but this pools two benchmarks. In the 82 events with three expirations, where BS-IVlong coincides with Mkt-ATM, Kou-SEJ-Q captured -0.23 $[-0.38, -0.12]$ of its error; in the 59 events with a longer expiration, where BS-IVlong is a forecast, it captured 0.43 [0.21, 0.58] and DJ-implied 0.44 [0.24, 0.58] (RMSE 2.67 and 2.69 against 3.43 for

Table 14: Pre-event pricing errors for recent announcements of 28 large firms with large past announcement moves (May 2025 to October 2026). Panel A: 3,184 quotes of the third-shortest expiration in 141 events, which no calibrated model uses; Mkt-ATM uses its ATM implied volatility, and so does BS-IVlong in the events whose chain has only three expirations, where the longest quoted expiration is the third. Panel B: 6,729 quotes of the two shortest expirations, whose ATM implied volatilities Mkt-ATM, DJ-implied and Kou-SEJ-Q fit by construction. Common sample in which every model is defined. Errors are model minus market, in implied-volatility points (RMSE pooled over quotes and averaged over per-event RMSEs; bias = mean signed error), in dollars (MAE), relative to the mid for mids of at least \$1 (MAPE), and the share of model prices inside the bid–ask spread.

Model	RMSE (pooled)	RMSE (event avg.)	Bias	MAE (\$)	MAPE (%)	Inside (%)
<i>A. Third expiration (out of sample except for Mkt-ATM, and for BS-IVlong in chains with three expirations)</i>						
BS-HV	13.48	11.93	−7.86	4.078	22.7	14.2
BS-IVlong	2.96	2.83	−1.43	0.645	3.5	64.0
Mkt-ATM	2.41	2.40	−0.55	0.459	2.6	73.1
DJ-hist	10.10	8.08	−0.95	3.049	16.2	22.5
DJ-implied	3.21	2.97	−0.90	0.576	3.8	62.8
Kou-SEJ-P	10.20	8.18	−1.30	2.993	16.2	23.8
Kou-SEJ-Q	3.34	3.00	−0.79	0.575	3.9	63.3
Kou-SEJ-Q-Am	3.15	2.83	−0.28	0.554	3.8	65.4
<i>B. Two shortest expirations (ATM fitted by Mkt-ATM, DJ-implied and Kou-SEJ-Q)</i>						
BS-HV		18.90	−15.96			9.6
BS-IVlong		10.50	−9.37			17.8
Mkt-ATM		2.75	−0.66			78.1
DJ-hist		10.07	−1.98			22.1
DJ-implied		2.71	−0.63			78.3
Kou-SEJ-P		10.41	−3.05			20.8
Kou-SEJ-Q		3.00	−0.10			73.4
Kou-SEJ-Q-Am		2.99	0.26			72.9

BS-HV: vendor 30-day historical volatility at the pre-event close; BS-IVlong: ATM IV of the longest quoted expiration; Mkt-ATM: ATM IV of the quote’s own expiration; DJ: Black–Scholes with announcement variance (historical or two-expiry implied); Kou-SEJ-P/Q: scheduled double-exponential announcement jump with historical (P) or option-implied (Q) variance; -Am: American (finite-difference early-exercise premium added to puts).

BS-IVlong; post-freeze, not pre-registered). Where a longer-dated at-the-money volatility was a genuine forecast, the announcement decomposition improved on it. Replacing the Gaussian announcement jump of DJ-implied by the double-exponential law of Kou-SEJ-Q, with the same fitted at-the-money levels, raised the out-of-sample RMSE by 5.6% in the wings and by 11.2% at the money, relative to the RMSE of DJ-implied. The 90% intervals of these increases, a post-freeze addition, are $[-0.4, 12.8]$ and $[3.6, 21.0]$ with events resampled and $[0.2, 11.2]$ and $[3.5, 20.0]$ with firms resampled. For the eight largest moves of the period the market smile before the announcement was flatter than the smile implied by a double-exponential law fitted to the firm’s past announcements, which rests on a median of 20 observations and is noisy (Figure 18 in Appendix F.6; Figure 19 there shows the same pattern across all events). At these maturities the evidence favored a Gaussian announcement jump calibrated to the term structure over a fat-tailed one whose shape is taken from history, clearly at the money and as a point estimate in the wings. For the puts the early-exercise premium had a median of \$0.171 (0.92% of the price) and a 95th percentile of \$1.225 (2.14%); on the third expiration, which the calibration does not use, it lowered the Kou-SEJ-Q error from 3.00 to 2.83 points, whereas on the puts of the main study, almost all of which lie in the calibration expirations, it did not (Section 8.7).

Pre-registered predictions. Table 54 in Appendix F.6 applies the predictions to the out-of-sample expiration, with the implied-volatility difference and a simpler reading than Tables 12 and 13: a prediction holds only if every stated component holds, without intervals, and the post-event half of P1 is not evaluated, since no post-event leg was priced. P2a, P2b, P4 and P6 are judged on Kou-SEJ-Q only, and P4 uses terciles of the recent sample, as the caption of the table states. Event-jump models cut the error of BS-HV by more than 30% (P1, holds), option-implied announcement variance beat historical variance (P5, holds), Mkt-ATM was not beaten at the money (P3, holds), and BS-HV under-priced across moneyness (P7, holds). The predictions about where the gains concentrate failed: the reduction was not larger in the wings (P2a, fails) or for firms with larger past moves (P2b, fails), residuals did not peak in the top VIX tercile (P4, fails), and the double-exponential shape lost 11.2% at the money, outside the 5% band (P6, fails; its wing-minus-ATM clause held, since the shape lost less in the wings, 5.6%, than at the money). The option-implied announcement standard deviation was close to the realized one on average, with a ratio of 1.01 [0.88, 1.16]; as a ratio of mean variances, the form of the main study’s ratio (49), it was 1.02 [0.78, 1.35], and 1.12 [0.84, 1.53] with the realized move corrected for the diffusion between the snapshots (post-freeze). The ratio of standard deviations was 0.77 in the lowest VIX tercile and 1.29 in the highest, but the historical law did not forecast the size of the move better than the option-implied one (mean absolute errors 4.3 and 4.3 percentage points), so P8 fails.

Relation to the main sample, and caveats. The recent sample agreed with the main study where the two can be compared: BS-HV was far too cheap, the historical announcement law removed part of its error but left the models far from the market, the option-implied models came within a few volatility points, and the double-exponential shape taken from history did not improve on a Gaussian announcement jump with the same at-the-money levels, here on an expiration that neither calibrated model uses, so the pattern also held outside the calibration expirations, in a sample that had informed the specification. On the out-of-calibration quotes of the main sample Kou-SEJ-Q was slightly ahead of DJ-implied (4.38 against 4.47 in the holdout, 50 events), here slightly behind (3.00 against 2.97); neither difference supported a role for the shape. The two samples differ on the premium (the variance ratio of the main study was 1.70 [1.60, 1.82] in the development period and 1.14 [1.09, 1.19] in the holdout, against 1.12 [0.84, 1.53] here in the same form) and on the VIX half of P4, which held in the main sample and not here; P1 held here because its post-event half was not evaluated, and P8 failed here because the move forecast of the historical law was no better, while in the main holdout it held on a payoff comparison that had little power. The firms were chosen, not sampled, and for their large past moves; the intervals resample events, and with firms resampled they are similar (mispricing captured by Kou-SEJ-Q relative to BS-HV [0.83, 0.88], relative to BS-IVlong [−0.07, 0.24]); the background volatility of BS-HV, DJ-hist and Kou-SEJ-P is the vendor’s historical volatility; dividends are set to zero; quotes are end-of-day; and the vendor’s constant-tenor expirations mean that the evaluation horizon is about two to four weeks rather than the earnings-week contract.

9 Discussion and limitations

This section reads the results of Sections 6 to 8 together. Unless stated otherwise, errors are implied-volatility differences on the common sample, as in the descriptive tables of Section 8, and reductions relative to BS-HV use the pre-registered vega-scaled metric.

What the scheduled jump buys. The cleanest measure of the announcement jump is the comparison that uses no option price. BS-HV, DJ-hist and Kou-SEJ-P share the background volatility σ_{HV} and differ before the announcement only by the announcement term, whose

variance DJ-hist and Kou-SEJ-P take from the firm’s past moves, DJ-hist as their mean square and Kou-SEJ-P net of the diffusion in the announcement window, so that the difference between the two is not the shape of the law alone; a Gaussian law with the variance of Kou-SEJ-P puts most of their holdout difference on the variance (Appendix F.1). That term cut the RMSE of BS-HV by 38.9% (DJ-hist) and 36.1% (Kou-SEJ-P) in the holdout and by 23.7% and 25.1% in the development sample (Table 9), out of sample in both. It also removed most of the dependence of the error on the maturity, which is the signature of a lump of variance that does not shrink with the time to expiration: in the holdout the RMSE of BS-HV was 16.9 points higher for 5–15 than for 31–45 days to expiration, that of DJ-hist 4.1 points and that of Kou-SEJ-P 5.3 (17.2, 7.0 and 6.4 in the development sample; Tables 42 and 43 in Appendix F.2). The gap of Kou-SEJ-Q, whose background and announcement variance are fitted to the two expirations of the event’s chain and so to its term structure, was smaller still (0.8 and 2.1), largely by construction. After the announcement the term is gone, and DJ-hist and Kou-SEJ-P priced exactly as BS-HV (Table 49 in Appendix F.4): for a fixed background, the fall of their advantage from the reductions just quoted to zero is the contribution of the jump. The option-implied models kept part of their advantage after the announcement (Kou-SEJ-Q 30.6% in the development sample and 20.6% in the holdout) because their backgrounds are option-implied, so their post-event comparison measures the background and not the jump, which is why the post-event half of P1 failed for them. The jump also predicts the fall of implied volatility across the announcement, the volatility crush (Table 50 in Appendix F.4). BS-HV, CRR-HV and BS-IVlong, which have no announcement component, predicted changes of less than one volatility point (Mkt-ATM uses the post-event quotes and is not a forecast), while the market’s at-the-money volatility fell by 6.69 points on average in the development sample and by 8.35 in the holdout. Every model with an announcement component predicted a fall of the market’s order, the calibrated option-implied ones (DJ-implied and Kou-SEJ-Q) too large on average (Kou-SEJ-Q 8.65 and 11.93), because their background lay below the post-event market level, and all with Mincer–Zarnowitz slopes below one. The scheduled jump thus gets the existence and the order of magnitude of the crush; it does not get its cross-section, which also depends on how the diffusive volatility moves across the announcement, and the model holds that constant.

Premium or mispricing. Three objects are easily confused here, and the design separates them only in part. The first is the earnings variance risk premium, a difference between two laws: the option-implied announcement variance exceeded the realized one by factors of 1.70 [1.60, 1.82] in the development period and 1.14 [1.09, 1.19] in the holdout (Table 11, Panel A), ratios that assume calendar-time maturities and a flat background term structure (in trading time the holdout ratio was 1.06, and a placebo decomposition without an announcement averaged 0.37 of the realized variance; Section 6.7), as Dubinsky and Johannes (2004) find for the risk-neutral announcement volatility and as the link to non-diversifiable announcement risk in Barth and So (2014) would lead one to expect. A model that prices with the historical law and is judged against market quotes inherits this premium as a bias. The second is the noise of a short history, which is larger: at $n = 12$, the Monte Carlo history length nearest the holdout median of 11 prior announcements, the historical variance of a single firm has a standard deviation of 0.66 times its true value (Table 23 in Appendix C.3), and Panel F of Table 6 shows the history erring in both directions, above the realized variance in the development period (ratio 1.50) and below it in the holdout (0.88). Accordingly the mean error of DJ-hist was small next to its RMSE (vega-scaled 1.33 against 8.62 in the development sample, -1.34 against 8.37 in the holdout): the distance of the historical models from the market was mainly a dispersion across events, not the shift of level that a premium alone would produce. The third object, model error, is what remains; the comparisons of Section 8 isolate it only for the models that use the event’s own chain, and for them only away from the at-the-money strike. None of the three is mispricing

in the sense of a profit opportunity, and the paper makes no claim of that kind: mispricing captured is a distance to a benchmark. The premium was also uneven within each period. On the priced events the ratio was 1.56 and 1.05 [0.99, 1.11], below its value on all events of the primary sample, so the events without a kept quote, mostly those of low-priced and wide-spread names (Appendix D.5), carried a larger premium; the holdout value on the priced events is closer to the positive returns of at-the-money straddles bought before announcements in Gao et al. (2018). For the recent large firms the ratio of the implied to the realized announcement standard deviation was 1.01 [0.88, 1.16] (Section 8.10), and 1.12 [0.84, 1.53] as a ratio of mean variances with the diffusion-corrected move, the form of (49).

Where the residual is largest. For the option-implied models the residual was largest in the money, in the top VIX tercile and, in the holdout, among the lowest-priced stocks (Appendix F.2). The in-the-money bin had the largest RMSE of every model (Kou-SEJ-Q 6.36 against 1.76 at the money in the holdout). Part of that gap comes from the binning itself: the bins use the delta at the quote’s own implied volatility, so which bin a quote falls in depends on the market volatility whose error is being measured. Binned instead by the delta at the at-the-money volatility of the expiration, a post-freeze check, the in-the-money RMSE of Kou-SEJ-Q in the holdout was 4.52, still the largest of the three bins (Appendix F.2). Beyond that, two readings compete. Prices in the money are dominated by intrinsic value and most exposed to early exercise and dividends, and the parity-implied spot applies European parity to American quotes; an error of the spot moves calls and puts of the same strike in opposite directions. On the American subsample (Panel B of Table 51 in Appendix F.5) the European Kou-SEJ-Q instead under-priced both the in-the-money calls (−5.02) and the out-of-the-money puts (−1.87), which lie below the forward, and over-priced both the out-of-the-money calls (2.42) and the in-the-money puts (1.41), which lie above it, in the holdout as in the development sample. Calls and puts of similar strikes erred in the same direction, which points to the tilt of the model’s smile rather than to the spot (see the shape paragraph below). In the holdout the RMSE of Kou-SEJ-Q was 5.58 in the top VIX tercile against 4.82 and 4.11 in the other two (5.99 against 4.11 and 4.15 in the development sample), the half of P4 that held, whose statistic is the difference from the larger of the other two; the top relative-spread tercile, the other half, was small and did not have the largest errors, and whether P4 holds there depends on whose residuals are tested (Section 8.9). The lowest price tercile, the pre-registered size proxy, had the largest holdout errors of every model (Kou-SEJ-Q 6.15), as wider spreads relative to the price and a less precise parity spot for low-priced stocks lead one to expect (Table 34 in Appendix D.3). The gains relative to BS-HV followed the omitted announcement variance: in the holdout the share of BS-HV’s error that Kou-SEJ-Q captured fell from 0.86 for 5–15 days to 0.77 for 31–45 days and rose from 0.67 to 0.85 from the bottom to the top tercile of past moves (Figure 16 in Appendix F.2). The past-move half of P2 held accordingly, except for DJ-hist in the development sample, while its wing half failed for DJ-hist, DJ-implied, Kou-SEJ-P and Kou-SEJ-Q and held only for the literal variant Kou-SEJ-Q-prereg, whose at-the-money fit was worse: the gains were largest at the money, where the option-implied models are fitted.

Poisson against scheduled timing. The verification study puts the same total variance into three models (Table 4; Figure 7 in Appendix B.19). Relative to the scheduled double-exponential jump, a Poisson model with the same jump law and one expected jump prices the at-the-money straddle about eight percent lower, and Black–Scholes with the variance spread over the option’s life about eight percent higher (at-the-money implied volatilities 0.62185 for the scheduled jump, 0.57298 for the Poisson model and 0.67185 for Black–Scholes). I did not price a Poisson specification on the data, but three results of the empirical study indicate what it would miss. First, the timing. The two-expiry decomposition presumes a lump of variance that does not grow with the maturity, and it gave a positive announcement variance for all but

2.9% of the kept events; the maturity profile of BS-HV's error, which the announcement term removed, is the one such a lump produces. A Poisson component with a constant intensity adds variance in proportion to the time to expiration and produces neither. Second, the crush: with a constant intensity no particular date is followed by a systematic fall of implied volatility, while the market's at-the-money volatility fell by 8.35 points on average in the holdout. Third, the level at equal variance. The literal pre-registered variant, which gives a double-exponential law the variance of the Gaussian decomposition, was biased low (-2.94 and -3.90 , against -1.18 and -1.11 for Kou-SEJ-Q), and the calibrated Kou-SEJ-Q needed a median of 1.16 times the decomposition's variance to match the at-the-money quotes: the at-the-money effect of the verification table, measured on the data. On the numbers of that table a Poisson law would put still less value at the money. The wings do not separate the laws. Relative to its own at-the-money level DJ-implied under-priced the wings (-0.84 and -0.83 points, vega-scaled), but this nets under-priced out-of-the-money puts against over-priced out-of-the-money calls, the familiar put skew, and the same sign appears after the announcement (-0.92 and -0.59 ; Section 8.9), so P7 measures the background smile, not the tails of the announcement law. What separates the scheduled from the Poisson specification is the timing, and the timing is what the term structure and the crush identify.

Why the double-exponential smile was too steep. In the wing bin Kou-SEJ-Q had the smaller mean error but the larger RMSE (2.97 against 2.50 for DJ-implied in the holdout, 2.42 against 2.01 in the development sample), and Figure 3 shows where. In the bins of standardized moneyness between 0.6 and 1.0 below and above the forward, the market's implied volatility minus the flat at-the-money level was 4.17 points on the side of low strikes and -1.23 on the side of high strikes in the holdout (3.27 and -0.27 in the development sample), the familiar equity skew. The smile of Kou-SEJ-Q lay above the flat level on both sides, by 0.63 and 1.39 points (0.39 and 0.91): it moved the low-strike wing toward the market, by too little, and the high-strike wing away from it. The double-exponential law thus did two things. Its fat tails raised both wings, which the market wanted on the side of low strikes and not on the other; and its shape, taken from the firm's past moves, tilted the smile toward high strikes, the opposite of the market's skew. Matching the at-the-money level with a fat-tailed law also takes more announcement variance than the Gaussian decomposition (previous paragraph), and that variance appears in the wings. Three findings of Section 6 explain why the firm's shape helped so little. Under $\mathbb{P}$ the pooled announcement moves were only slightly negatively skewed (skewness -0.34 , mean up and down moves in the ratio 0.965), and a firm's own share of up moves varied across firms by little more than binomial noise (p -value 0.10 for homogeneity), so the firm-specific asymmetry is mostly estimation noise. In the predictive comparison of Table 24 in Appendix C.7, fat tails helped (the Student t_4 and Laplace laws beat the Gaussian, with intervals above zero), the Kou-SEJ-P law with the firm's own shape added nothing distinguishable from zero to a Gaussian with the same corrected variance (-0.007 nats per event in the holdout, with an interval that contains zero), and the same law with a shape pooled across firms beat its firm-shape version by $+0.121$ nats and the corrected Gaussian as well. Under $\mathbb{Q}$ the pre-event skew was steeper than any zero-mean double-exponential jump on a flat diffusion can produce in 90% of the events where no admissible p_E of Kou-SEJ-Q-prereg matched the market skew (Appendix C.12), so part of the market's skew presumably belongs to the background, which a constant diffusion cannot carry; and the zero-mean law fitted to that skew had a median share of up moves of 0.71, against 0.50 in the historical laws, which with a zero mean means rarer and larger down moves. Read together, the announcement law appears fat-tailed under $\mathbb{P}$ and more negatively skewed under $\mathbb{Q}$ than the moves under $\mathbb{P}$, while the pre-event smile does not identify its tails under $\mathbb{Q}$ (Section 8.9), and it is poorly estimated in its shape from one firm's few moves. On the recent announcements, on an expiration that neither calibrated model uses, the double-exponential shape raised the RMSE by 5.6% in the wings and 11.2% at the money relative to the Gaussian

(Section 8.10), so the pattern also held outside the calibration expirations, in a sample that had informed the specification. The evidence speaks against the firm-by-firm shape more than against the double-exponential family, whose pooled version predicted the moves better than the Gaussian; whether it would also price better is untested. The family has a further limit: a law whose density peaks at a zero move, as the double exponential does, tends to give a smile that rises on both sides of the money (Table 4), while the concave pre-announcement curves that [Alexiou et al. \(2025\)](#) document point to a law with two modes. The average market smile of this sample was skewed rather than concave, so a symmetric two-mode law would not by itself fit the low-strike wing, and whether a bimodal announcement law prices the event-level smiles better is untested here. A pooled shape, or a shape calibrated to the pre-event smile together with a skewed background, is the next specification to try, and a skewed law with two modes the one after it. I have not priced either: the comparisons that suggest the first include the holdout, and a fair test of both needs announcements that have not been used, such as those after the snapshot of Appendix E.

Early exercise at these maturities. Section 8.7 and Appendix F.5 give the numbers: the American price helped for calls with an ex-dividend date in their life, also after the post-freeze recalibration to American quotes, and not for puts, and the median put premium stayed below one percent of the European price even at the holdout's rates. Two points go beyond them. First, the put result says more about the smile than about early exercise. The double count of the frozen calibration explains only part of the rise of the put error; most of the rise lies in the money, where the European model already over-priced the puts (mean error 1.41 in the holdout), on the high-strike side of the tilted smile of the previous paragraph, and a positive premium added to an over-priced put raises its error further. On the recent sample, whose evaluation expiration lies outside the calibration, the premium instead lowered the error of Kou-SEJ-Q from 3.00 to 2.83 points. Second, at equal variance the premium depends on the timing of the jump, by a factor of forty between the scheduled and the Poisson specification in the verification study, although both premia are small there. In the data the premium rarely mattered against the cost of trading: it exceeded the quote's half-spread for 2.7% of the holdout puts and 12.4% of the calls with a dividend, and the gain for those calls, which the binomial model shares, rests on a continuous dividend yield (Section 8.7). The American machinery therefore matters chiefly with dividends, at higher rates, in the money and at maturities longer than those of the main sample.

What the failures of the pre-registration teach. Of the eight predictions, four failed in both samples (P1, P2, P4 and P6); in the development sample P3 also failed and P8 was inconclusive. The failures are of three kinds (Section 8.9). Some misdescribed the design. The post-event half of P1 presumed that the event models share BS-HV's background after the announcement, which only DJ-hist and Kou-SEJ-P do; the development failure of P3 is a difference of -0.0003 points in a comparison that is nearly an identity, since DJ-implied equals Mkt-ATM inside the calibration expirations, and it rests on the out-of-calibration at-the-money quotes of 4 events; and P8 could not be judged in the development sample, which has 19 weekly events. One was underspecified: P4 named no model, and its spread half holds or fails depending on whose residuals are tested. And two failed for substance: the wing half of P2 and P6 failed because the gains of the event models concentrated at the money, where the option-implied models are fitted, and because the double-exponential shape taken from the firm's past moves did not improve on a Gaussian law away from it. Three lessons follow for pre-registering a study of this kind. A prediction about a calibrated model should be stated for what the calibration leaves free, here the smile away from the calibration strikes and the expirations outside the pair, not for what it pins down. Each prediction should name its model, its statistic and its rule for combining components, which this plan left open and which I fixed after the headline errors had been seen (deviation (v) of Section 8.1). And each component needs a minimum sample,

or a prediction about a secondary sample is decided by a handful of events. Against these weaknesses stand the parts of the design that held: the freeze of the specification before any holdout quote was priced, the hashes of the pricing code and, later, of its inputs, the report of every component under both metrics (4 development and 2 holdout components change with the metric, no prediction does), the literal pre-registered variant beside the primary one (it changes only P4, which holds with it in both samples), and the check with other bootstrap draws.

9.1 Limitations

- **End-of-day quotes.** Every price is an end-of-day bid–ask snapshot. Intraday dynamics on the announcement day, the opening auction after an after-close report and execution inside the spread are not observed, and a mid-quote is not a transaction price; [Muravyev and Pearson \(2020\)](#) show that the cost of trading options can be well below what quoted spreads suggest, so the quoted spread is only a loose measure of the uncertainty of the mid. The paper reports no trading exercise, and none of its statements concerns returns net of costs.
- **Two to three expirations per snapshot.** The database carries 2.35 expirations per symbol and date on average from 2020 on, so the term-structure decomposition of the announcement variance uses one pair of maturities, weekly expirations drop out of the feed in their last days, and the earnings-week weeklies enter only as a secondary sample evaluated on settlement payoffs. Almost all common-sample quotes lie in the two calibration expirations (99.3% in the holdout), so the main sample tests the option-implied models out of calibration on only 624 holdout quotes in 50 events; the recent sample supplies that test on more chains but fewer firms.
- **Parity-implied spot.** The underlying price is implied from at-the-money call–put pairs of the nearest expiration rather than taken from a closing-price feed, by a parity relation that ignores dividends and the early-exercise premia of American quotes. The tolerance rule that accepts a spot is itself a deviation from the plan (Section 7.3), and the check against closing prices covers only the 60 most-quoted symbols (median absolute deviation 0.05%); for the other names the dispersion of the parity values is the only check. A spot error moves calls and puts of a strike in opposite directions, and the mean errors of the American subsample do not show that pattern, but this does not rule out spot errors in single events.
- **Another underlying’s chain under a ticker.** On a few dates the database files the chain of another underlying under a ticker, and the parity pairs then agree closely on the wrong price. The spot-consistency screen that I added after the main study had been priced (Section 7.3 and Appendix D.3; deviation (viii) of Section 8.1) rejects a spot more than a factor of 3 above or below both neighboring spots of the ticker. A wrong chain whose price lies within that factor of the right one, a wrong chain that persists over several snapshots, and the first or last spot of a series pass it.
- **$\mathbb{P}$ -measure against $\mathbb{Q}$ -measure announcement law.** The historical law of an event rests on the firm’s prior announcements, a median of 11 in the holdout and 5 in the development sample, so its scale is imprecise and its shape is not identified at the firm level (Section 6.3). It is also not the pricing law: its difference from the option-implied law contains a risk premium, estimation error and model error, in proportions that this study bounds but does not separate. Every double-exponential law priced in the main study took its shape from history or from the skew slope alone: Kou-SEJ-Q the asymmetry and tail ratio of the firm’s past moves (a pooled shape where the firm’s law is undefined), Kou-SEJ-Q-prereg an asymmetry fitted to the skew slope with the decomposition’s variance.

The family with a pooled shape, or with a shape calibrated jointly with the level to the pre-event smile, was not priced, so the shape results speak against these shapes rather than against the family; the planned test uses the announcements of Appendix E.

- **No stochastic volatility.** The background diffusion coefficient is constant over the option’s life and across the announcement, so time variation and skew of the diffusive volatility are absorbed by the announcement parameters or left in the residual. The events where no admissible p_E of Kou-SEJ-Q-prereg matched the market skew and the over-predicted crush of the option-implied models suggest that both matter here. Bakshi et al. (1997) find stochastic volatility to be of first-order importance for pricing index options, and Dubinsky and Johannes (2004) combine it with an announcement jump; the model and the solver of this paper do not include it.
- **American premium small at these maturities.** For near-the-money options with a few weeks to expiration and no dividend in their life, the early-exercise premium is a small fraction of the price, so the empirical results speak mostly to the European part of the model. The American solver was applied to a stratified tenth of the events (142 development and 545 holdout events), and its consistent calibration to American quotes is a post-freeze analysis (Appendix F.5). In the data the American price mattered for calls with a dividend, where the binomial model gains as much under the same continuous dividend yield, and it would matter more at the longer maturities that the main sample excludes.
- **Vendor implied volatilities not used.** The database’s implied-volatility and Greek columns are not used; implied volatilities are recomputed from bid and ask with the paper’s own conventions. They have a correlation of 0.979 with the vendor’s and a median absolute difference of 1.48 points, which falls to 0.62 when time is measured in calendar days (Appendix D.5), so the results are not directly comparable with studies that take vendor volatilities as given.
- **Dividends.** The dividend yield of an option is set from the realized dividend table, $q = D/(ST)$ for the dividends going ex in its life: a known-dividend assumption, since declaration dates are not in the data, and a continuous yield in place of a discrete payment. It affects 35,430 of the 151,969 kept quotes. The post-freeze calibration showed that with a dividend in the option’s life the continuous yield does not reproduce the split between call and put prices at the calibration strikes (Appendix F.5), and the at-the-money implied volatilities of the decomposition use $q = 0$ (Appendix D.4). The solver could instead apply the shift $S \rightarrow S - D$ at the ex-dividend date as a second jump condition, the standard finite-difference treatment of a discrete dividend (Leentvaar and Oosterlee, 2006), with the same restart as at the announcement; the frozen specification does not.
- **The calendar clock of the decomposition.** Option-implied volatilities accrue in calendar time, the clock of the two-expiry decomposition, while σ_{HV} accrues in trading time, a recorded deviation from the trading-time convention of the plan (deviation (iii) of Section 8.1). The decomposition depends on the ratio of the two maturities, which weekends and holidays change, so the option-implied announcement variance, the premium ratio and the calibration of Kou-SEJ-Q depend on the clock. A post-freeze recomputation of the decomposition in trading time (Appendix C.11) lowered the premium ratio from 1.70 to 1.52 in the development sample and from 1.14 to 1.06 in the holdout, whose interval with announcement-week clusters, $[0.99, 1.13]$, contains one; the priced models and the calibration of Kou-SEJ-Q were not re-estimated in trading time.
- **The flat background of the decomposition.** The two-expiry decomposition counts any difference between the background implied variances of its two expirations as announce-

ment variance. Applied to the same two expirations after the announcement, where it should average zero, it averaged 0.48 and 0.37 times the realized announcement variance (development, holdout) and rose with the VIX (Appendix C.11). The premium ratio and the option-implied announcement variance of the option-implied models therefore contain part of the background slope; the development premium exceeds both placebos, the holdout one does not.

- **Verdict rules fixed after the headline errors were seen.** The pre-registration fixed the predictions and the failure rule but not how components combine, which statistic represents each prediction, the minimum number of events of a component or the treatment of ties. I coded these rules after the headline common-sample RMSEs of both samples had been displayed (deviation (v) of Section 8.1). No model or parameter changed afterwards. The later edits of frozen files changed a table caption, non-numerical code of the package and the scripts (paths, interfaces, documentation) and, as a recorded bug fix, the block that writes the ex-ante table of Appendix E. The re-pricing after the data correction above used the frozen code, and the re-pricings after the code edits reproduced every frozen price; but the order of these steps cannot be verified from the files. Every component is reported with its interval, so a reader can apply another aggregation, and neither the implied-volatility difference nor two other bootstrap seeds changed an overall verdict.
- **Dependence across announcements.** The pre-registered intervals resample announcements as independent draws, while announcements cluster in weeks with shared market news and within firms. With announcement weeks resampled the intervals of the predictions are wider (median ratio 1.4 and 1.4), one component verdict changes in each sample, and no prediction's verdict changes (Section 8.9).
- **A development sample of 2021 announcements.** DJ-hist and Kou-SEJ-P need 4 prior announcements of the firm and the calendar starts on 22 January 2020, so no 2020 announcement enters the common sample, and every development comparison across all nine models refers to the 1,386 announcements of 2021. The pandemic quarters of 2020 enter only the broad sample, the premium and the estimation (Table 41 in Appendix F.1 and Table 6). The specification was also frozen on a period with rates near zero, which is why the early-exercise questions of Section 8.7 arose only in the holdout.
- **Selection toward liquid names.** The relative-spread filter removes most quotes of low-priced and illiquid names: by development-sample tercile of the spot, 29%, 42% and 64% of the events with a primary expiration kept a quote, and by tercile of the relative spread 90%, 49% and 6% (Table 39 in Appendix D.6). The pricing results therefore concern the liquid, higher-priced part of the optionable cross-section, the top spread tercile of the pre-registered cuts is thin, and the events without a kept quote carried a larger premium than the priced ones.
- **Sample period.** The sample starts in January 2020 and includes the pandemic quarters; the holdout contains the rate cycle that began in 2022 and the gap in the database between 29 July 2024 and 12 August 2024. The two periods differ in the rate regime, the VIX and the background volatility (Section 7.5), and a little over five years is a short history for a premium that varies with the state of the economy (Kelly et al., 2016).
- **The recent sample.** The firms of Section 8.10 were chosen, not sampled, from large firms with large past announcement moves. Their background volatility is the vendor's historical volatility, dividends are set to zero, and the evaluation expiration is a constant-tenor contract rather than the earnings-week one, so the recent results complement the main study rather than replicate it. It was also priced before the freeze, and its calibration,

estimator and metric were adopted for the main study (deviations (i), (ii) and (iv) of Section 8.1), so it informed the specification rather than testing it on unused data.

10 Conclusion

I have priced American equity options around scheduled earnings announcements with a Kou jump-diffusion that carries a double-exponential jump at the known announcement time, solved the resulting partial integro-differential linear complementarity problem with an IMEX-CNAB finite-difference scheme that treats the announcement as a jump condition inside the time stepping, and tested the solver against Fourier-cosine prices and published American benchmarks. With space and time refined together, the European problem converged at observed order two: from the second refinement on, 2.00 to 2.01 at the money against the Kou cosine benchmark and 1.9 to 2.0 off the money, and 2.00 to 2.01 at every level against the Black–Scholes closed form and the product-characteristic-function benchmark. In the tested range, where the spatial error dominates, the American problem converged at joint order about two (ratios of successive differences between 3.6 and 4.5 on the nested grids away from the exercise boundary, 3.9 to 4.1 on the at-the-money ladder of the Kou put and 3.7 to 3.8 through the announcement step). Refined in time alone, the exact discrete complementarity problem converged at order 1.2 to 1.5, and the early-exercise premium that the empirical study uses converged at joint order 1.7 to 1.8 (Section 5.5). The American put values that [Salmi and Toivanen \(2011\)](#) list for the Kou test case were matched to within 2.3×10^{-5} on the finest grid, and the two-decimal American prices with an announcement of [Leung and Santoli \(2014\)](#) in 54 of 54 cases after extrapolation.

The estimation results set the limits of what a firm’s history can reveal. At the history lengths of the sample a firm’s announcement variance is unbiased but known only to within a large fraction of its size, its double-exponential shape is not identified, and a background jump intensity is not identified from a trailing window of returns, so the models carry no background jumps (Section 6). Over the announcements of the primary sample, the option-implied announcement variance exceeded the realized one by factors of 1.70 [1.60, 1.82] in the development period and 1.14 [1.09, 1.19] in the holdout, ratios that assume calendar-time maturities and a flat background term structure: in trading time they were 1.52 and 1.06, and the decomposition applied after the announcement, where it should average zero, averaged 0.48 and 0.37 of the realized variance, so the development premium survives these checks while the holdout ratio is within their uncertainty of one. The ratio rose with the VIX in the development period, from 1.14 to 2.16 across its terciles, partly through the slope of the background term structure, but hardly in the holdout, where it fell with the firm’s past announcement moves, from 1.39 to 1.04, and was 0.98 in the highest price tercile. The median expected absolute move was 3.91% under the double-exponential law fitted to the firm’s past moves, lower partly by construction, 5.40% under the option-implied Gaussian law and 4.21% for the realized move (Section 6.7).

Priced at the last close before the announcement with the information of that close, Kou-SEJ with the option-implied announcement variance (Kou-SEJ-Q) removed 0.79 [0.76, 0.82] of the absolute dollar error of Black–Scholes with historical volatility in the development sample and 0.80 [0.79, 0.81] in the holdout, and the variant with the historical announcement law (Kou-SEJ-P) removed 0.28 [0.21, 0.34] and 0.42 [0.40, 0.44]. The historical variant uses no option price, so its share is an out-of-sample measure of the announcement jump, net of the variance risk premium and of the noise of short histories. Most of the share of the option-implied variant is the at-the-money level to which it is fitted, which the market’s at-the-money volatility, held flat across strikes, captured as well (0.83 and 0.82). The gains were largest at the money, at the shortest maturities and for the firms with the largest past moves, and the residual of the option-implied models was largest in the money, also when the quotes are binned by the delta at the expiration’s at-the-money volatility rather than their own, and in the top VIX tercile. Away from the money the double-exponential law with the firm’s historical shape did worse

than a Gaussian announcement law with the same level, because its smile lay above the flat level on both sides while the market’s was skewed toward low strikes, and the same held on the recent announcements, on an expiration that neither calibrated model was fitted to; on the few quotes of the main sample outside the calibration expirations it was slightly better, though not significantly. After the announcement the models that share the historical background of Black–Scholes coincided with it, and the others differed from it through their option-implied background, not through the jump; every model with an announcement component predicted an implied-volatility crush of the market’s order, the calibrated option-implied ones (DJ-implied and Kou-SEJ-Q) too large on average. Early exercise mattered for calls with a dividend in their life and not for puts, also after a post-freeze recalibration to American quotes. Of the eight pre-registered predictions, P5 and P7 held in both samples, P3 and P8 in the holdout only, and P1, P2, P4 and P6 failed in both.

The results do not show that earnings-week options are mispriced. They show how much of the distance between a diffusion benchmark and market prices a jump of known timing and random size explains, that for pricing the variance of that jump is better taken from the option chain than from the firm’s history, that its shape is poorly estimated from one firm’s few announcements, and how much of the remainder a variance risk premium would have to account for. The ex-ante model prices of Appendix E will be compared with the outcomes in a later version. The solver, the estimators and the scripts that regenerate every table and figure are available in the public code repository given in the data and code availability statement, and the frozen derived inputs will be archived with a DOI.

Acknowledgments

I thank Malcolm Wardlaw (Terry College of Business, University of Georgia), in whose derivatives course this work became a master’s project, for his guidance and advice throughout, and Steven Malliaris (Terry College of Business, University of Georgia) for reading an earlier draft and for suggesting the empirical validation design of Sections 7 and 8.

Funding

This research received no external funding.

Disclaimer

This article is academic research and does not constitute investment advice or a recommendation to buy or sell any security. The paper evaluates pricing errors against end-of-day quotes; it reports no trading strategy, and the historical comparisons it reports are not indicative of future performance.

Data and code availability

Option quotes are end-of-day bid and ask snapshots from the public DoltHub `options` database of the user `post-no-preference` ([post-no-preference, 2025](#)), table `option_chain`, in the state of 22 April 2025, which I used as a single frozen file (79,614,858 rows, snapshots from 9 February 2019 to 22 April 2025); the database’s own implied-volatility and Greek columns are not used. For the announcements after 22 April 2025 (Section 8.10) I queried the live `option_chain` table of the same database through its SQL interface on 13 October 2026; those chains, the volatility history and the rates of that period are frozen in `data/raw/`, listed with their checksums in `data/raw/MANIFEST_recent.json`, and every query is logged with its time

in `output/recent_api_calls.jsonl`. Earnings announcement dates, with their before-open or after-close timing, come from the DoltHub `earnings` database ([post-no-preference, 2026a](#)), table `earnings_calendar`; daily Treasury constant-maturity yields from the DoltHub `rates` database ([post-no-preference, 2026b](#)), table `us_treasury`, completed with the one-month series DGS1MO from FRED ([Board of Governors of the Federal Reserve System, 2026](#)) where the DoltHub table has no value; and the closing prices of the cross-check of the parity-implied spot, the cash dividends and the splits from the `ohlcv`, `dividend` and `split` tables of the DoltHub `stocks` database ([post-no-preference, 2026c](#)). The VIX close is the FRED series VIXCLS ([Chicago Board Options Exchange, 2026](#)). The earnings, rates and stocks databases and the FRED series were retrieved on 13 October 2026; every retrieved file is frozen in `data/raw/` with its query or URL, retrieval time and checksum in `data/raw/MANIFEST.json`, so that the results do not depend on later revisions of the sources.

The solver, the estimators, the data-processing pipeline and the scripts that regenerate every table, figure and quoted number of this paper from the frozen raw data (`code/scripts/build_all.py` runs them in order and checks that a re-run reproduces the files) are available in the public code repository <https://github.com/GlebSoftware/American-Option-Pricing-Around-Scheduled-Earnings-Announcements>.

A replication archive of the version used for this paper, with the frozen derived inputs (coverage index, parity-implied spots, at-the-money implied volatilities, event sample and priced sample), the pre-registration `output/preregistration.md`, the freeze record `output/spec_freeze.json` and the hash of the ex-ante prices of Appendix E, will be deposited at Zenodo, and its DOI will be added in a later version of this paper. The options database itself is distributed by its maintainer on DoltHub and is not part of the archive; the archive will record the snapshot dates, the row count and the SHA-256 checksum of the file I used, so that the same state of the database can be identified. The derived data sets and the other downloads will be included as far as the terms of their sources allow redistribution; for any that are not, the archive will keep the query and the retrieval record of the manifest, and the scripts regenerate the derived data from the sources.

References

- Y. Aït-Sahalia. Disentangling diffusion from jumps. *Journal of Financial Economics*, 74(3): 487–528, 2004. doi: 10.1016/j.jfineco.2003.09.005.
- Y. Aït-Sahalia and J. Jacod. Testing for jumps in a discretely observed process. *Annals of Statistics*, 37(1):184–222, 2009. doi: 10.1214/07-AOS568.
- L. Alexiou, A. Goyal, A. Kostakis, and L. Rompolis. Pricing event risk: Evidence from concave implied volatility curves. *Review of Finance*, 29(4):963–1007, 2025. doi: 10.1093/rof/rfaf016.
- A. Almendral and C. W. Oosterlee. Numerical valuation of options with jumps in the underlying. *Applied Numerical Mathematics*, 53(1):1–18, 2005. doi: 10.1016/j.apnum.2004.08.037.
- K. I. Amin. Jump diffusion option valuation in discrete time. *Journal of Finance*, 48(5): 1833–1863, 1993. doi: 10.1111/j.1540-6261.1993.tb05130.x.
- L. Andersen and J. Andreasen. Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing. *Review of Derivatives Research*, 4(3):231–262, 2000. doi: 10.1023/A: 1011354913068.
- U. M. Ascher, S. J. Ruuth, and B. T. R. Wetton. Implicit-explicit methods for time-dependent partial differential equations. *SIAM Journal on Numerical Analysis*, 32(3):797–823, 1995. doi: 10.1137/0732037.

- G. Bakshi, C. Cao, and Z. Chen. Empirical performance of alternative option pricing models. *Journal of Finance*, 52(5):2003–2049, 1997. doi: 10.1111/j.1540-6261.1997.tb02749.x.
- O. E. Barndorff-Nielsen and N. Shephard. Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics*, 4(1):1–30, 2006. doi: 10.1093/jjfinec/nbi022.
- M. E. Barth and E. C. So. Non-diversifiable volatility risk and risk premiums at earnings announcements. *The Accounting Review*, 89(5):1579–1607, 2014. doi: 10.2308/accr-50758.
- A. Beber and M. W. Brandt. The effect of macroeconomic news on beliefs and preferences: Evidence from the options market. *Journal of Monetary Economics*, 53(8):1997–2039, 2006. doi: 10.1016/j.jmoneco.2006.05.013.
- Y. Z. Bergman, B. D. Grundy, and Z. Wiener. General properties of option prices. *Journal of Finance*, 51(5):1573–1610, 1996. doi: 10.1111/j.1540-6261.1996.tb05218.x.
- A. Berman and R. J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. Academic Press, New York, 1979.
- F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973. doi: 10.1086/260062.
- Board of Governors of the Federal Reserve System. Market yield on U.S. Treasury securities at 1-month constant maturity, quoted on an investment basis (DGS1MO). FRED, Federal Reserve Bank of St. Louis, <https://fred.stlouisfed.org/series/DGS1MO>, 2026. Daily, percent; retrieved 13 October 2026.
- M. J. Brennan and E. S. Schwartz. The valuation of American put options. *Journal of Finance*, 32(2):449–462, 1977. doi: 10.2307/2326779.
- M. Briani, R. Natalini, and G. Russo. Implicit-explicit numerical schemes for jump-diffusion processes. *Calcolo*, 44(1):33–57, 2007. doi: 10.1007/s10092-007-0128-x.
- M. Broadie and J. Detemple. American option valuation: New bounds, approximations, and a comparison of existing methods. *Review of Financial Studies*, 9(4):1211–1250, 1996.
- A. C. Cameron, J. B. Gelbach, and D. L. Miller. Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2):238–249, 2011. doi: 10.1198/jbes.2010.07136.
- Chicago Board Options Exchange. CBOE Volatility Index: VIX (VIXCLS). FRED, Federal Reserve Bank of St. Louis, <https://fred.stlouisfed.org/series/VIXCLS>, 2026. Daily close; retrieved 13 October 2026.
- P. Christoffersen and K. Jacobs. The importance of the loss function in option valuation. *Journal of Financial Economics*, 72(2):291–318, 2004. doi: 10.1016/j.jfineco.2003.02.001.
- R. Cont and P. Tankov. *Financial Modelling with Jump Processes*. Chapman & Hall/CRC, Boca Raton, 2004.
- R. Cont and E. Voltchkova. A finite difference scheme for option pricing in jump diffusion and exponential Lévy models. *SIAM Journal on Numerical Analysis*, 43(4):1596–1626, 2005. doi: 10.1137/S0036142903436186.
- R. W. Cottle, J.-S. Pang, and R. E. Stone. *The Linear Complementarity Problem*. Academic Press, Boston, 1992.

- C. W. Cryer. The solution of a quadratic programming problem using systematic overrelaxation. *SIAM Journal on Control*, 9(3):385–392, 1971. doi: 10.1137/0309028.
- Y. d’Halluin, P. A. Forsyth, and G. Labahn. A penalty method for American options with jump diffusion processes. *Numerische Mathematik*, 97(2):321–352, 2004. doi: 10.1007/s00211-003-0511-8.
- Y. d’Halluin, P. A. Forsyth, and K. R. Vetzal. Robust numerical methods for contingent claims under jump diffusion processes. *IMA Journal of Numerical Analysis*, 25(1):87–112, 2005. doi: 10.1093/imanum/drh011.
- F. X. Diebold and R. S. Mariano. Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263, 1995. doi: 10.1080/07350015.1995.10524599.
- M. W. M. Donders, R. Kouwenberg, and T. C. F. Vorst. Options and earnings announcements: An empirical study of volatility, trading volume, open interest and liquidity. *European Financial Management*, 6(2):149–171, 2000. doi: 10.1111/1468-036X.00118.
- A. Dubinsky and M. Johannes. Earnings announcements and equity options. Working paper, Graduate School of Business, Columbia University, Sept. 2004. First draft November 2003. Later versions circulated from 2006 as “Anticipated uncertainty, earning announcements and equity options” and “Fundamental uncertainty, earning announcements and equity options”.
- A. Dubinsky, M. Johannes, A. Kaeck, and N. J. Seeger. Option pricing of earnings announcement risks. *Review of Financial Studies*, 32(2):646–687, 2019. doi: 10.1093/rfs/hhy060.
- L. H. Ederington and J. H. Lee. The creation and resolution of market uncertainty: The impact of information releases on implied volatility. *Journal of Financial and Quantitative Analysis*, 31(4):513–539, 1996. doi: 10.2307/2331358.
- E. Ekström and J. Tysk. Properties of option prices in models with jumps. *Mathematical Finance*, 17(3):381–397, 2007. doi: 10.1111/j.1467-9965.2007.00308.x.
- N. El Karoui, M. Jeanblanc-Picqué, and S. E. Shreve. Robustness of the Black and Scholes formula. *Mathematical Finance*, 8(2):93–126, 1998. doi: 10.1111/1467-9965.00047.
- F. Fang and C. W. Oosterlee. A novel pricing method for European options based on Fourier-cosine series expansions. *SIAM Journal on Scientific Computing*, 31(2):826–848, 2008. doi: 10.1137/080718061.
- F. Fang and C. W. Oosterlee. Pricing early-exercise and discrete barrier options by Fourier-cosine series expansions. *Numerische Mathematik*, 114(1):27–62, 2009. doi: 10.1007/s00211-009-0252-4.
- L. Feng and V. Linetsky. Pricing options in jump-diffusion models: An extrapolation approach. *Operations Research*, 56(2):304–325, 2008. doi: 10.1287/opre.1070.0419.
- P. A. Forsyth and K. R. Vetzal. Quadratic convergence for valuing American options using a penalty method. *SIAM Journal on Scientific Computing*, 23(6):2095–2122, 2002. doi: 10.1137/S1064827500382324.
- J. Frank, W. Hundsdorfer, and J. G. Verwer. On the stability of implicit-explicit linear multistep methods. *Applied Numerical Mathematics*, 25(2–3):193–205, 1997. doi: 10.1016/S0168-9274(97)00059-7.
- C. Gao, Y. Xing, and X. Zhang. Anticipating uncertainty: Straddles around earnings announcements. *Journal of Financial and Quantitative Analysis*, 53(6):2587–2617, 2018. doi: 10.1017/S0022109018000285.

- R. Giacomini and H. White. Tests of conditional predictive ability. *Econometrica*, 74(6): 1545–1578, 2006. doi: 10.1111/j.1468-0262.2006.00718.x.
- M. B. Giles and R. Carter. Convergence analysis of Crank–Nicolson and Rannacher time-marching. *Journal of Computational Finance*, 9(4):89–112, 2006. doi: 10.21314/JCF.2006.152.
- P. R. Hansen, A. Lunde, and J. M. Nason. The model confidence set. *Econometrica*, 79(2): 453–497, 2011. doi: 10.3982/ECTA5771.
- C. R. Harvey, Y. Liu, and H. Zhu. ... and the cross-section of expected returns. *Review of Financial Studies*, 29(1):5–68, 2016. doi: 10.1093/rfs/hhv059.
- E. G. Haug, J. Haug, and A. Lewis. Back to basics: A new approach to the discrete dividend problem. *Wilmott Magazine*, pages 37–47, Sept. 2003. doi: 10.1002/wilm.42820030514.
- J. E. Hilliard and A. Schwartz. Pricing European and American derivatives under a jump-diffusion process: A bivariate tree approach. *Journal of Financial and Quantitative Analysis*, 40(3):671–691, 2005. doi: 10.1017/S0022109000001915.
- P. Honoré. Pitfalls in estimating jump-diffusion models. Working paper, Aarhus School of Business, 1998. SSRN abstract 61998, doi: 10.2139/ssrn.61998.
- W. Hundsdorfer and J. G. Verwer. *Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations*, volume 33 of *Springer Series in Computational Mathematics*. Springer, Berlin, 2003.
- S. Ikonen and J. Toivanen. Operator splitting methods for American option pricing. *Applied Mathematics Letters*, 17(7):809–814, 2004. doi: 10.1016/j.aml.2004.06.010.
- S. Ikonen and J. Toivanen. Operator splitting methods for pricing American options under stochastic volatility. *Numerische Mathematik*, 113(2):299–324, 2009. doi: 10.1007/s00211-009-0227-5.
- K. R. Jackson, S. Jaimungal, and V. Surkov. Fourier space time-stepping for option pricing with Lévy models. *Journal of Computational Finance*, 12(2):1–29, 2008. doi: 10.21314/JCF.2008.178.
- P. Jaillet, D. Lamberton, and B. Lapeyre. Variational inequalities and the pricing of American options. *Acta Applicandae Mathematicae*, 21(3):263–289, 1990. doi: 10.1007/BF00047211.
- B. Kelly, L. Pástor, and P. Veronesi. The price of political uncertainty: Theory and evidence from the option market. *Journal of Finance*, 71(5):2417–2480, 2016. doi: 10.1111/jofi.12406.
- S. G. Kou. A jump-diffusion model for option pricing. *Management Science*, 48(8):1086–1101, 2002. doi: 10.1287/mnsc.48.8.1086.166.
- S. G. Kou and H. Wang. Option pricing under a double exponential jump diffusion model. *Management Science*, 50(9):1178–1192, 2004. doi: 10.1287/mnsc.1030.0163.
- Y. Kwon and Y. Lee. A second-order finite difference method for option pricing under jump-diffusion models. *SIAM Journal on Numerical Analysis*, 49(6):2598–2617, 2011a. doi: 10.1137/090777529.
- Y. Kwon and Y. Lee. A second-order tridiagonal method for American options under jump-diffusion models. *SIAM Journal on Scientific Computing*, 33(4):1860–1872, 2011b. doi: 10.1137/100806552.

- J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright. Convergence properties of the Nelder–Mead simplex method in low dimensions. *SIAM Journal on Optimization*, 9(1):112–147, 1998. doi: 10.1137/S1052623496303470.
- S. S. Lee and P. A. Mykland. Jumps in financial markets: A new nonparametric test and jump dynamics. *Review of Financial Studies*, 21(6):2535–2563, 2008. doi: 10.1093/rfs/hhm056.
- C. C. W. Leentvaar and C. W. Oosterlee. American options with discrete dividends solved by highly accurate discretizations. In A. Di Bucchianico, R. M. M. Mattheij, and M. A. Peletier, editors, *Progress in Industrial Mathematics at ECMI 2004*, Mathematics in Industry, pages 427–431. Springer, Berlin, Heidelberg, 2006. doi: 10.1007/3-540-28073-1_67.
- T. Leung and M. Santoli. Accounting for earnings announcements in the pricing of equity options. *Journal of Financial Engineering*, 1(4):1450031, 2014. doi: 10.1142/S2345768614500317. arXiv:1412.8414v2; Table 4 of this version gives the American prices with an announcement.
- C. Mancini. Non-parametric threshold estimation for models with stochastic diffusion coefficient and jumps. *Scandinavian Journal of Statistics*, 36(2):270–296, 2009. doi: 10.1111/j.1467-9469.2008.00622.x.
- R. C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1–2):125–144, 1976. doi: 10.1016/0304-405X(76)90022-2.
- D. Muravyev and N. D. Pearson. Options trading costs are lower than you think. *Review of Financial Studies*, 33(11):4973–5014, 2020. doi: 10.1093/rfs/hhaa010.
- J. A. Nelder and R. Mead. A simplex method for function minimization. *The Computer Journal*, 7(4):308–313, 1965. doi: 10.1093/comjnl/7.4.308.
- J. M. Patell and M. A. Wolfson. Anticipated information releases reflected in call option prices. *Journal of Accounting and Economics*, 1(2):117–140, 1979. doi: 10.1016/0165-4101(79)90003-X.
- J. M. Patell and M. A. Wolfson. The ex ante and ex post price effects of quarterly earnings announcements reflected in option and stock prices. *Journal of Accounting Research*, 19(2):434–458, 1981. doi: 10.2307/2490874.
- G. Peskir and A. Shiryaev. *Optimal Stopping and Free-Boundary Problems*. Lectures in Mathematics, ETH Zürich. Birkhäuser, Basel, 2006.
- H. Pham. Optimal stopping, free boundary, and American option in a jump-diffusion model. *Applied Mathematics and Optimization*, 35(2):145–164, 1997. doi: 10.1007/BF02683325.
- M. Piazzesi. An econometric model of the yield curve with macroeconomic jump effects. NBER Working Paper 8246, National Bureau of Economic Research, 2001.
- M. Piazzesi. Bond yields and the Federal Reserve. *Journal of Political Economy*, 113(2):311–344, 2005. doi: 10.1086/427466.
- D. M. Pooley, K. R. Vetzal, and P. A. Forsyth. Convergence remedies for non-smooth payoffs in option pricing. *Journal of Computational Finance*, 6(4):25–40, 2003. doi: 10.21314/JCF.2003.101.
- post-no-preference. options: US stock option chains with Greeks and volatility history (DoltHub database). <https://www.dolthub.com/repositories/post-no-preference/options>, 2025. Snapshot of 22 April 2025 used in this paper; end-of-day quotes 2019-02-09 to 2025-04-22.

- post-no-preference. earnings: US company earnings calendar, estimates and financial statements (DoltHub database). <https://www.dolthub.com/repositories/post-no-preference/earnings>, 2026a. Table `earnings_calendar` (symbol, date, before-open/after-close flag); retrieved 13 October 2026.
- post-no-preference. rates: US Treasury constant-maturity yields (DoltHub database). <https://www.dolthub.com/repositories/post-no-preference/rates>, 2026b. Table `us_treasury`; retrieved 13 October 2026.
- post-no-preference. stocks: US stock prices, dividends and splits (DoltHub database). <https://www.dolthub.com/repositories/post-no-preference/stocks>, 2026c. Tables `ohlcv`, `dividend` and `split`; retrieved 13 October 2026.
- C. A. Ramezani and Y. Zeng. Maximum likelihood estimation of the double exponential jump-diffusion process. *Annals of Finance*, 3(4):487–507, 2007. doi: 10.1007/s10436-006-0062-y.
- R. Rannacher. Finite element solution of diffusion problems with irregular data. *Numerische Mathematik*, 43(2):309–327, 1984. doi: 10.1007/BF01390130.
- S. Salmi and J. Toivanen. An iterative method for pricing American options under jump-diffusion models. *Applied Numerical Mathematics*, 61(7):821–831, 2011. doi: 10.1016/j.apnum.2011.02.002.
- S. Salmi and J. Toivanen. Comparison and survey of finite difference methods for pricing American options under finite activity jump-diffusion models. *International Journal of Computer Mathematics*, 89(9):1112–1134, 2012. doi: 10.1080/00207160.2012.669475.
- S. Salmi and J. Toivanen. IMEX schemes for pricing options under jump-diffusion models. *Applied Numerical Mathematics*, 84:33–45, 2014. doi: 10.1016/j.apnum.2014.05.007.
- J. Toivanen. Numerical valuation of European and American options under Kou’s jump-diffusion model. *SIAM Journal on Scientific Computing*, 30(4):1949–1970, 2008. doi: 10.1137/060674697.
- J. Toivanen. A high-order front-tracking finite difference method for pricing American options under jump-diffusion models. *Journal of Computational Finance*, 13(3):61–79, 2010. doi: 10.21314/JCF.2010.222.
- R. E. Whaley and J. K. Cheung. Anticipation of quarterly earnings announcements: A test of option market efficiency. *Journal of Accounting and Economics*, 4(2):57–83, 1982. doi: 10.1016/0165-4101(82)90014-3.

A Proofs and derivations

A.1 The compensators κ and κ_E

For the density (2),

$$\mathbb{E}[e^Y] = \int_0^\infty p \eta_1 e^{-(\eta_1-1)y} dy + \int_{-\infty}^0 (1-p) \eta_2 e^{(\eta_2+1)y} dy = \frac{p \eta_1}{\eta_1 - 1} + \frac{(1-p) \eta_2}{\eta_2 + 1},$$

where the first integral converges iff $\eta_1 > 1$ and the second iff $\eta_2 > -1$, a weaker requirement than the $\eta_2 > 0$ needed for f_Y to be a density. Subtracting $1 = p + (1-p)$ term by term, $\kappa = p(\frac{\eta_1}{\eta_1-1} - 1) + (1-p)(\frac{\eta_2}{\eta_2+1} - 1) = \frac{p}{\eta_1-1} - \frac{1-p}{\eta_2+1}$, which is (3). The compensation formula for the compound Poisson process gives $\mathbb{E}[\sum_{k \leq N_t} (e^{Y_k} - 1)] = \lambda \kappa t$, so $\sum_{k \leq N_t} (e^{Y_k} - 1) - \lambda \kappa t$ is a martingale (Cont and Tankov, 2004, Chapter 3); this is why the drift in (1) carries $-\lambda \kappa$. The same computation with $(p_E, \eta_{1,E}, \eta_{2,E})$ gives $\mathbb{E}[e^Z] = 1 + \frac{p_E}{\eta_{1,E}-1} - \frac{1-p_E}{\eta_{2,E}+1}$, a positive number because $e^Z > 0$, so $\kappa_E = \ln \mathbb{E}[e^Z]$ in (5) is well defined and $\mathbb{E}[e^{Z-\kappa_E}] = 1$, which is (6).

A.2 Proof of Proposition 2.3

Exponential form of the price. Let $X_t = \mu t + \sigma W_t + \sum_{k \leq N_t} Y_k$ with μ as in (7). By Itô's formula for semimartingales with jumps (Cont and Tankov, 2004, Chapter 8), $de^{X_t} = e^{X_t-}[(\mu + \frac{1}{2}\sigma^2)dt + \sigma dW_t + d(\sum_{k \leq N_t}(e^{Y_k} - 1))]$, and $\mu + \frac{1}{2}\sigma^2 = r - q - \lambda\kappa$, so $S_0 e^{X_t}$ solves (1); the solution of this linear equation is unique, hence $S_t = S_0 e^{X_t}$ for $t < t_E$. At t_E the price is multiplied by $e^{Z - \kappa_E}$ and the same dynamics restart from S_{t_E} . Therefore, for $t_E \in (0, T)$ and with the convention of Section 2 for $t_E \notin (0, T)$,

$$\varrho_T = \ln \frac{S_T}{S_0} = \mu T + \sigma W_T + \sum_{k=1}^{N_T} Y_k + (Z - \kappa_E) \mathbf{1}_{\{t_E \in (0, T)\}}. \quad (50)$$

(i) and (v). The four summands in (50) are independent, so ϕ_{ϱ_T} is the product of their characteristic functions: $e^{iv\mu T}$; $e^{-\sigma^2 v^2 T/2}$; conditioning on N_T ,

$$\mathbb{E}[e^{iv \sum_{k \leq N_T} Y_k}] = \sum_{n \geq 0} e^{-\lambda T} \frac{(\lambda T)^n}{n!} \phi_Y(v)^n = \exp\{\lambda T(\phi_Y(v) - 1)\};$$

and $e^{-iv\kappa_E} \phi_Z(v)$. This is (9), in which t_E does not appear except through the indicator; hence (v). At $v = -i$, $\phi_Y(-i) = \mathbb{E}[e^Y] = 1 + \kappa$, so the exponent equals $T(\mu + \frac{1}{2}\sigma^2 + \lambda\kappa) = (r - q)T$, and $e^{-\kappa_E} \phi_Z(-i) = e^{-\kappa_E} \mathbb{E}[e^Z] = 1$. The one-sided integrals $\int_0^\infty e^{ivy} p\eta_1 e^{-\eta_1 y} dy = p\eta_1/(\eta_1 - iv)$ and $\int_{-\infty}^0 e^{ivy} (1 - p)\eta_2 e^{\eta_2 y} dy = (1 - p)\eta_2/(\eta_2 + iv)$ give (8).

(ii). The cumulant generating function is $\ln \phi_{\varrho_T}(v) = T[iv\mu - \frac{1}{2}\sigma^2 v^2 + \lambda(\phi_Y(v) - 1)] - iv\kappa_E + \ln \phi_Z(v)$, and cumulants of independent summands add. Expanding $\phi_Y(v) - 1 = \sum_{n \geq 1} \mathbb{E}[Y^n](iv)^n/n!$ with $\mathbb{E}[Y^n] = \int_0^\infty y^n p\eta_1 e^{-\eta_1 y} dy + \int_{-\infty}^0 y^n (1 - p)\eta_2 e^{\eta_2 y} dy = n!(p\eta_1^{-n} + (1 - p)(-1)^n \eta_2^{-n})$ gives the Kou cumulants $T(\mu + \lambda \mathbb{E}[Y])$, $T(\sigma^2 + \lambda \mathbb{E}[Y^2])$ and $\lambda T \mathbb{E}[Y^n]$ for $n \geq 3$, which are the first terms in (10). The random variable $Z - \kappa_E$ is a single jump, not a compound Poisson sum, so its cumulants are those of Z shifted in the first: $k_1 = m_1 - \kappa_E$, and for $n \geq 2$ the moment-cumulant relations $k_2 = m_2 - m_1^2$, $k_4 = m_4 - 4m_3 m_1 - 3m_2^2 + 12m_2 m_1^2 - 6m_1^4$ with $m_n = \mathbb{E}[Z^n]$ computed as for Y .

(iii). Let $\lambda = 0$ and condition on $Z = z$. By (50), $\varrho_T = (r - q - \frac{1}{2}\sigma^2)T + \sigma W_T + z - \kappa_E$, so S_T has the Black-Scholes terminal law with spot $S_0 e^{z - \kappa_E}$, volatility σ and maturity T , and $\mathbb{E}[e^{-rT} \varphi(S_T) | Z = z] = V^{\text{BS}}(S_0 e^{z - \kappa_E}, 0)$. Integrating against f_E gives (11). For $\lambda > 0$ the conditional law is the Kou terminal law with the shifted spot.

(iv). $\mathbb{E}[e^{-rT}(S_T - K)] = e^{-rT}(S_0 \phi_{\varrho_T}(-i) - K) = S_0 e^{-qT} - K e^{-rT}$ by (i), and $(s - K)^+ - (K - s)^+ = s - K$. $\square$

A.3 The PIDE and the jump condition

For $V \in C^{2,1}$ the generator of the Kou process acting on functions of S is $\mathcal{L}_S V = \frac{1}{2}\sigma^2 S^2 V_{SS} + (r - q - \lambda\kappa)SV_S + \lambda \int_{\mathbb{R}} [V(Se^y) - V(S)]f_Y(y) dy$, and the martingale property of $e^{-rt}V(S_t, t)$ on each interval without announcement gives $\partial_t V + \mathcal{L}_S V - rV = 0$ (Cont and Tankov, 2004, Chapter 12). With $x = \ln S$, $\tau = T - t$ and $u(x, \tau) = V(e^x, T - \tau)$ one has $SV_S = u_x$, $S^2 V_{SS} = u_{xx} - u_x$ and $\partial_t = -\partial_\tau$, so

$$\partial_\tau u = \frac{1}{2}\sigma^2(u_{xx} - u_x) + (r - q - \lambda\kappa)u_x - ru + \lambda \int_{\mathbb{R}} u(x + y, \tau) f_Y(y) dy - \lambda u,$$

which is (12) after collecting the u_x and u terms. The jump condition: by the Markov property and (4), for a claim that is not exercised at the announcement, $V(S, t_E^-) = \mathbb{E}[V(S_{t_E}, t_E) | S_{t_E^-} = S] = \int_{\mathbb{R}} V(Se^{z - \kappa_E}, t_E) f_E(z) dz$, no discount factor being involved; in (x, τ) this is $u(x, \tau_E^+) = \int u(x + z - \kappa_E, \tau_E^-) f_E(z) dz$, and the substitution $y = z - \kappa_E$ gives the second form in (13). Finally $\mathcal{P}_E \mathbf{1} = \int f_E = 1$ and $\mathcal{P}_E e^x = e^x \int e^{z - \kappa_E} f_E(z) dz = e^x$ by (6).

A.4 Proof of Lemma 3.2

Let $L_t = \ln(S_t/S_0)$ for $t \in [0, T]$, which by the argument that gives (50) equals $X_t + (Z - \kappa_E)\mathbf{1}_{\{t \geq t_E\}}$ when $t_E \in (0, T)$ and X_t otherwise. The process L has independent increments and does not depend on S_0 , so if $S_t = S$ then $S_s = Se^{L_s - L_t}$ for $s \geq t$, and the increments $(L_s - L_t)_{s \geq t}$ are independent of $\mathcal{F}_t$. Information in $\mathcal{F}_t$ beyond S_t is therefore independent of the future payoff and cannot raise the value, so the supremum in (14) over $(\mathcal{F}_t)$ -stopping times equals the supremum over the stopping times of the filtration generated by the increments (Peskir and Shiryaev, 2006, Chapter I), and (14) can be written as

$$V(S, t) = \sup_{\vartheta \in \mathcal{T}_{t,T}^L} \mathbb{E} \left[e^{-r(\vartheta-t)} \varphi(Se^{L_\vartheta - L_t}) \right], \quad (51)$$

where $\mathcal{T}_{t,T}^L$ is the set of stopping times of the filtration generated by $(L_s - L_t)_{s \geq t}$ with values in $[t, T]$; this set does not depend on S . Each expectation is finite: $\varphi(s) \leq a + bs$ and $e^{-(r-q)s}e^{L_s}$ is a martingale, so $\mathbb{E}[e^{L_\vartheta - L_t}] \leq e^{(r-q)(T-t)}$ for every $\vartheta \in \mathcal{T}_{t,T}^L$. For fixed ϑ and on each path the map $S \mapsto e^{-r(\vartheta-t)}\varphi(Se^{L_\vartheta - L_t})$ is convex, as the composition of the convex φ with a linear map; expectation preserves convexity; and the supremum of a family of convex functions is convex. Thus $V(\cdot, t)$ is convex. The European value is the case $\vartheta \equiv T$, and $C_E(S) = \mathbb{E}[V(Se^{Z - \kappa_E}, t_E)]$ is convex as a mixture of convex functions of S . $\square$

A.5 Proof of Theorem 3.3

Write $\tilde{S} = Se^{Z - \kappa_E}$, so that $\mathbb{E}[\tilde{S}] = S$ by (6). (i). Stopping immediately is admissible in (14), so $V(s, t_E) \geq \varphi(s)$ for every $s > 0$, and therefore $C_E(S) = \mathbb{E}[V(\tilde{S}, t_E)] \geq \mathbb{E}[\varphi(\tilde{S})]$. Jensen's inequality for the convex φ gives $\mathbb{E}[\varphi(\tilde{S})] \geq \varphi(\mathbb{E}\tilde{S}) = \varphi(S)$. Since $C_E \geq \varphi$, the maximum in (18) is C_E ; this part holds for every law of Z . For the strictness, let the law of Z have full support, so that $\tilde{S}$ falls in every open subinterval of $(0, \infty)$ with positive probability; for $\text{DE}(p_E, \eta_{1,E}, \eta_{2,E})$ with $0 < p_E < 1$ the density f_E is positive on $\mathbb{R}$, so $\tilde{S}$ has a density that is positive on $(0, \infty)$. Let ℓ be a supporting line of φ at S , so that $\varphi \geq \ell$ on $(0, \infty)$ and $\varphi(S) = \ell(S)$; then $\mathbb{E}[\varphi(\tilde{S})] - \varphi(S) = \mathbb{E}[\varphi(\tilde{S}) - \ell(\tilde{S})]$ because $\mathbb{E}[\ell(\tilde{S})] = \ell(\mathbb{E}\tilde{S}) = \ell(S)$. If φ is not affine, $\varphi - \ell$ is a nonnegative continuous function that is not identically zero, so it is positive on some open interval, which $\tilde{S}$ hits with positive probability; hence $\mathbb{E}[\varphi(\tilde{S}) - \ell(\tilde{S})] > 0$. For the put $(K - s)^+$ the supporting line at $S < K$ is $K - s$ and the two functions differ on (K, ∞) ; at $S > K$ it is 0 and they differ on $(0, K)$; at $S = K$ every supporting line has a slope in $[-1, 0]$ and differs from the payoff on at least one side. The call is symmetric. Since then $C_E(S) > \varphi(S)$ for every $S > 0$, $\mathcal{E}_{t_E} = \emptyset$. (ii). By Lemma 3.2 the function $V(\cdot, t_E)$ is convex, so Jensen's inequality gives $C_E(S) = \mathbb{E}[V(\tilde{S}, t_E)] \geq V(\mathbb{E}\tilde{S}, t_E) = V(S, t_E)$. $\square$

A.6 Proof of Corollary 3.4

Structure of $\mathcal{E}_t$. Fix $t < T$. For $S \geq K$, $V(S, t) > 0 = (K - S)^+$ because $\sigma > 0$ makes the put finish in the money with positive probability, so $\mathcal{E}_t \subset (0, K)$. On $(0, K)$ the function $g(S) = V(S, t) - (K - S)$ is convex by Lemma 3.2, nonnegative, and satisfies $0 \leq g(S) \leq S$ because $V \leq K$; hence $g(S) \rightarrow 0$ as $S \downarrow 0$. If $g(S_1) = 0$ for some $S_1 \in (0, K)$, convexity gives $g(S) \leq \frac{S_1 - S}{S_1 - \varepsilon}g(\varepsilon) + \frac{S - \varepsilon}{S_1 - \varepsilon}g(S_1) \rightarrow 0$ as $\varepsilon \downarrow 0$ for every $S \in (0, S_1)$, so $g = 0$ on $(0, S_1]$. Therefore $\mathcal{E}_t$ is empty or an interval $(0, S^*(t)]$, closed at $S^*(t)$ by continuity of g .

Collapse. Since $\mathcal{E}_t \subset (0, K)$, it suffices to consider $S \in (0, K)$. Fix such an S and $t < t_E$, and let $\Delta = t_E - t$. The strategy “hold until t_E and exercise immediately after the announcement” is admissible in (14) and gives

$$V(S, t) \geq e^{-r\Delta} \mathbb{E} \left[(K - Se^{X_\Delta} e^{Z - \kappa_E})^+ \right], \quad (52)$$

where X_Δ is the Kou log increment over $[t, t_E]$, independent of Z , with characteristic function $\exp\{\Delta[\nu\mu - \frac{1}{2}\sigma^2\nu^2 + \lambda(\phi_Y(\nu) - 1)]\} \rightarrow 1$ as $\Delta \downarrow 0$. Thus $X_\Delta \rightarrow 0$ in distribution, $(X_\Delta, Z) \rightarrow (0, Z)$ in distribution, and since $(x, z) \mapsto (K - Se^{x+z-\kappa_E})^+$ is bounded and continuous, the expectation in (52) converges to $\mathbb{E}[(K - Se^{Z-\kappa_E})^+]$. Since $(K - s)^+ = K - s + (s - K)^+$ and $\mathbb{E}[Se^{Z-\kappa_E}] = S$, $\mathbb{E}[(K - Se^{Z-\kappa_E})^+] = K - S + \mathbb{E}[(Se^{Z-\kappa_E} - K)^+] > K - S$, the inequality being strict because $p_E > 0$ makes Z unbounded above, so that $Se^{Z-\kappa_E} > K$ with positive probability. Hence $\liminf_{t \uparrow t_E} [V(S, t) - (K - S)] \geq \mathbb{E}[(K - Se^{Z-\kappa_E})^+] - (K - S) > 0$, so there is $\Delta(S) > 0$ with $V(S, t) > K - S$ for all $t \in (t_E - \Delta(S), t_E)$. For such t , $S \notin \mathcal{E}_t$, and by the interval structure $S^*(t) < S$. Since $S \in (0, K)$ was arbitrary, $\limsup_{t \uparrow t_E} S^*(t) \leq S$ for every $S \in (0, K)$, which is (20). The same argument with $(s - K)^+$, with the bound $V \leq S$ in place of $V \leq K$, with $p_E < 1$ in place of $p_E > 0$ and with the uniform integrability of $Se^{X_\Delta + Z - \kappa_E}$ in place of boundedness, shows that the American call with $q > 0$ is not exercised shortly before the announcement. $\square$

The proof isolates the economics: holding through the announcement forgoes at most $K(1 - e^{-r\Delta})$ of interest, which vanishes with Δ , and gains the announcement optionality $\mathbb{E}[(K - Se^{Z-\kappa_E})^+] - (K - S)$, which does not. Nothing is said about how fast $S^*(t)$ falls; Figure 2 shows the observed shape.

A.7 Proof of Proposition 3.5

Let $0 < t'_E < t_E < T$ and realize both models on one probability space with the same Kou process X and the same announcement variable Z : model 1 has log price $L_t^{(1)} = X_t + (Z - \kappa_E)\mathbf{1}_{\{t \geq t'_E\}}$ and model 2 has $L_t^{(2)} = X_t + (Z - \kappa_E)\mathbf{1}_{\{t \geq t_E\}}$. Let $\mathcal{F}_t^{(1)} = \sigma(X_s, s \leq t) \vee \sigma(Z)$ for $t \geq t'_E$ and $\mathcal{F}_t^{(1)} = \sigma(X_s, s \leq t)$ before, and define $\mathcal{F}^{(2)}$ likewise with t_E ; then $\mathcal{F}_t^{(2)} \subset \mathcal{F}_t^{(1)}$ for all t , so every $\mathcal{F}^{(2)}$ -stopping time is an $\mathcal{F}^{(1)}$ -stopping time. Write $V^{(1)} = V^{t'_E}$ and $V^{(2)} = V^{t_E}$ for the two value functions and V^X for the American value in the Kou model without announcement.

Step 1 (after t_E). For $t \geq t_E$ neither model has a pending announcement, so $V^{(1)}(\cdot, t) = V^{(2)}(\cdot, t) = V^X(\cdot, t)$; likewise $V^{(1)}(\cdot, t) = V^X(\cdot, t)$ for $t \geq t'_E$.

Step 2 (at t'_E). In model 1 the announcement occurs at t'_E , and by Theorem 3.3 $V^{(1)}(S, t'_E) = \mathbb{E}[V^X(Se^{Z-\kappa_E}, t'_E)]$. In model 2 nothing happens at t'_E , and by (51)

$$V^{(2)}(S, t'_E) = \sup_{\vartheta \in \mathcal{T}^{(2)}} \mathbb{E}[e^{-r(\vartheta - t'_E)} \varphi(S_\vartheta^{(2)})], \quad S_\vartheta^{(2)} = Se^{X_\vartheta - X_{t'_E}} e^{(Z - \kappa_E)\mathbf{1}_{\{\vartheta \geq t_E\}}},$$

where $\mathcal{T}^{(2)}$ is the set of $\mathcal{F}^{(2)}$ -stopping times with values in $[t'_E, T]$. Define $\hat{S}_\vartheta = Se^{X_\vartheta - X_{t'_E}} e^{Z - \kappa_E}$, the price in a model in which the announcement has already happened at t'_E . Fix $\vartheta \in \mathcal{T}^{(2)}$. On $\{\vartheta \geq t_E\}$, $S_\vartheta^{(2)} = \hat{S}_\vartheta$. On $\{\vartheta < t_E\}$ the random variables ϑ and $S_\vartheta^{(2)}$ are measurable with respect to $\mathcal{G} = \sigma(X_s, s < t_E)$, because $\{\vartheta < t_E\} = \bigcup_{s < t_E} \{\vartheta \leq s\}$ and $\{\vartheta \leq s\} \in \mathcal{F}_s^{(2)} = \sigma(X_u, u \leq s)$ for $s < t_E$; and Z is independent of $\mathcal{G}$. Hence, conditioning on $\mathcal{G}$ and applying Jensen's inequality to the convex φ with $\mathbb{E}[e^{Z-\kappa_E}] = 1$,

$$\mathbb{E}[\varphi(\hat{S}_\vartheta)\mathbf{1}_{\{\vartheta < t_E\}} | \mathcal{G}] = \mathbf{1}_{\{\vartheta < t_E\}} \int_{\mathbb{R}} \varphi(S_\vartheta^{(2)} e^{z - \kappa_E}) f_E(z) dz \geq \mathbf{1}_{\{\vartheta < t_E\}} \varphi(S_\vartheta^{(2)}).$$

Multiplying by the $\mathcal{G}$ -measurable discount factor and taking expectations, $\mathbb{E}[e^{-r(\vartheta - t'_E)} \varphi(\hat{S}_\vartheta)] \geq \mathbb{E}[e^{-r(\vartheta - t'_E)} \varphi(S_\vartheta^{(2)})]$ for every $\vartheta \in \mathcal{T}^{(2)}$. Taking the supremum over $\mathcal{T}^{(2)}$ and then enlarging to the $\mathcal{F}^{(1)}$ -stopping times,

$$\begin{aligned} V^{(2)}(S, t'_E) &\leq \sup_{\vartheta \in \mathcal{T}^{(2)}} \mathbb{E}[e^{-r(\vartheta - t'_E)} \varphi(\hat{S}_\vartheta)] \leq \sup_{\vartheta \in \mathcal{T}^{(1)}} \mathbb{E}[e^{-r(\vartheta - t'_E)} \varphi(\hat{S}_\vartheta)] \\ &= \mathbb{E}[V^X(Se^{Z-\kappa_E}, t'_E)] = V^{(1)}(S, t'_E), \end{aligned}$$

where the last supremum is evaluated by conditioning on Z , which is $\mathcal{F}_{t'_E}^{(1)}$ -measurable: given Z , it is the Kou American problem started at $Se^{Z-\kappa_E}$.

Step 3 (before t'_E). On $[t, t'_E)$ both models follow the same Kou dynamics. By the dynamic programming principle at the deterministic time t'_E (Peskir and Shiryaev, 2006, Chapter I),

$$V(S, t) = \sup_{\vartheta \in \mathcal{T}_{t, t'_E}} \mathbb{E} \left[e^{-r(\vartheta-t)} \varphi(S_{\vartheta}) \mathbf{1}_{\{\vartheta < t'_E\}} + e^{-r(t'_E-t)} V(S_{t'_E}^-, t'_E^-) \mathbf{1}_{\{\vartheta = t'_E\}} \right],$$

for $V = V^{(1)}$ and for $V = V^{(2)}$, with the same S on $[t, t'_E)$ for both models and $V^{(2)}(\cdot, t'_E^-) = V^{(2)}(\cdot, t'_E)$. The right-hand side is monotone in the terminal function, and Step 2 gives $V^{(1)}(\cdot, t'_E^-) \geq V^{(2)}(\cdot, t'_E^-)$; hence $V^{(1)}(S, t) \geq V^{(2)}(S, t)$ for all $t < t'_E$ and in particular at $t = 0$. $\square$

A.8 Proof of Proposition 4.1

From (23) and (24), for $1 \leq i \leq M-1$, $a_{i,i+1} = -\frac{\Delta\tau}{2} \left(\frac{\sigma^2}{2\Delta x^2} \pm \frac{\mu}{2\Delta x} \right)$ and $a_{ii} = 1 + \frac{\Delta\tau}{2} \left(\frac{\sigma^2}{\Delta x^2} + r + \lambda \right) > 0$. Both off-diagonal entries are non-positive iff $\sigma^2/\Delta x^2 \geq |\mu|/\Delta x$, which is (25). In that case $\sum_{j \neq i} |a_{ij}| = \frac{\Delta\tau}{2} \sigma^2/\Delta x^2$, so the row excess is $1 + \frac{\Delta\tau}{2} (r + \lambda)$; the boundary rows have excess 1. A matrix with positive diagonal and non-positive off-diagonal entries that is strictly diagonally dominant is a nonsingular M-matrix, and the inverse of a nonsingular M-matrix is entrywise nonnegative (Berman and Plemmons, 1979, Chapter 6). $\square$

Symmetrization for projected SOR. Let (25) hold strictly, so that $a_{i,i+1} < 0$ and $a_{i+1,i} < 0$, and let A_I be the interior block of A (rows and columns $1, \dots, M-1$; the Dirichlet values enter the right-hand side). With $\Pi = \text{diag}(\pi_i)$, $\pi_1 = 1$ and $\pi_{i+1}/\pi_i = (a_{i+1,i}/a_{i,i+1})^{1/2}$, the matrix $\Pi^{-1}A_I\Pi$ is symmetric, with the diagonal of A and off-diagonal entries $-(a_{i,i+1}a_{i+1,i})^{1/2}$. By the inequality of the arithmetic and geometric means the two off-diagonal moduli of a row sum to at most $|a_{i,i-1}| + |a_{i,i+1}|$, so $\Pi^{-1}A_I\Pi$ is strictly diagonally dominant with positive diagonal, hence positive definite. Under $\mathbf{V} = \Pi\mathbf{y}$ the complementarity problem (37) becomes the one with matrix $\Pi^{-1}A_I\Pi$, right-hand side $\Pi^{-1}\mathbf{b}$ and obstacle $\Pi^{-1}\varphi$, because dividing row i by $\pi_i > 0$ preserves the inequalities and divides the i -th product in the complementarity condition by π_i^2 . The iterates of (40) correspond in the same way: applied to the scaled problem, the update of node i is that of the original problem divided by π_i , and the projection commutes with the positive scaling. The theorem of Cryer (1971) for symmetric positive definite matrices therefore gives convergence of (40) for every $\omega_{\text{SOR}} \in (0, 2)$.

A.9 Hat-function weights: proof of Lemma 4.2

Partial moments. With $\Phi_0(y) = \int_0^y f_Y$ and $\Phi_1(y) = \int_0^y s f_Y(s) ds$ (signed integrals),

$$\begin{aligned} \Phi_0(y) &= p(1 - e^{-\eta_1 y}), & \Phi_1(y) &= p \left(\frac{1}{\eta_1} - \left(y + \frac{1}{\eta_1} \right) e^{-\eta_1 y} \right), & y &\geq 0, \\ \Phi_0(y) &= -(1-p)(1 - e^{\eta_2 y}), & \Phi_1(y) &= (1-p) \left(\frac{1}{\eta_2} + \left(y - \frac{1}{\eta_2} \right) e^{\eta_2 y} \right), & y &< 0, \end{aligned} \tag{53}$$

by direct integration (integrate $ye^{\mp\eta y}$ by parts), and $G_0(a, b) = \Phi_0(b) - \Phi_0(a)$, $G_1(a, b) = \Phi_1(b) - \Phi_1(a)$. For a shifted density $\tilde{f}(y) = f_E(y + \kappa_E)$ the partial moments are $\tilde{G}_0(a, b) = G_0^E(a + \kappa_E, b + \kappa_E)$ and $\tilde{G}_1(a, b) = \int_{a+\kappa_E}^{b+\kappa_E} (s - \kappa_E) f_E(s) ds = G_1^E(a + \kappa_E, b + \kappa_E) - \kappa_E G_0^E(a + \kappa_E, b + \kappa_E)$, with $G_{0,1}^E$ built from (53) with the announcement parameters.

Half weights, (29). On $[y_{k-1}, y_k]$ the hat ψ_k equals $(y - y_{k-1})/\Delta x$, so $w_k^- = (G_1(y_{k-1}, y_k) - y_{k-1}G_0(y_{k-1}, y_k))/\Delta x$; on $[y_k, y_{k+1}]$ it equals $(y_{k+1} - y)/\Delta x$, giving w_k^+ .

(i) *Mass and interpolation.* The hats form a partition of unity, $\sum_k \psi_k \equiv 1$, so $\sum_k w_k = \int f_Y = 1$. The piecewise-linear interpolant of a grid function is $I_h \mathbf{V}(x) = \sum_j V_j \psi_0(x - x_j)$, and $x_i + y - x_j = y - (j - i)\Delta x$, so $I_h \mathbf{V}(x_i + y) = \sum_j V_j \psi_{j-i}(y)$ and $\int I_h \mathbf{V}(x_i + y) f_Y(y) dy = \sum_j V_j w_{j-i}$.

(ii) *Geometric structure.* For $k \geq 1$ the support $[y_{k-1}, y_{k+1}]$ of ψ_k lies in $[0, \infty)$, where $f_Y(y) = p\eta_1 e^{-\eta_1 y}$. Substituting $y = k\Delta x + s$,

$$\begin{aligned} w_k &= p\eta_1 e^{-\eta_1 k\Delta x} \int_{-\Delta x}^{\Delta x} \left(1 - \frac{|s|}{\Delta x}\right) e^{-\eta_1 s} ds = p\eta_1 \rho_1^k \cdot \frac{2}{\Delta x} \int_0^{\Delta x} (\Delta x - s) \cosh(\eta_1 s) ds \\ &= p\eta_1 \rho_1^k \cdot \frac{2(\cosh(\eta_1 \Delta x) - 1)}{\eta_1^2 \Delta x} = \frac{4p \sinh^2(\eta_1 \Delta x/2)}{\eta_1 \Delta x} \rho_1^k, \end{aligned}$$

using $\int_0^h (h-s) \cosh(\eta s) ds = (\cosh(\eta h) - 1)/\eta^2$ and $\cosh z - 1 = 2 \sinh^2(z/2)$. The case $k \leq -1$ is the mirror image with $(1-p, \eta_2)$. The half weights have the same form with $\int_{-\Delta x}^{\Delta x}$ replaced by $\int_{-\Delta x}^0$ or $\int_0^{\Delta x}$, which still factor as a constant times $\rho^{|k|}$, so the ratios $w_k^\pm/w_k$ are constant on each side; this is what allows the recursion (56) to absorb the edge columns by scaling V_0 and V_M . For w_0 the hat straddles the origin: $\int_0^{\Delta x} (1-y/\Delta x) \eta e^{-\eta y} dy = 1 - (1 - e^{-\eta \Delta x})/(\eta \Delta x)$, applied with (p, η_1) on the right and $(1-p, \eta_2)$ on the left, gives (28). As a check, the geometric sum $\sum_{k \geq 1} w_k = A_1 \rho_1 / (1 - \rho_1) = p(1 - e^{-\eta_1 \Delta x})/(\eta_1 \Delta x)$, and adding the right half of w_0 gives exactly p ; the left side gives $1-p$.

(iii) *Trapezoid mass.* With $z = \eta \Delta x$, $\Delta x \eta (\frac{1}{2} + \sum_{k \geq 1} e^{-kz}) = z(\frac{1}{2} + \frac{1}{e^z - 1}) = \frac{z}{2} \coth \frac{z}{2}$, applied to each side with weights p and $1-p$; and $\frac{z}{2} \coth \frac{z}{2} = 1 + \frac{z^2}{12} - \frac{z^4}{720} + \dots$ $\square$

A.10 Tail integrals

For $a \leq 0$, $\int_{-\infty}^a f_Y(y) dy = (1-p)e^{\eta_2 a}$ and $\int_{-\infty}^a e^y f_Y(y) dy = (1-p)\eta_2 \int_{-\infty}^a e^{(\eta_2+1)y} dy = (1-p)\eta_2 e^{(\eta_2+1)a}/(\eta_2+1)$. With $a_i = x_0 - x_i \leq 0$ and the lower far field $A(\tau) - B(\tau)e^x$ of the put,

$$\begin{aligned} \int_{-\infty}^{a_i} (A - B e^{x_i+y}) f_Y(y) dy &= A(1-p)e^{\eta_2 a_i} - B e^{x_i} (1-p) \frac{\eta_2}{\eta_2+1} e^{(\eta_2+1)a_i} \\ &= (1-p)e^{\eta_2 a_i} \left(A - B \frac{\eta_2}{\eta_2+1} e^{x_0} \right), \end{aligned}$$

since $e^{x_i} e^{a_i} = e^{x_0}$; this is (31). The European put has $(A, B) = (K e^{-r\tau}, e^{-q\tau})$. The two branches of the American far field differ by $(K - e^x) - (K e^{-r\tau} - e^{x-q\tau}) = K(1 - e^{-r\tau}) - e^x(1 - e^{-q\tau})$, which is nonincreasing in x and nonnegative exactly for $x \leq x_\tau^*$ of Section 4.2. The solver switches to $(K, 1)$ when $K - e^{x_0}$ exceeds $K e^{-r\tau} - e^{x_0-q\tau}$, that is, when $x_0 < x_\tau^*$, and the payoff branch then dominates on the whole lower tail; this is always the case for $q \leq r$, $r > 0$. Otherwise the tail omits $\int_{-\infty}^{x_\tau^*-x_i} [K(1 - e^{-r\tau}) - e^{x_i+y}(1 - e^{-q\tau})] f_Y(y) dy \leq K(1 - e^{-r\tau})(1-p)e^{\eta_2(x_\tau^*-x_i)}$. For $b \geq 0$, $\int_b^\infty f_Y = p e^{-\eta_1 b}$ and $\int_b^\infty e^y f_Y(y) dy = p\eta_1 e^{-(\eta_1-1)b}/(\eta_1-1)$, which converges because $\eta_1 > 1$; with $b_i = x_M - x_i$ and the upper far field $B e^x - A$ of the call this gives $\text{tail}_i^{\text{call}} = p e^{-\eta_1 b_i} (B \frac{\eta_1}{\eta_1-1} e^{x_M} - A)$, with the same switch to $(K, 1)$ for the American call when $e^{x_M} - K$ exceeds the discounted forward, that is, when $x_M > x_\tau^*$, which is then exact on the whole upper tail; otherwise, possible only for $q < r$, the tail omits the integral of the positive difference $e^{x_i+y}(1 - e^{-q\tau}) - K(1 - e^{-r\tau})$ over $y > x_\tau^* - x_i$, which is at most $(1 - e^{-q\tau}) e^{x_i} p\eta_1 e^{-(\eta_1-1)(x_\tau^*-x_i)}/(\eta_1-1)$. The upper tail of the put and the lower tail of the call vanish because the far-field value is zero there. For the announcement operator the same integrals are taken against the shifted density $f_E(\cdot + \kappa_E)$, with the same switch: with F_E the distribution function of Z and $H_E(y) = \int_{-\infty}^y e^s f_E(s) ds$,

$$\int_{-\infty}^a f_E(y + \kappa_E) dy = F_E(a + \kappa_E), \quad \int_{-\infty}^a e^y f_E(y + \kappa_E) dy = e^{-\kappa_E} H_E(a + \kappa_E),$$

where $H_E(y) = (1-p_E)\eta_{2,E} e^{(\eta_{2,E}+1)y}/(\eta_{2,E}+1)$ for $y \leq 0$ and $H_E(y) = (1-p_E)\eta_{2,E}/(\eta_{2,E}+1) + p_E \eta_{1,E} (1 - e^{-(\eta_{1,E}-1)y})/(\eta_{1,E}-1)$ for $y > 0$, and $H_E(\infty) = \mathbb{E}[e^Z] = e^{\kappa_E}$; the upper tails use $1 - F_E$ and $e^{-\kappa_E}(H_E(\infty) - H_E)$. The identity $\mathcal{P}_E^h e^x = e^x$ of Section 4.4 is the statement that the on-grid weights applied to e^{x_j} plus these two tails reproduce $e^{-\kappa_E} H_E(\infty) e^{x_i} = e^{x_i}$, up to the interpolation error of e^x .

A.11 Proof of Lemma 4.3

Applying (33) to the scalar model gives $(1 + \frac{\alpha}{2})v^{n+1} = (1 - \frac{\alpha}{2})v^n + \beta(\frac{3}{2}v^n - \frac{1}{2}v^{n-1})$, whose characteristic polynomial is (35). Write $\nu = \lambda\Delta\tau$, so that $\alpha \geq \nu$ and $|\beta| \leq \nu$.

Necessity. Take $a = 0$ and $\theta = -1$: $\alpha = \nu$, $\beta = -\nu$, and (35) becomes $(1 + \frac{\nu}{2})z^2 - (1 - 2\nu)z - \frac{\nu}{2} = 0$, with the real roots $z_{\pm} = ((1 - 2\nu) \pm \sqrt{5\nu^2 - 2\nu + 1})/(2 + \nu)$. Because $5\nu^2 - 2\nu + 1 - (1 - 2\nu)^2 = \nu^2 + 2\nu > 0$, $z_- < 0 < z_+$, so $|z_-| = (2\nu - 1 + \sqrt{5\nu^2 - 2\nu + 1})/(2 + \nu)$, and $|z_-| \geq z_+$ if and only if $2\nu - 1 \geq 1 - 2\nu$, i.e. $\nu \geq \frac{1}{2}$. The root z_- satisfies $z_- < -1$ iff $2\nu - 1 + \sqrt{5\nu^2 - 2\nu + 1} > 2 + \nu$, i.e. iff $\sqrt{5\nu^2 - 2\nu + 1} > 3 - \nu$; for $\nu \geq 3$ this holds, and for $\nu < 3$ squaring gives $4\nu^2 + 4\nu - 8 > 0$, i.e. $(\nu + 2)(\nu - 1) > 0$, i.e. $\nu > 1$. At $\nu = 1$, $z_- = -1$. Hence for $\nu > 1$ a root lies outside the closed disc, and the stated formula for $|z_-|$, and the range $\nu \geq \frac{1}{2}$ on which it is the larger modulus, follow.

Sufficiency. Let $\nu < 1$ and suppose $z = e^{i\omega}$ is a root of (35) for some $a \geq 0$ and $|\theta| \leq 1$. Then $(1 + \frac{\alpha}{2})e^{2i\omega} - (1 - \frac{\alpha}{2})e^{i\omega} = \frac{\beta}{2}(3e^{i\omega} - 1)$, and taking moduli, with $\gamma = \cos \omega$,

$$\left| (1 + \frac{\alpha}{2})e^{i\omega} - (1 - \frac{\alpha}{2}) \right|^2 = 2(1 - \gamma) + \frac{\alpha^2}{2}(1 + \gamma) \leq \frac{|\beta|^2}{4} |3e^{i\omega} - 1|^2 = \frac{|\beta|^2}{4} (10 - 6\gamma) \leq \frac{\nu^2}{4} (10 - 6\gamma).$$

Using $\alpha \geq \nu$ on the left, $2(1 - \gamma) + \frac{\nu^2}{2}(1 + \gamma) - \frac{\nu^2}{4}(10 - 6\gamma) = 2(1 - \gamma)(1 - \nu^2) \leq 0$, which for $\nu < 1$ forces $\gamma = 1$, i.e. $z = 1$. Substituting $z = 1$ into (35) gives $\alpha = \beta$, i.e. $(a + \lambda) = \lambda\theta$, which with $a \geq 0$ and $|\theta| \leq 1$ forces $a = 0$ and $\theta = 1$: the neutral mode, whose other root is $\nu/(2 + \nu) < 1$. So for $\nu < 1$ a root can lie on the unit circle only at the single parameter point $(a, \theta) = (0, 1)$. Now let $R(a, \theta)$ be the larger root modulus, a continuous function on the connected set $\mathcal{U} = \{a \geq 0, |\theta| \leq 1\}$, and note that $R < 1$ at $(a, \theta) = (1, 0)$, where the roots are $(1 - \frac{\alpha}{2})/(1 + \frac{\alpha}{2})$ and 0. If $R > 1$ somewhere on $\mathcal{U}$, a path in $\mathcal{U} \setminus \{(0, 1)\}$ (which is connected, since $\mathcal{U}$ has nonempty interior in $\mathbb{R} \times \mathbb{C}$) from $(1, 0)$ to that point would contain a point with $R = 1$, i.e. a root on the unit circle away from $(0, 1)$, contradicting the above. Hence $R \leq 1$ on $\mathcal{U}$, with $R < 1$ off the neutral mode. For $\nu = 1$ the roots are limits of roots with $\nu < 1$, so $R \leq 1$ there as well. This is the root-locus form of the Schur–Cohn criterion (Hundsdoerfer and Verwer, 2003): the reduced polynomial of (35) is of degree one, and the argument above locates the only parameter at which its root reaches the unit circle. $\square$

B Implementation details and further verification

This appendix records the solver as it is implemented in the package `kousej` (`code/kousej/fd.py`, `cos.py`, `bench.py`, `model.py`), so that the formulas of Section 4 can be matched line by line with the code, and it states the defaults that every run of the paper uses. Where the code departs from the wording of Section 4 I say so. Appendices B.9 to B.20 then report the Black–Scholes limit and the verification checks that support Section 5 without being benchmarks of their own, with their tables and figures.

B.1 Grid

The grid is $x_i = x_c + (i - M/2)\Delta x$, $i = 0, \dots, M$, with M even and at least four. The center x_c defaults to $\ln K$; the verification ladders at $S_0 \neq K$ pass $x_c = \ln S_0$ instead (Tables 1 and 2, Panel B), and in every other verification run $S_0 = K$, so that both are the center node. The half-width is

$$L = 6\sigma\sqrt{T} + \frac{4}{\min(\eta_1, \eta_2)} \mathbf{1}_{\{\lambda > 0\}} + \frac{c_E}{\min(\eta_{1,E}, \eta_{2,E})} \mathbf{1}_{\{0 < t_E < T\}}, \quad c_E = 6 \text{ by default}, \quad (54)$$

unless a value is passed explicitly, and the step starts as $\Delta x = 2L/M$. If $x_c \neq \ln K$, the step is then snapped so that $\ln K$ is a node: with $d = |\ln K - x_c|$ and $m = \max\{1, \text{round}(d/\Delta x)\}$,

$\Delta x \leftarrow d/m$ and $L \leftarrow M\Delta x/2$, after which the node nearest to $\ln K$ is set to $\ln K$ exactly. This is why the half-width varies along the $S_0 = 110$ ladder of Table 1 and why the coarse grids there are not nested. The announcement reach c_E is the parameter varied in the `jump_reach` check of Table 20, and the domain checks there and in Table 1 double L at fixed Δx . Prices at a spot that is a node are read off the node; at any other spot they are a cubic spline in x through the grid values (`FDResult.price`), which is how the $S = 85, \dots, 105$ rows of Table 3 are obtained on a grid centered at $\ln K$. Delta and gamma are central differences in x at the interior nodes, $\Delta = u_x/S$ and $\Gamma = (u_{xx} - u_x)/S^2$, and the COS reference of Figure 5 is differentiated with the same stencil on the same nodes.

B.2 Weights as implemented

The weights of Lemma 4.2 are computed from the signed partial moments of the kernel, $\Phi_0(y) = \int_0^y f(s) ds$ and $\Phi_1(y) = \int_0^y s f(s) ds$, which are elementary for the double-exponential density and are coded once for f_Y and once for f_E . For a kernel shifted by s (the announcement uses $f_E(\cdot + \kappa_E)$, so $s = \kappa_E$; the Kou kernel has $s = 0$) and a lag $y = k\Delta x$, the code forms

$$\begin{aligned} G_0(a, b) &= \Phi_0(b + s) - \Phi_0(a + s), & G_1(a, b) &= \Phi_1(b + s) - \Phi_1(a + s) - s G_0(a, b), \\ w_k^- &= \frac{G_1(y - \Delta x, y) - (y - \Delta x) G_0(y - \Delta x, y)}{\Delta x}, \\ w_k^+ &= \frac{(y + \Delta x) G_0(y, y + \Delta x) - G_1(y, y + \Delta x)}{\Delta x}, & w_k &= w_k^- + w_k^+, \end{aligned} \quad (55)$$

which is (29) with the shift carried inside the primitives, so that the shift is exact and not snapped to the grid. The full weights w_k for $k = -M, \dots, M$ fill the $(M+1)^2$ matrix $\tilde{w}_{ij} = w_{j-i}$ of (32); the column $j = 0$ is replaced by the right half weights w_{-i}^+ and the column $j = M$ by the left half weights w_{M-i}^- . The trapezoid alternative used as a foil in Section 5 takes $w_k = \Delta x f(k\Delta x)$, $w_0 = \frac{1}{2}\Delta x (f(0^+) + f(0^-))$ and half of each weight at the two edge columns; its mass error is recorded as the maximum over all rows of the operator, including the two Dirichlet rows, which is what Appendix B.10 explains.

B.3 Tails

Outside the grid the solution is replaced by $A(\tau) - B(\tau)e^x$ for the put and by $B(\tau)e^x - A(\tau)$ for the call, and the integral of the kernel against this form over the missing half-line is evaluated in closed form at every time level, exactly as in (31). With $a_i = x_0 - x_i \leq 0$ and $b_i = x_M - x_i \geq 0$ the Kou tails are $\lambda(1-p)e^{\eta_2 a_i} (A - B \frac{\eta_2}{\eta_2+1} e^{x_0})$ for the put and $\lambda p e^{-\eta_1 b_i} (B \frac{\eta_1}{\eta_1-1} e^{x_M} - A)$ for the call, with $(A, B) = (K e^{-r\tau}, e^{-q\tau})$ for a European option and, for an American option, $(K, 1)$ whenever the payoff exceeds the European far field at the edge node. The announcement tails use the distribution function F_E and the partial exponential moment $\int_{-\infty}^c e^s f_E(s) ds$ of the law of Z , both closed-form, evaluated at $a_i + \kappa_E$ and $b_i + \kappa_E$ and multiplied by $e^{-\kappa_E}$ where e^x appears. Three identities check the tails together with the weights and are reported in Section 5: $J\mathbf{1} + \text{tails} = \lambda$, $J e^x + \text{tails} = \lambda(1 + \kappa) e^x$ up to interpolation error, and $\mathcal{P}_E^h e^x + \text{tails} = e^x$ up to interpolation error.

B.4 Three evaluations of the jump operator

The recursion and the FFT form. By (27), with $\hat{\mathbf{V}}$ equal to $\mathbf{V}$ except that $\hat{V}_0$ and $\hat{V}_M$ are scaled by the constant ratios w_{-1}^+/w_{-1} and w_1^-/w_1 that convert full hats into half hats, the discrete jump operator (32) is

$$\begin{aligned} U_i &= \rho_1(\hat{V}_{i+1} + U_{i+1}), \quad U_M = 0; & L_i &= \rho_2(\hat{V}_{i-1} + L_{i-1}), \quad L_0 = 0; \\ (J\mathbf{V})_i &= \lambda(w_0 V_i + A_1 U_i + A_2 L_i) + \lambda \text{tail}_i(\tau) \end{aligned} \quad (56)$$

for $0 < i < M$, with the diagonal weight at the two edge rows replaced by its inside half. Each recursion is a first-order linear filter and costs $O(M)$ operations. Alternatively, let $w_m^r = w_{M-m}$, $m = 0, \dots, 2M$, be the reversed kernel. Then the sum $\sum_j w_{j-i} V_j$ equals $(\mathbf{V} * \mathbf{w}^r)_{i+M}$, a linear convolution, and

$$\begin{aligned} (J\mathbf{V})_i &= \lambda \left[\text{IFFT} \left(\text{FFT}(\mathbf{V}_{\text{pad}}) \odot \text{FFT}(\mathbf{w}_{\text{pad}}^r) \right) \right]_{i+M} \\ &\quad + \lambda(w_{-i}^+ - w_{-i})V_0 + \lambda(w_{M-i}^- - w_{M-i})V_M + \lambda \text{tail}_i(\tau), \end{aligned} \quad (57)$$

where both vectors are zero-padded to a common length of at least $3M + 1$, so that the circular convolution coincides with the linear one on the $M + 1$ entries that are read out. The reversal of the kernel (equivalently, conjugating its transform) implements the correlation.

As implemented. The on-grid part $\sum_j \tilde{w}_{ij} V_j$ is evaluated in one of three ways, selected by `jump_method`. *Dense*: the matrix product, $O(M^2)$ per step and $(M + 1)^2$ doubles of memory (313 megabytes at $M = 6400$). *FFT*: the kernel $w_{-M}, \dots, w_M$ is reversed and both vectors are zero-padded to $n = 2^{\lceil \log_2(3M+1) \rceil}$ points; the real FFT of the padded reversed kernel is stored at set-up, each step costs one forward and one inverse real FFT of length n , entries $M, \dots, 2M$ of the circular convolution are read out, and the two edge columns are corrected by adding $(w_{-i}^+ - w_{-i})V_0 + (w_{M-i}^- - w_{M-i})V_M$. *Recursion*: with $\rho_1 = e^{-\eta_1 \Delta x}$, $\rho_2 = e^{-\eta_2 \Delta x}$ and the amplitudes $A_1 = w_1/\rho_1$ and $A_2 = w_{-1}/\rho_2$ (those of (27) for the hat weights), the sums over $j > i$ and $j < i$ are two first-order linear filters, $U_i = \rho_1(\hat{V}_{i+1} + U_{i+1})$ and $L_i = \rho_2(\hat{V}_{i-1} + L_{i-1})$ (`scipy.signal.lfilter`), where $\hat{V}$ equals V except that the edge entries are scaled by w_{-1}^+/w_{-1} and w_1^-/w_1 , and the result is $\lambda(A_1 U_i + A_2 L_i + w_0 V_i)$ with the diagonal weight replaced by its inside half in the two edge rows. This is (56), and it is exact for both quadratures because the weight at lag k is a pure geometric sequence on each side of zero. The three forms agree to 1.2×10^{-17} (dense against FFT) and 8.7×10^{-14} (dense against recursion) on a random vector (Table 1); the verification ladders use the dense form below $M = 3200$ and the recursion from there on, and the timings of Table 22 show where each is cheapest. The announcement operator is implemented in the dense and FFT forms only; the recursion of Section 4.4 is not coded for it, since it is applied once per solve. Its grid step is taken as $x_1 - x_0$ rather than the stored Δx , a rounding-level difference.

B.5 Time stepping and the two legs

With an announcement in the life of the option the march is split into the legs $[0, \tau_E]$ and $[\tau_E, T]$ with $N_A = \max\{n_h/2, \text{round}(N\tau_E/T)\}$ and $N_B = \max\{n_h/2, \text{round}(N(T - \tau_E)/T)\}$ steps, $n_h = 4$ being the number of Rannacher half steps, so that each leg has at least two full steps and $N_A + N_B$ equals N up to the rounding ($800 = 400 + 400$ at $t_E = T/2$; $756 + 44$ at $t_E = 0.5$ days in Panel C of Table 3). Each leg has its own $\Delta\tau$ and its own tridiagonal matrix $A = I - \frac{\Delta\tau}{2}D$, stored in banded form and solved by `scipy.linalg.solve_banded`. The Rannacher start of each leg consists of $n_h = 4$ IMEX-Euler half steps, $A\mathbf{V}^{k+1/2} = \mathbf{V}^k + \frac{\Delta\tau}{2}J\mathbf{V}^k$ with the same matrix A as the Crank–Nicolson step, which is the point of taking half steps; they produce $\mathbf{V}^1, \mathbf{V}^2$ and the history $J\mathbf{V}^1$, and the CNAB recursion (33) runs from there with J evaluated once per step and reused. The scalar stability scan of Lemma 4.3 (`cnab_stability_scan`) takes a on $\{0\} \cup [10^{-3}, 10^4]$ and θ on 719 points of the unit circle; the kernel-specific scan (`cnab_stability_scan_kernel`) uses the discrete symbols $a(\xi) = \sigma^2(1 - \cos \xi \Delta x)/\Delta x^2 + (r + \lambda) - i\mu \sin(\xi \Delta x)/\Delta x$ and $\theta(\xi) = \sum_k w_k e^{ik\xi \Delta x}$ on 399 frequencies in $(0, \pi/\Delta x]$.

B.6 The three LCP solvers

Operator splitting (`'it'`, the default): the right-hand side receives $\Delta\tau\boldsymbol{\mu}^n$ in the interior rows, the tridiagonal system is solved, $\mathbf{V}^{n+1} = \max(\tilde{\mathbf{V}} - \Delta\tau\boldsymbol{\mu}^n, \boldsymbol{\varphi})$, the Dirichlet rows are reset, and

$\mu^{n+1} = \max(\mu^n + (\mathbf{V}^{n+1} - \tilde{\mathbf{V}})/\Delta\tau, 0)$ with zero in the Dirichlet rows; the maximum with zero is a guard, since the update is nonnegative by construction. In the Rannacher half steps $\Delta\tau/2$ takes the place of $\Delta\tau$. *Brennan–Schwartz* ('bs'): the eliminated diagonal b'_i depends only on the three constant bands of A and is computed once per leg; for a put the super-diagonal is eliminated downward from $i = M - 1$ and the back-substitution runs upward from $i = 1$ with $V_i = \max\{(r'_i - aV_{i-1})/b'_i, \varphi_i\}$; for a call both sweeps are mirrored. *Projected SOR* ('psor'): relaxation $\omega_{\text{SOR}} = 1.3$, warm start from $\max(\mathbf{V}^n, \varphi)$ with the Dirichlet rows set, sweeps in increasing i , and the stopping test $\max_i |V_i^{(k+1)} - V_i^{(k)}| < 10^{-10}$ in the absolute max norm, with at most 20 000 sweeps. Section 4.5 writes the test in the relative form $\varepsilon \max\{1, \|\mathbf{V}\|_\infty\}$; the code uses the absolute form, which for prices below 100 is the stricter of the two, and the sweep counts of Tables 2 and 18 refer to it. The PSOR and Brennan–Schwartz sweeps are pure-Python loops, which is why PSOR is slow in Table 22 and why Brennan–Schwartz, with one sweep, is not. Projected SOR is the cross-check and is never used in production.

B.7 The announcement step and the restart

At the end of the first leg, $\tau = \tau_E$, the solver stores $\mathbf{V}_{\text{pre}} = \mathbf{V}$ (the value immediately after the announcement in calendar time), computes $\mathbf{V}_{\text{cont}} = \mathcal{P}_E^h \mathbf{V} + \text{tails}^E(\tau_E)$ with the dense matrix of (36), sets $\mathbf{V} \leftarrow \max(\mathbf{V}_{\text{cont}}, \varphi)$ for an American option and $\mathbf{V} \leftarrow \mathbf{V}_{\text{cont}}$ otherwise, resets the multiplier μ to zero, records the exercise boundary once more at τ_E (the record from the end of leg A is the boundary at t_E^+ , the record after the projection is the boundary at t_E^-), and starts leg B with its own Rannacher half steps and an empty Adams–Bashforth history. The three vectors $\mathbf{V}_{\text{pre}}$, $\mathbf{V}_{\text{cont}}$ and $\mathbf{V}_{\text{post}}$ are returned in `FDRresult.info`, and $\min_i ((\mathbf{V}_{\text{cont}})_i - \varphi_i)$ and $(\mathbf{V}_{\text{cont}} - \mathbf{V}_{\text{pre}})_{M/2}$ are the two numbers of Panel B in Table 3. The exercise boundary after every full step is the largest node with $V_i \leq \varphi_i + 10^{-10}$ and $S_i < K$ for a put (the smallest node with $S_i > K$ for a call), and is recorded as not-a-number when no such node exists, which is how the empty region just before the announcement appears in Figure 2(b). The announcement is applied only if $0 < t_E < T$; otherwise the solver runs a single leg and the model is Kou's. Passing `compensated=False` drops κ_E from the shifted density and from the tails and is used only for the parity-violation diagnostic of Table 20.

B.8 Orders, Richardson extrapolation and the benchmarks

Observed orders along a ladder are $\log(e_k/e_{k+1})/\log(\Delta x_k/\Delta x_{k+1})$; because $N = M/2$ on every ladder, this is a joint order in Δx and $\Delta\tau$. Richardson extrapolation is $(4v_{\text{fine}} - v_{\text{coarse}})/3$ from the two finest grids, and the “error estimate” attached to it in Table 2 is $|v_{\text{fine}} - v_{\text{Richardson}}|$. The cosine benchmark (Fang and Oosterlee, 2008) truncates the log-return density to $[c_1 - L_c\sqrt{c_2 + \sqrt{c_4}}, c_1 + L_c\sqrt{c_2 + \sqrt{c_4}}]$ from the cumulants (10), uses $N = 2^{12}$ terms and $L_c = 10$ by default, expands the put payoff and prices calls by parity. The reference values use $N = 2^{14}$ with $L_c = 20$ for the Kou test case of Table 1, whose rare large jumps a narrower interval cuts off (at $L_c = 12$ the value is off by up to 5.9×10^{-7}), and $L_c = 12$ for the scheduled-jump references, which change by less than 10^{-12} when N is raised to 2^{16} and L_c to 30. The mixture benchmark of Proposition 2.3(iii) integrates $V^{\text{BS}}(S_0 e^{z - \kappa_E}) f_E(z)$ by adaptive quadrature (`scipy.integrate.quad`) separately on $[-z_{\text{max}}, 0]$ and $[0, z_{\text{max}}]$ with $z_{\text{max}} = 50/\min(\eta_{1,E}, \eta_{2,E})$, absolute tolerance 10^{-13} and relative tolerance 10^{-12} ; the error estimate it returns is the one quoted in Table 20. The binomial benchmark of Tables 2 and 17 is the binomial Black–Scholes method with Richardson extrapolation (BBSR) of Broadie and Detemple (1996), coded in the verification script: the Cox–Ross–Rubinstein tree of n steps whose continuation value at the last step before expiry is the Black–Scholes European value over one step, evaluated at $n = 10\,000$, 20 000 and 40 000 with backward induction vectorized over the nodes, and extrapolated as $2V_n - V_{n/2}$; the value at 40 000 steps is the reference. For the puts at $S_0 = K$, where the spot is

Table 15: IMEX-CNAB with four Rannacher half steps at $\lambda = 0$ against the Black–Scholes closed form: $S_0 = K = 100$, $r = 0.05$, $q = 0$, $\sigma = 0.3$, $T = 0.5$. The error is $|V_{\text{FD}}(S_0) - V_{\text{BS}}(S_0)|$; the observed order is $\log_2$ of the ratio of successive errors.

M	N	Δx	$\Delta \tau$	V_{FD}	error	order
<i>European put, Black–Scholes value 7.16586783</i>						
100	50	0.02546	0.01000	7.15084739	1.50×10^{-2}	–
200	100	0.01273	0.00500	7.16212055	3.75×10^{-3}	2.00
400	200	0.00636	0.00250	7.16493150	9.36×10^{-4}	2.00
800	400	0.00318	0.00125	7.16563378	2.34×10^{-4}	2.00
Richardson from the two finest grids				7.16586787	4.04×10^{-8}	–
<i>European call, Black–Scholes value 9.63487663</i>						
100	50	0.02546	0.01000	9.61999255	1.49×10^{-2}	–
200	100	0.01273	0.00500	9.63116344	3.71×10^{-3}	2.00
400	200	0.00636	0.00250	9.63394882	9.28×10^{-4}	2.00
800	400	0.00318	0.00125	9.63464471	2.32×10^{-4}	2.00
Richardson from the two finest grids				9.63487667	4.04×10^{-8}	–

Grid centered at $\ln K$ with half-width $6\sigma\sqrt{T}$; time steps $N = M/2$ so that $\Delta\tau/\Delta x^2$ is not fixed. Richardson extrapolation assumes order two in Δx and $\Delta\tau$ jointly.

a node of every even-step tree, the Richardson value $2 \text{CRR}_{40000} - \text{CRR}_{20000}$ of the plain tree is a second, independent estimate. Terminal densities in Figure 7 are obtained by numerical inversion of the characteristic function on $[-1, 1]$ at 2001 points.

B.9 The differential operator: the Black–Scholes limit

Table 15 prices the at-the-money European put and call with $\sigma = 0.3$, $r = 0.05$ and $T = 0.5$ on four grids from $(M, N) = (100, 50)$ to $(800, 400)$. The put error falls from 1.50×10^{-2} to 2.34×10^{-4} and the call error from 1.49×10^{-2} to 2.32×10^{-4} ; the observed order is 2.00 at every level for both (the unrounded put values are 2.003, 2.001 and 2.000). Richardson extrapolation of the two finest grids gives 7.16586787 against the closed form 7.16586783 and 9.63487667 against 9.63487663, a residual of 4.0×10^{-8} in both cases, which is what a pure second-order error leaves behind. This is the operator D , the matrix A of (24) and the far-field data working as intended; nothing here involves jumps.

B.10 Operator identities and the ends of the Kou ladders

The cosine benchmark. For the Kou test case of Table 1 the characteristic function satisfies the martingale identity $\phi_{\varrho_T}(-i) = e^{rT}$ to rounding, and with λ set to zero the cosine expansion returns the Black–Scholes put to 3×10^{-15} at the three spots. The direct cosine expansion of the call payoff at the default setting differs from the parity value by -9.4×10^{-4} to -1.7×10^{-3} , a known weakness of the expansion for the unbounded call payoff, and is not used anywhere.

The ends of the ladders. At $S_0 = 110$ the orders of Table 1 are 1.88, 1.97, 1.99, 2.00, 2.00 and 1.98, and at $S_0 = 90$ they are 1.62, 1.92, 1.98, 1.99, 1.99 and 1.96. The start of these ladders is distorted: the grid is centered at $\ln S_0$ and the step is snapped so that $\ln K$ is a node, which changes the half-width from level to level (L runs from 1.59 to 1.91 along the $S_0 = 110$ ladder; Appendix B.1), so the coarse grids are not nested and the first order mixes the discretization error with the change of domain. From the second refinement on the orders are close to two at both spots, down to the finest grids, where the errors are 1.77×10^{-6} at $S_0 = 90$ and 4.55×10^{-6} at $S_0 = 110$; the reference is settled far below these errors (Section 5.2).

Operator self-checks. Lemma 4.2(i) says that the hat weights conserve mass exactly. On the (400, 200) grid the maximum over rows of $|(J\mathbf{1})_i + \text{tail}_i - \lambda|$ is 3.4×10^{-15} with hat weights and 2.2×10^{-4} with trapezoid weights. The trapezoid figure needs a comment. The maximum is attained in the two boundary rows, where the trapezoid rule puts the half weight $\frac{1}{4}\Delta x (f_Y(0^+) + f_Y(0^-))$ on the node at lag zero although only one side of the kink of f_Y lies inside the grid, an $O(\Delta x)$ error; the interior rows obey the $(\eta\Delta x)^2/12$ law of Lemma 4.2(iii), which is smaller here. The boundary rows are overwritten by the Dirichlet data, so this does not reach the prices: on the first four grids at $S_0 = 100$ the trapezoid errors are 8.39×10^{-2} , 2.02×10^{-2} , 5.01×10^{-3} and 1.25×10^{-3} against 8.37×10^{-2} , 2.01×10^{-2} , 4.99×10^{-3} and 1.25×10^{-3} with hat weights. On these grids $\eta_2\Delta x$ falls from 0.11 to 0.014 and $\lambda = 0.1$; the trapezoid error exceeds the hat-weight error by a small and almost constant fraction at every level, and both fall at order two, so at this intensity and resolution the choice of quadrature does not show in the price. The regime in which it does, a large intensity on a coarsely resolved kernel, is the subject of Appendix B.14. The dense matrix, the FFT (57) and the recursion (56) applied to the same random vector agree to 1.2×10^{-17} and 8.7×10^{-14} respectively. The orientation check of Section 4.2 gives $(Je^x)_{M/2}/(\lambda e^{x_{M/2}}) = 1.007582$ against $1 + \kappa = 1.007576$, the gap being the interpolation error of e^x on the grid; a reflected kernel would give 1.230388, so an orientation error could not be missed.

Stability. The amplification polynomial (35), scanned over $a \geq 0$ and over the unit circle $|\theta| = 1$ with the point $\theta = 1$ left out (at $a = 0$ it is the neutral mode of Lemma 4.3, whose root is exactly one), has largest root modulus 0.999986, 0.999993 and 0.999999 at $\lambda\Delta\tau = 0.5, 0.9$ and 0.99 , and 1.0100, 1.0975 and 1.4404 at $\lambda\Delta\tau = 1.01, 1.1$ and 1.5 , which is Lemma 4.3 in numbers on the circle; the interior of the disc is covered by the proof of the lemma, not by the scan. With the discrete symbols of D and J actually used on the (400, 200) grid, including the imaginary part that the drift adds, the largest modulus is 0.9986, 0.9993, 0.9997, 0.9998 and 0.9999 at $\lambda\Delta\tau = 0.5, 1, 2, 4$ and 8 . For these parameters and this grid the explicit treatment is therefore stable far beyond the generic bound, a property of the pair $(a(\xi), \theta(\xi))$, namely $|\theta(\xi)| < 1$ off $\xi = 0$ together with the diffusion damping, and not of the sign of $\text{Re } \theta$; the scan says nothing about other parameters or grids.

B.11 The American ladders: the binomial reference and the $S_0 = 90$ ladder

The binomial reference of Panel A of Table 2 (Appendix B.8) is converged well below the finite-difference errors it measures: the BBSR values at 20 000 and 40 000 steps differ by at most 1×10^{-6} , and the Richardson value of the even-step tree agrees with the reference to 5×10^{-7} , while the finest finite-difference errors are -7.0×10^{-5} (case 1) and -2.8×10^{-5} (case 2).

The spot $S_0 = 90$ of Table 2, Panel B, lies 0.48, or 0.54 percent, above the boundary 89.517, and the ladder behaves accordingly. On the two coarsest grids ($\Delta x = 0.035$ and 0.018 , that is, 3.2 and 1.6 currency units at $S = 90$) the node $S_0 = 90$ sits inside the discrete exercise region and the value is the payoff 10 exactly; from (400, 200) on the value rises through 10.00273, 10.00489, 10.00497, 10.00506, 10.00507 to 10.00507 on the (12800, 6400) grid with ratios 1.27, 27.08, 0.86, 9.22 and 3.31. The ratios are not those of a smooth error expansion because the distance between S_0 and the nearest node of the discrete boundary changes from grid to grid, and Richardson extrapolation, which assumes such an expansion, is not justified here; I report it anyway, 10.005072 against the published 10.005071 (difference $+9.2 \times 10^{-7}$), with the finest-grid estimate 1.0×10^{-6} attached, so that the reader can see that the agreement at this spot is of the same size as elsewhere but rests on the finest grid rather than on the extrapolation. Near the boundary the convergence is governed by the kink of the solution's second derivative, and any one-point measure of it is irregular.

Table 16: The time step: fixed- Δx temporal ladders, joint ladders of the early-exercise premium, and the temporal part of the error.

Panel A: Kou put (Toivanen parameters), $M = 3200$ fixed, grid centered at $\ln S_0$; ratios of successive differences in N

N	Eu. (100)	IT (90)	IT (100)	IT (110)	BS (90)	BS (100)	BS (110)
100	4.00	-2.61	3.15	2.18	1.58	2.54	2.00
200	4.00	3.59	3.26	2.27	3.09	2.48	2.08
400	4.00	4.41	3.46	2.35	5.33	2.46	2.14
800	4.00	4.11	4.35	2.60	2.54	2.49	2.29
1600	4.00	2.22	4.95	2.45	2.43	2.36	2.24
3200	4.00	3.44	4.11	3.02	2.79	2.61	2.48

Deeper BS time ladders at $S_0 = 100$: $M = 3200$, N to 25600: ratios 2.57, 2.64, 2.81; $M = 6400$, N to 25600: ratios 2.58, 2.59

Panel B: scheduled jump (base case, $S_0 = K = 100$), $M = 1600$ fixed; ratios of successive differences in N

N	Eu.	IT	BS	EEP (IT)	EEP (BS)
100	4.46	3.20	2.90	2.86	2.56
200	4.00	2.94	2.62	2.75	2.43
400	4.00	2.98	2.53	2.85	2.41
800	4.00	3.19	2.60	3.11	2.53
1600	4.00	3.37	2.55	3.32	2.51
3200	4.00	2.67	2.69	2.60	2.66

Panel C: scheduled jump, $N = M/2$; ratios of successive differences and observed orders against the references

M	Eu.	IT	BS	EEP (IT)	EEP (BS)	ord. EEP (IT)	ord. EEP (BS)	ord. BS
400	—	—	—	—	—	1.72	1.67	1.88
800	4.01	3.72	3.69	3.25	3.19	1.78	1.68	1.86
1600	4.00	3.74	3.66	3.37	3.22	1.87	1.66	1.84
3200	4.00	3.81	3.59	3.58	3.17	1.94	1.66	1.82

Panel D: temporal part of the error

Kou American put, $S_0 = 100$, IT: error at (800, 400) -1.462×10^{-3} ; time refined (800, 3200) -1.446×10^{-3} ; space refined (6400, 400) -3.226×10^{-5}

Kou American put, $S_0 = 100$, BS: error at (800, 400) -1.492×10^{-3} ; time refined (800, 3200) -1.447×10^{-3} ; space refined (6400, 400) -8.594×10^{-5}

Kou American put on (3200, 1600): error (temporal part) IT -9.26×10^{-5} (-1.04×10^{-6}), BS -1.01×10^{-4} (-9.45×10^{-6})

Scheduled jump on (1600, 800), error (temporal part): Eu. -9.42×10^{-6} (-1.13×10^{-7}); IT -1.71×10^{-5} (-1.79×10^{-6}); BS -1.96×10^{-5} (-4.28×10^{-6}); EEP (IT) -7.68×10^{-6} (-1.68×10^{-6}); EEP (BS) -1.02×10^{-5} (-4.17×10^{-6})

IT is the operator splitting, BS the Brennan–Schwartz solution of the discrete LCP of each step, EEP the finite-difference American minus the finite-difference European value on the same grid. A ratio of four is second order, two is first order. The references are time-converged exact-LCP values: Brennan–Schwartz with $N = 25600$ on $M = 3200$ and 6400, extrapolated in space, 2.8078787 for the Kou put at $S_0 = 100$; with $N = 12800$ on $M = 1600$ and 3200, 4.6342069 for the scheduled-jump put, whose European reference is the COS value 4.6112540. The temporal part is the difference between the value on the grid and the value at the same M with the time step taken to zero (the deepest BS time ladder, extended by its last ratio).

B.12 The time step

Table 16 holds the measurements behind Section 5.5. Panels A and B refine time alone at fixed Δx , so the ratio of successive differences in N isolates the temporal order: four for the European solver at every level, 2.36 to 2.61 for the exact discrete LCP at $S_0 = 100$ and 2.00 to 2.48 at $S_0 = 110$, and irregular values for the splitting, whose first-order splitting error and the time error of the projected step have opposite signs and change their balance from level to level. At

Table 17: American calls with a dividend yield: against a binomial reference, and the put–call symmetry with and without an announcement.

<i>Panel A: $\lambda = 0$, no announcement; American call against the binomial reference (BBSR)</i>								
case	M	N	C_{am}	error	order	EEP error	$C - P^*$	
1	400	200	7.650213	-1.20×10^{-3}	–	-2.03×10^{-4}	-6.12×10^{-5}	
	800	400	7.651105	-3.04×10^{-4}	1.98	-5.58×10^{-5}	-1.53×10^{-5}	
	1600	800	7.651333	-7.59×10^{-5}	2.00	-1.39×10^{-5}	-3.82×10^{-6}	
2	400	200	7.415802	-5.84×10^{-4}	–	-1.02×10^{-4}	-3.90×10^{-6}	
	800	400	7.416243	-1.43×10^{-4}	1.93	-3.08×10^{-5}	-9.08×10^{-7}	
	1600	800	7.416349	-3.74×10^{-5}	1.94	-9.28×10^{-6}	-2.27×10^{-7}	
3	400	200	2.420945	-6.10×10^{-4}	–	-1.34×10^{-4}	-1.19×10^{-5}	
	800	400	2.421389	-1.65×10^{-4}	1.89	-4.58×10^{-5}	-2.97×10^{-6}	
	1600	800	2.421510	-4.44×10^{-5}	1.84	-1.57×10^{-5}	-7.17×10^{-7}	
<i>Panel B: announcement DE(0.4, 15, 10) at $t_E = 3$ trading days, $\sigma = 0.3$, $r = 0.04$, $q = 0.08$, $T = 20/252$, $K = 100$</i>								
S_0	M	N	C_{am}	$C_{\text{am}} - P_{\text{am}}^*$	$C_{\text{eu}} - \text{COS}$	$C_{\text{eu}} - P_{\text{eu}}^*$	EEP	
95	400	200	3.095164	$+1.00 \times 10^{-4}$	-1.19×10^{-4}	$+1.15 \times 10^{-4}$	0.032441	
	800	400	3.095354	$+3.85 \times 10^{-5}$	-2.70×10^{-5}	$+3.79 \times 10^{-5}$	0.032538	
	1600	800	3.095400	$+7.72 \times 10^{-6}$	-7.09×10^{-6}	$+8.32 \times 10^{-6}$	0.032565	
100	400	200	5.395761	$+2.00 \times 10^{-4}$	-8.51×10^{-5}	$+2.12 \times 10^{-4}$	0.062213	
	800	400	5.395950	$+5.01 \times 10^{-5}$	-2.13×10^{-5}	$+5.30 \times 10^{-5}$	0.062338	
	1600	800	5.396000	$+1.25 \times 10^{-5}$	-5.31×10^{-6}	$+1.32 \times 10^{-5}$	0.062372	
105	400	200	8.475419	$+3.19 \times 10^{-4}$	$+6.37 \times 10^{-5}$	$+3.13 \times 10^{-4}$	0.107781	
	800	400	8.475505	$+6.85 \times 10^{-5}$	$+1.59 \times 10^{-5}$	$+7.19 \times 10^{-5}$	0.107915	
	1600	800	8.475529	$+1.81 \times 10^{-5}$	$+4.22 \times 10^{-6}$	$+1.90 \times 10^{-5}$	0.107950	

Panel A cases: (1) $\sigma = 0.3$, $r = 0.05$, $q = 0.08$, $T = 0.5$, $S_0 = 100$, $K = 100$; (2) $\sigma = 0.4$, $r = 0.04$, $q = 0.06$, $T = 20/252$, $S_0 = 105$, $K = 100$; (3) $\sigma = 0.35$, $r = 0.05$, $q = 0.1$, $T = 30/252$, $S_0 = 95$, $K = 100$; grids centered at $\ln S_0$. The error is $C_{\text{am}} - V_{\text{ref}}$ with V_{ref} the BBSR value; “EEP error” is the finite-difference premium (American minus European on the same grid) minus the reference premium (BBSR minus the Black–Scholes call). P^* is the put of the symmetric problem $P(K, S_0; r \leftrightarrow q)$, solved on its own grid: $C(S_0, K, r, q) = P(K, S_0, q, r)$. Panel B: P^* carries the law of $-Z$ under the share measure, $\text{DE}(p^*, \eta_{2,E} + 1, \eta_{1,E} - 1)$ with $p^* = (1 - p_E)\eta_{2,E}/((\eta_{2,E} + 1)\mathbb{E}[e^Z])$, here $\text{DE}(0.5600, 11, 14)$; the two COS values agree to 5.0×10^{-12} . Splitting (IT) throughout, as in the empirical study.

$S_0 = 90$, next to the exercise boundary, the American ratios are irregular for both LCP solvers, as the joint ladder of Appendix B.11 is. The deeper Brennan–Schwartz ladders continue with ratios of 2.57 to 2.81 at $S_0 = 100$, so the time-converged reference, the spatial Richardson value of the $N = 25600$ solutions on $M = 3200$ and 6400, carries a temporal residual that extending each time ladder by its last ratio changes by only 3×10^{-7} . Panel C gives the joint ladders of the scheduled-jump case for the European put, the American put by both LCP solvers and the early-exercise premium, with observed orders against the references, and Panel D the split of the error into its temporal and spatial parts.

B.13 American calls with a dividend yield

Table 17 checks the code paths of the American call: the upper far field $\max\{S_M e^{-q\tau} - K e^{-r\tau}, S_M - K\}$, the factors $e^{-q\tau}$ in the boundary data and in the tails, the call’s exercise in the splitting, and, with an announcement, the call’s announcement tails. Panel A uses no jumps, so the binomial reference applies; the BBSR values there change by at most 1×10^{-6} from 20 000 to 40 000 steps. Panel B uses the symmetry $C(S, K, r, q; Z) = P(K, S, q, r; Z^*)$, where Z^* is the law of $-Z$ under the measure with density $e^Z/\mathbb{E}[e^Z]$: for $Z \sim \text{DE}(p_E, \eta_{1,E}, \eta_{2,E})$ it is $\text{DE}(p^*, \eta_{2,E} + 1, \eta_{1,E} - 1)$ with $p^* = (1 - p_E)\eta_{2,E}/((\eta_{2,E} + 1)\mathbb{E}[e^Z])$, and the compensated jump of the symmetric problem, $-Z + \kappa_E$, is the compensated jump of Z^* . The cosine prices of the call and of the symmetric put agree to 5×10^{-12} , which checks the identity, and the finite-difference

Table 18: Volatility sweep with background jumps: FD against COS for the European put, Ikonen–Toivanen against PSOR for the American put, monotonicity and constraint checks, and $\lambda\Delta\tau$.

<i>Panel A: European and American put prices, $(M, N) = (800, 200)$</i>							
σ	Δx	$\lambda\Delta\tau$	$\sigma^2\Delta\tau/\Delta x^2$	$V_{\text{eu}}^{\text{FD}}$	$V_{\text{eu}}^{\text{COS}}$	FD–COS	$V_{\text{am}}^{\text{IT}}$
0.3	0.00135	5.36×10^{-5}	8.81	2.182484	2.182648	-1.64×10^{-4}	2.193066
0.6	0.00220	5.36×10^{-5}	13.27	4.435148	4.435370	-2.22×10^{-4}	4.444842
1.0	0.00333	5.36×10^{-5}	16.06	7.436626	7.436938	-3.11×10^{-4}	7.445962
1.5	0.00475	5.36×10^{-5}	17.79	11.174751	11.175178	-4.28×10^{-4}	11.184074
<i>Panel B: diagnostics; IT and PSOR on the $(200, 50)$ grid, monotonicity on the $(800, 200)$ grid</i>							
σ	IT(200, 50)	PSOR(200, 50)	$\max \text{IT–PSOR} $	sweeps	$\max \Delta_i V_{\text{eu}}$	$\max \Delta_i V_{\text{am}}$	mass err. trap
0.3	2.190213	2.190138	2.56×10^{-4}	13.1/19	-7.79×10^{-8}	-7.83×10^{-8}	2.1×10^{-4}
0.6	4.441007	4.440950	3.49×10^{-4}	14.8/20	-9.86×10^{-10}	-9.87×10^{-10}	3.6×10^{-4}
1.0	7.440689	7.440652	2.73×10^{-4}	17.1/22	-2.12×10^{-11}	-2.12×10^{-11}	5.7×10^{-4}
1.5	11.176985	11.176924	3.74×10^{-4}	18.1/22	-2.75×10^{-11}	-2.75×10^{-11}	8.5×10^{-4}

$\lambda = 0.3$, $p = 0.4$, $\eta_1 = 25$, $\eta_2 = 20$, $T = 9/252$, $S_0 = K = 100$, $r = 0.05$, $q = 0$. $\max \Delta_i V$ is the largest increase of the put value between adjacent nodes (a put is non-increasing in S , so a non-positive value is correct); the payoff constraint $V_{\text{am}} \geq \varphi$ and the European lower bound hold on every grid (maximum violations 0 , 8.5×10^{-8} , 0 , 8.5×10^{-14}). The hat-weight mass error is below 10^{-13} in all four cases; the trapezoid value is listed. Kernel-resolution demonstration ($\lambda = 40$, $p = 24/45$ so $\kappa = 0$, $\eta_1 = 25$, $\eta_2 = 20$, $\sigma = 0.1$, $T = 1$, $L = 3$, $N = 200$, COS value 13.509972): relative mass error trapezoid vs hat $M = 100$: 1.5×10^{-1} vs 1.6×10^{-15} , $M = 200$: 4.4×10^{-2} vs 1.6×10^{-15} , $M = 400$: 1.8×10^{-2} vs 5.2×10^{-15} , $M = 800$: 8.4×10^{-3} vs 7.5×10^{-15} , $M = 1600$: 4.0×10^{-3} vs 8.7×10^{-15} ; price error $M = 100$: 6.1×10^3 vs 5.8×10^{-1} , $M = 200$: 5.2×10^1 vs 1.5×10^{-1} , $M = 400$: 6.6×10^0 vs 3.7×10^{-2} , $M = 800$: 1.4×10^0 vs 7.0×10^{-3} , $M = 1600$: 3.4×10^{-1} vs 4.4×10^{-4} ; hat weights with $N = M/2$ instead of $N = 200$: 5.5×10^{-1} , 1.4×10^{-1} , 3.7×10^{-2} , 9.2×10^{-3} , 2.3×10^{-3} . Stability demonstration ($\lambda = 400$, $M = 200$, no Rannacher, COS value 43.746436): $\lambda\Delta\tau = 0.4$: $\max|z| = 0.9999$, $V(100) = 43.7524$, $\lambda\Delta\tau = 0.8$: $\max|z| = 0.9998$, $V(100) = 43.7512$, $\lambda\Delta\tau = 1.3$: $\max|z| = 0.9996$, $V(100) = 43.7483$, $\lambda\Delta\tau = 2.0$: $\max|z| = 0.9995$, $V(100) = 43.7426$, $\lambda\Delta\tau = 4.0$: $\max|z| = 0.9989$, $V(100) = 43.7106$; no blow-up at $\lambda\Delta\tau > 1$, as the kernel-specific scan predicts.

values, computed on two different grids, approach each other as the grids are refined.

B.14 Volatility sweep, kernel resolution and stability

Table 18 moves the background parameters toward the regime of the empirical study: $T = 9/252$, $\lambda = 0.3$, $p = 0.4$, $\eta_1 = 25$, $\eta_2 = 20$ (mean jump sizes of 4 and 5 percent), and σ from 0.3 to 1.5, on a single $(800, 200)$ grid so that $\sigma^2\Delta\tau/\Delta x^2$ grows from 8.8 to 17.8 and Δx from 0.00135 to 0.00475 as the half-width grows with σ . On this one grid the European put differs from the COS value by -1.64×10^{-4} , -2.22×10^{-4} , -3.11×10^{-4} and -4.28×10^{-4} , so the error is at most 4.3×10^{-4} up to $\sigma = 1.5$; since Δx and the parabolic ratio rise together, the sweep does not say which of them drives the error, and it is not a convergence test. The hat-weight mass error stays below 1.6×10^{-14} while the trapezoid value grows from 2.1×10^{-4} to 8.5×10^{-4} with $\eta_2\Delta x$. The put is nonincreasing in S on every grid, the largest increase between neighbors being -7.8×10^{-8} (a negative number, that is, no increase) for the European and American values at $\sigma = 0.3$ and smaller at higher σ ; the American value never falls below the payoff; and the European value falls below its lower bound $Ke^{-rT} - S$ by at most 8.5×10^{-8} , at $\sigma = 0.6$. On the coarse $(200, 50)$ grid the splitting and PSOR solutions differ by 2.6×10^{-4} to 3.7×10^{-4} over the grid, the splitting error of that grid (Section 4.5), with 13.1 to 18.1 sweeps per step, the count rising with σ . Nothing in the sweep degrades at $\sigma = 1.5$.

The kernel-resolution demonstration in the note to the table and in Figure 4 is where the two quadratures part company. With $\lambda = 40$, $T = 1$ and $L = 3$ on $M = 100$ to 1600 intervals, $\eta_2\Delta x$ runs from 1.2 to 0.075, so the kernel is unresolved on the coarse grids. The hat weights conserve mass to 8.7×10^{-15} at every M . With $N = 200$ fixed, the setting of the comparison with the




Figure 4: Hat weights against trapezoid weights in the kernel-resolution demonstration of Table 18 ($\lambda = 40$, $p = 24/45$, $\eta_1 = 25$, $\eta_2 = 20$, $\sigma = 0.1$, $T = 1$, $L = 3$, $N = 200$). (a) Relative mass error $|\sum_k w_k - 1|$ of the discrete jump operator, maximum over rows, against $\eta_2 \Delta x$: the hat weights conserve mass to 10^{-14} at every resolution; the trapezoid weights do not, and the dotted line is the interior law $(\eta_2 \Delta x)^2/12$ of Lemma 4.2(iii) (at small $\eta_2 \Delta x$ the trapezoid curve lies above it because the maximum is attained in the two boundary rows, Appendix B.10). (b) Price error against the COS value 13.509972 at $S = K$: the spurious growth rate $\lambda \varepsilon$ of the trapezoid weights produces 6114 at $M = 100$ and 65.2 at $M = 200$ for an option worth 13.5. With hat weights and $N = 200$ fixed, as for the trapezoid, the error falls from 0.58 to 4.4×10^{-4} but not at a constant rate, since the fixed time step leaves a temporal error that the finer grids expose; with $N = M/2$ it falls from 0.55 to 2.3×10^{-3} at observed orders 1.92 to 2.00.

trapezoid, their price error falls from 0.58 to 4.4×10^{-4} toward the COS value 13.509972, by factors of 3.82, 4.17, 5.23 and 16.0, which is not order two: the fixed step leaves a temporal error of the opposite sign, which cancels a growing part of the spatial error and on the finest grid outweighs it. With $N = M/2$ the hat-weight error falls from 0.55 to 2.3×10^{-3} at observed orders 1.92 to 2.00. The trapezoid weights carry a relative mass excess of 0.151, 0.044, 0.018, 0.0084 and 0.0040, which enters the scheme as a growth rate $\lambda \varepsilon$ of 6 per year on the coarsest grid: the price there is 6114, the maximum of the grid function 33238, and even at $M = 1600$ the value 13.85 is off by 0.34. The point of Lemma 4.2(i) is that the hat weights remove this failure mode without any condition on $\eta \Delta x$; what remains is the ordinary second-order interpolation error.

The stability demonstration takes $\lambda = 400$ on an $M = 200$ grid with $L = 3$, without the Rannacher start, and reduces N from 1000 to 100 so that $\lambda \Delta \tau$ rises from 0.4 to 4, four times the generic bound of Lemma 4.3. The kernel-specific scan predicts largest moduli 0.9999, 0.9998, 0.9996, 0.9995 and 0.9989, and the solver agrees: no entry overflows, the maximum of the grid function stays at 90.1, and the price moves smoothly from 43.7524 to 43.7106 against the COS value 43.746436 as the time step coarsens. The cell Péclet number of this run is 0.135, well inside the M-matrix condition (25).

B.15 Sign and time direction

Remark B.1 (Pitfalls). Each of the following produces a scheme that runs and returns plausible-looking numbers, and each was present in an earlier implementation of this problem; the verification suite of Section 5 was designed to catch them. (1) *Time direction.* With $\tau = T - t$ the implicit matrix is $I - \frac{\Delta \tau}{2} D$ and the inequality in (15) is $\partial_\tau u - \mathcal{L}u \geq 0$; $I + \frac{\Delta \tau}{2} D$ marches the

Table 19: Sign and time-direction variants of the Crank–Nicolson step on the coarsest grid of Table 15 ($M = 100$, $N = 50$, no Rannacher start, $\lambda = 0$; Black–Scholes put value 7.165868).

variant	Eur. $V(100)$	$\max V $	Am. IT $V(100)$	Am. PSOR $V(100)$	sweeps	LHS diag.
correct	7.15128	7.0×10^1	7.37781	7.37675	14.2/20	1.6947
wrong RHS sign	-7.18×10^{-17}	7.1×10^1	$+2.24 \times 10^{-18}$	$+4.85 \times 10^{-16}$	1.0/1	1.6947
time reversed	-2.96×10^{113}	3.0×10^{113}	0	0	30.9/33	0.3053

Schemes: correct: $(I - \frac{\Delta\tau}{2}D)\mathbf{V}^{n+1} = (I + \frac{\Delta\tau}{2}D)\mathbf{V}^n$; wrong RHS sign: $(I - \frac{\Delta\tau}{2}D)\mathbf{V}^{n+1} = (I - \frac{\Delta\tau}{2}D)\mathbf{V}^n$; time reversed: $(I + \frac{\Delta\tau}{2}D)\mathbf{V}^{n+1} = (I - \frac{\Delta\tau}{2}D)\mathbf{V}^n$. The wrong RHS sign cancels the differential terms and collapses the European solution to zero (and the American one to the payoff); the time-reversed scheme is anti-diffusive and blows up. PSOR capped at 200 sweeps; an M-matrix needs a positive, dominant diagonal.

backward equation forward, amplifies the stiff modes and is not an M-matrix, and the reversed inequality empties the exercise region. (2) *Splitting*. The drift and $-(r + \lambda)u$ belong to the implicit part; treating $-(r + \lambda)u$ explicitly applies Adams–Bashforth to $v' = -(r + \lambda)v$, whose stability interval is bounded, and nothing is gained since the implicit matrix is tridiagonal either way. (3) *Log transform*. The drift coefficient contains $-\frac{1}{2}\sigma^2$; dropping it biases the price by an amount that grows with $\sigma^2 T$. (4) *FFT evaluation*. Sampled kernel values need the factor Δx (integrated weights do not); the kernel must be reversed; the transforms must be zero-padded to at least $3M + 1$ points, or in-the-money values wrap around onto the out-of-the-money end; and the solution outside the grid is the far-field form, not zero. (5) *Lower boundary*. On a log grid the lower edge is at a positive price, so the put is worth $K - e^{x_0}$ (American) or $Ke^{-r\tau} - e^{x_0 - q\tau}$ (European) there, not K . (6) *PSOR test*. Stop on the change of the iterate in a relative max norm, as in (40), not on the projection amount. (7) *Units*. λ, r, q per year and T in years; report λT with every parameter set. (8) *Start-up*. Crank–Nicolson without (34) keeps the oscillation excited by the kink: the temporal order fluctuates and the gamma oscillates near the strike (Figure 5).

Table 19 is a diagnostic rather than a test. It records what the scheme produces when the Crank–Nicolson step is written in the wrong direction, on the coarsest grid of Table 15 and without the Rannacher start, so that a reader who adapts the code can recognize the symptoms. The correct step gives 7.15128 for the European put (the closed form is 7.165868 and the grid is coarse), 7.37781 for the American put by splitting and 7.37675 by PSOR, the difference being the splitting error of this coarse grid, with 14.2 sweeps per step on average. With the sign of the right-hand side reversed, $(I - \frac{\Delta\tau}{2}D)$ on both sides, the differential terms cancel: the European value at the money collapses to -7.2×10^{-17} , and the American solvers return the payoff, PSOR stopping after a single sweep because the iterate does not move. With the time direction reversed, $(I + \frac{\Delta\tau}{2}D)$ on the left, the diagonal of the implicit matrix is 0.3053 instead of 1.6947, the matrix is no longer an M-matrix, and the march amplifies the stiff modes: after fifty steps the European value at the money is -2.96×10^{113} , the largest entry of the grid function is 3.0×10^{113} , and the projected solvers return zero at the money with entries up to 5.5×10^{85} elsewhere, PSOR taking 30.9 sweeps per step. Each of these mistakes announces itself on the coarsest grid in a fraction of a second once the three numbers in the table are looked at; the point of recording them is that the pitfalls of Remark B.1 do not announce themselves in a code that does not compute such diagnostics.

B.16 The start-up and the gamma

Figure 5 reports what the Rannacher start buys. On the $(400, 200)$ grid, where $\sigma^2 \Delta\tau / \Delta x^2 = 0.36$, the gamma from second differences agrees with the COS reference differentiated with the same stencil to 1.51×10^{-4} whether or not the start-up is used, the second difference of the gamma along the grid is 7.3×10^{-4} in both and 7.2×10^{-4} in the reference, and the gamma at the strike is 0.051237 against 0.051089. On the $(800, 25)$ grid, where $\sigma^2 \Delta\tau / \Delta x^2 = 11.6$, the plain




Figure 5: Gamma of the Kou European put (parameters of Table 1) from second differences on the grid, Crank–Nicolson with a plain start and with four Rannacher half steps, against the COS price differentiated with the same stencil. Left: $(M, N) = (400, 200)$, $\sigma^2 \Delta \tau / \Delta x^2 = 0.36$; the three curves coincide (largest gamma error 1.51×10^{-4} with or without the start-up). Right: $(800, 25)$, $\sigma^2 \Delta \tau / \Delta x^2 = 11.6$; the plain start leaves an oscillation at the strike (gamma error 7.35×10^{-3} , gamma at K of 0.043777 against 0.051132), which the Rannacher start removes (error 5.73×10^{-5} , gamma at K of 0.051187).

Crank–Nicolson start leaves the oscillation that [Giles and Carter \(2006\)](#) describe: the gamma error is 7.35×10^{-3} , the second difference along the grid 2.71×10^{-2} against 1.82×10^{-4} in the reference, and the gamma at the strike 0.043777 against 0.051132. Four implicit-Euler half steps bring these to 5.73×10^{-5} , 1.83×10^{-4} and 0.051187. The delta is hardly affected (3.4×10^{-4} against 2.5×10^{-4} on the coarse time grid), and the price is not (Section 5.2): the start-up is a matter of the derivatives, and it is kept on in every run of the paper, in both legs when there is an announcement.

With the announcement. The solver starts the second leg with an empty Adams–Bashforth history whatever the number of half steps, because the history from before the jump condition is invalid after it. The half steps themselves are hardly needed there to smooth the gamma: the value after the announcement step is a convolution with the jump density, which largely removes the oscillation that a plain start leaves. On the $(800, 25)$ grid without any half steps the second difference of the gamma along the grid is 4.5×10^{-5} against 2.4×10^{-5} in the reference, where without an announcement a plain start leaves 149 times the reference value (Figure 5). The errors agree. For the European put of the base case, against the COS reference differentiated with the same stencil on the same nodes, the largest delta, gamma and theta errors for $90 < S < 110$ on the $(800, 400)$ grid are 5.0×10^{-6} , 1.6×10^{-6} and 3.2×10^{-4} with four half steps at the start of each leg and 5.0×10^{-6} , 1.6×10^{-6} and 3.2×10^{-4} with none. On the $(800, 25)$ grid, where $\sigma^2 \Delta \tau / \Delta x^2 = 24.2$, no half steps give 3.0×10^{-6} , 9.6×10^{-6} and 4.5×10^{-4} , and four give larger errors, 3.0×10^{-5} , 1.0×10^{-5} and 2.9×10^{-3} , the first-order error of the implicit-Euler half steps at a coarse time step. The theta is the backward difference $-(\mathbf{V}^N - \mathbf{V}^{N-1}) / \Delta \tau$ over the last step, compared with the same difference of the COS prices (about -10.40 per year at the money). For the American put at $S = K$, checked by self-convergence, the delta, gamma and theta move from -0.482296 , 0.040218 and -10.459 on the $(400, 200)$ grid to -0.482319 , 0.040219 and -10.455 on the $(3200, 1600)$ grid, the last changes being 1×10^{-6} , 1×10^{-8} and 7×10^{-4} .

Table 20: Scheduled announcement jump, European options: IMEX-CNAB with the jump condition against the product-CF COS benchmark, put–call parity, and the parity violation of the uncompensated variant.

M	N	Δx	European put			European call		
			V_{FD}	error	order	V_{FD}	error	order
200	100	0.00783	4.61064991	6.04×10^{-4}	–	4.75391735	9.18×10^{-5}	–
400	200	0.00392	4.61110319	1.51×10^{-4}	2.00	4.75398642	2.27×10^{-5}	2.01
800	400	0.00196	4.61121632	3.77×10^{-5}	2.00	4.75400349	5.68×10^{-6}	2.00
1600	800	0.00098	4.61124459	9.42×10^{-6}	2.00	4.75400775	1.42×10^{-6}	2.00
3200	1600	0.00049	4.61125166	2.36×10^{-6}	2.00	4.75400881	3.55×10^{-7}	2.00
COS, product CF			4.61125402			4.75400917		
mixture (BS)			4.61125402			4.75400917		
BS, no event			1.81278775			1.95554290		
event premium			+2.798466			+2.798466		
<i>Put–call parity</i>								
compensated, grid (1600, 800): $C - P = 0.14276316$, $S_0 - Ke^{-rT} = 0.14275515$, difference $+8.01 \times 10^{-6}$								
uncompensated (κ_E dropped), grid (800, 400): put 4.285398, call 5.127486, $C - P = 0.842088$, parity violation $+0.6993$								

$\sigma = 0.25$, $r = 0.04$, $q = 0$, $T = 9/252$, $S_0 = K = 100$, one compensated jump $Z \sim \text{DE}(0.5, 12, 12)$ at $t_E = T/2$ ($\mathbb{E}|Z| = 0.083333$, $\text{sd}(Z) = 0.117851$, $\kappa_E = 0.00696867$); two IMEX-CNAB legs with Rannacher restarts on the same grid; the error is against the product-CF COS value ($N = 2^{14}$, $L_c = 12$). “mixture (BS)” is the adaptive quadrature of Black–Scholes prices over the jump law, “BS, no event” the Black–Scholes value without the announcement, and “event premium” their difference COS – BS. Jump-condition operator self-checks (grid centered at $\ln K$; mass = $\max|P\mathbf{1} + \text{tails} - 1|$, martingale = $\max|(Pe^x + \text{tails})/e^x - 1|$, quadrature = $|P[V_{\text{BS}}(\cdot, \tau_E)](\ln K) - \text{mixture}|$): $M = 200$: mass 1.3×10^{-15} , martingale 5.1×10^{-6} , quadrature 2.03×10^{-3} ; $M = 400$: mass 2.4×10^{-14} , martingale 1.3×10^{-6} , quadrature 5.07×10^{-4} ; $M = 800$: mass 2.4×10^{-14} , martingale 3.2×10^{-7} , quadrature 1.27×10^{-4} ; $M = 1600$: mass 6.9×10^{-14} , martingale 8.0×10^{-8} , quadrature 3.17×10^{-5} . Asymmetric law (0.3, 20, 10): $(Pe^x)/e^x$ at $\ln K$ is 1.00000041 (a reflected kernel would give 1.012608 = $\mathbb{E}[e^Z]\mathbb{E}[e^{-Z}]$); without compensation the ratio is 1.00699333 = $\mathbb{E}[e^Z]$. The uncompensated parity violation $+0.6993$ equals $S_0(\mathbb{E}[e^Z] - 1) = +0.6993$. Moving t_E/T over $\{0.10, 0.25, 0.50, 0.75, 0.90\}$ changes the European put on the (800, 400) grid by at most 8.0×10^{-11} ; doubling the domain half-width changes it by -6.5×10^{-10} .

B.17 The scheduled jump: the European benchmark, operator identities and skew orientation

Table 20 gives the European ladders of the base case of Section 5.4 against the product-characteristic-function benchmark, put–call parity on the grid, and the parity violation of the uncompensated variant.

The announcement operator (36) of Table 20 is also tested on its own, without time stepping. Applied to the constant function it returns one to 6.9×10^{-14} or better on all four grids (Lemma 4.2(i) for the shifted density); applied to e^x it returns e^x with a maximum relative error of 5.1×10^{-6} , 1.3×10^{-6} , 3.2×10^{-7} and 8.0×10^{-8} on $M = 200$ to 1600, falling by a factor of 4.00 per halving of Δx , the interpolation error of e^x promised in Section 4.4; and applied to the exact Black–Scholes value at τ_E it reproduces the exact mixture value at $\ln K$ to 2.03×10^{-3} , 5.07×10^{-4} , 1.27×10^{-4} and 3.17×10^{-5} , again second order. With the asymmetric law (0.3, 20, 10) the ratio $(\mathcal{P}_E^h e^x)/e^x$ at $\ln K$ is 1.00000041, where the reflected kernel applied to e^x would give 1.012608 = $\mathbb{E}[e^Z]\mathbb{E}[e^{-Z}]$, a deviation more than four orders of magnitude above the grid error. Because $\ln(\mathbb{E}[e^Z]\mathbb{E}[e^{-Z}])$ involves only the even cumulants of Z , this identity catches a reflection through the reversed compensator but cannot see the asymmetry of the law, which Figure 6 checks. Dropping the compensator makes the ratio 1.00699333 = $\mathbb{E}[e^Z]$ and, on the (800, 400) grid, prices the put at 4.285398 and the call at 5.127486, so that $C - P = 0.842088$ against the parity value 0.142755: the violation $+0.6993$ equals $S_0(\mathbb{E}[e^Z] - 1) = +0.6993$, which is the failure that (6) forbids, now in numbers. Moving t_E/T over 0.10, 0.25, 0.50, 0.75 and 0.90




Figure 6: Black-Scholes implied volatility of European options with the asymmetric announcement law $Z \sim \text{DE}(0.3, 20, 10)$ (down-move probability 0.70, mean down move 0.100, mean up move 0.050, skewness -1.256) at $t_E = T/2$, with $\sigma = 0.25$, $r = 0.04$, $T = 9/252$, $S_0 = 100$; puts for $K \leq 100$, calls above. Lines are COS values from the product characteristic function, markers the finite-difference values on the $(800, 400)$ grid, which agree with them to within 10^{-5} in implied volatility. The law of Z produces a put skew, $\text{IV}(90) - \text{IV}(110) = +0.133$; the mirrored law of $-Z$ produces the opposite, -0.153 ; the symmetric law $\text{DE}(0.5, 12, 12)$ produces a smile whose lowest value among the plotted strikes, 0.622, is at the money. That the markers follow the lines of the correct law, and not of the mirrored one, checks the orientation of the announcement operator (36).

changes the European put on the $(800, 400)$ grid by at most 8.0×10^{-11} , Proposition 2.3(v) on the grid; doubling the half-width changes it by -6.5×10^{-10} ; and setting the announcement reach to 4, 6 or 8 mean moves at fixed $M = 800$ gives errors of -2.35×10^{-5} , -3.77×10^{-5} and -5.52×10^{-5} with $\Delta x = 0.00154$, 0.00196 and 0.00238, that is, the error follows Δx^2 and the reach itself is immaterial.

Figure 6 uses the asymmetric law $\text{DE}(0.3, 20, 10)$: a down move with probability 0.70 of mean size 0.100 against an up move of mean size 0.050, so $\mathbb{E}[Z] = -0.055$, the skewness is -1.256 and $\mathbb{E}[e^Z] = 0.952153$, which makes $\kappa_E = -0.049029$ negative and the compensation an upward shift. The implied volatilities of the European options at $S_0 = 100$ show a put skew, $\text{IV}(90) - \text{IV}(110) = +0.13346$ and $\text{IV}(95) - \text{IV}(105) = +0.07179$; the mirrored law of $-Z$, with $\kappa_E = +0.061558$, gives -0.15333 and -0.08565 ; and the symmetric base law gives a smile with 0.65113, 0.62185 and 0.66321 at $K = 90, 100$ and 110. The finite-difference prices on the $(800, 400)$ grid differ from the COS prices by at most 4.8×10^{-5} over the fourteen options, which is 7×10^{-6} or less in implied volatility. A reflected announcement kernel would have produced the law of $-Z$ shifted by $+\kappa_E$ instead of by its own compensator, that is, the mirrored skew with the forward price multiplied by $\mathbb{E}[e^Z]\mathbb{E}[e^{-Z}]$; the solver produces the correct curve.

A sweep over the scale of the announcement law that is in the JSON but not tabulated records that the at-the-money straddle of the base case rises from 3.8 to 17.4 percent of S_0 , and its implied volatility from 0.25 to 1.16, as the mean absolute move rises from zero to 0.167.

Table 21: The published benchmark with background jumps and an announcement: American and European puts against [Leung and Santoli \(2014\)](#).

<i>Set 1: $\sigma = 0.2$, $\lambda = 252$, $p = 0.5$, $\eta_1 = \eta_2 = 300$; announcement</i>									
DE(0.5, 30, 30)									
K	LS, $T - 2d$	FD	LS, $T/2$	FD	LS, 3d	FD	LS, Eu.	COS	
80	0.10	0.1042	0.10	0.1045	0.10	0.1045	0.10	0.1040	
85	0.37	0.3672	0.37	0.3686	0.37	0.3686	0.37	0.3665	
90	1.02	1.0177	1.02	1.0222	1.02	1.0224	1.02	1.0155	
95	2.30	2.3006	2.31	2.3124	2.31	2.3132	2.29	2.2946	
100	4.40	4.3991	4.42	4.4229	4.43	4.4262	4.38	4.3832	
105	7.35	7.3549	7.39	7.3921	7.40	7.4029	7.32	7.3163	
110	11.07	11.0685	11.11	11.1121	11.14	11.1397	10.98	10.9841	
115	15.36	15.3627	15.40	15.3990	15.45	15.4510	15.20	15.1980	
120	20.06	20.0560	20.07	20.0731	20.14	20.1374	19.77	19.7665	
<i>Set 2: $\sigma = 0.07$, $\lambda = 200$, $p = 0.5$, $\eta_1 = \eta_2 = 350$; announcement</i>									
DE(0.5, 25, 25)									
K	LS, $T - 2d$	FD	LS, $T/2$	FD	LS, 3d	FD	LS, Eu.	COS	
80	0.01	0.0104	0.01	0.0106	0.01	0.0107	0.01	0.0103	
85	0.05	0.0501	0.05	0.0511	0.05	0.0515	0.05	0.0499	
90	0.22	0.2198	0.22	0.2243	0.23	0.2263	0.22	0.2192	
95	0.84	0.8423	0.86	0.8599	0.87	0.8681	0.84	0.8400	
100	2.54	2.5380	2.60	2.5958	2.63	2.6264	2.53	2.5317	
105	5.68	5.6820	5.81	5.8088	5.91	5.9058	5.66	5.6644	
110	10.01	10.0112	10.10	10.0976	10.28	10.2798	9.87	9.8689	
115	15.00	15.0000	15.00	15.0000	15.08	15.0779	14.57	14.5719	
120	20.00	20.0000	20.00	20.0000	20.01	20.0093	19.45	19.4539	

LS: the values of Table 4 of the arXiv version of [Leung and Santoli \(2014\)](#) (Fourier transform method for the American put, “EU” for the European put), two decimals. Parameters in the notation of this paper: $S_0 = 100$, $r = 0.02$, $q = 0$, $T = 0.25$; their jump intensity κ is λ here, their rates λ_1, λ_2 are η_1, η_2 , and their announcement law (u, η_1, η_2) is $(p_E, \eta_{1,E}, \eta_{2,E})$; the announcement is 2 trading days before expiry, at $T/2$, or 3 trading days from today, with one trading day = $1/252$ year. FD: Richardson extrapolation of the (1600, 800) and (3200, 1600) grids (centered at $\ln S_0$, splitting). COS: the product characteristic function, $N = 2^{14}$, $L_c = 12$. Largest $|\text{FD} - \text{LS}|$ 0.0049 (54 of 54 within 0.005); largest $|\text{COS} - \text{LS}|$ 0.0046; finite-difference European minus COS at most 3.3×10^{-4} on the finer grid; a day of $1/365$ year changes the American values by at most 8.6×10^{-3} .

B.18 The published benchmark with background jumps and an announcement

Table 4 of the arXiv version of [Leung and Santoli \(2014\)](#) prices American puts under Kou dynamics with a double-exponential announcement jump by a Fourier transform method, for two parameter sets, nine strikes from 80 to 120 and three announcement dates, and lists the European put of each strike; Table 21 transcribes these values from the printed table and sets the solver beside them. The source writes the Poisson intensity as κ , the Kou rates as λ_1, λ_2 and the announcement law as $\text{DE}(u, \eta_1, \eta_2)$, and compensates the announcement jump by $\log \mathbb{E}[e^Z]$, which is the compensator κ_E of (6); the mapping to the notation of this paper is in the table note. The announcement dates are written there as “ $T - 2D$ ”, “1.5M” and “3D”, and I take a day to be one trading day of $1/252$ year and 1.5 months to be $T/2$; with calendar days the comparison is worse (Section 5.4). The cosine values from the product characteristic function reproduce every published European value within its rounding (18 of 18), so the mapping and the compensator are those of the source. The American values are compared at the resolution of the printed digits, half a unit in the second decimal: 54 of the 54 extrapolated values are within it, and the largest deviation, 0.0049, is at $K = 105$ in set 1. The test is weaker than those of Section 5.3, whose references carry six or more digits, but it is the only external check of the combination of background jumps, an announcement and early exercise.




Figure 7: Three models with the same variance 0.016121 of $\ln S_T$ over $T = 9/252$ (Table 4): (a) one compensated jump $\text{DE}(0.5, 12, 12)$ at $t_E = T/2$ on $\sigma = 0.25$; (b) Kou with $\lambda T = 1$ and the same jump law; (c) Black-Scholes with $\sigma = 0.672$. Left: terminal densities of $\ln(S_T/S_0)$ by Fourier inversion of the characteristic functions; the Poisson model has the highest peak (5.215 against 4.007 and 3.138) because it spends probability e^{-1} on no jump at all. Right: European implied-volatility smiles from the COS prices on a grid of unit strike steps; the Gaussian is flat at 0.672; the scheduled-jump smile has its minimum, 0.621, at $K = 99$, one strike below the money, and is 0.622 at $K = 100$; the Poisson model is lower at the money (0.573) and higher at $K = 80$ and $K = 120$ (0.765 and 0.767 against 0.730 and 0.728).

B.19 Three models with the same variance: densities and smiles

Figure 7 shows the terminal laws and the smiles of the three models of Table 4. The three terminal laws differ in shape at equal variance: the excess kurtosis is 2.227, 4.453 and 0; the density at zero is 4.007, 5.215 and 3.138; the probability of a terminal log return within 0.02 of zero is 0.151, 0.195 and 0.119; and the probability of a move beyond 0.05, 0.10 and 0.20 is 0.620, 0.351, 0.106 for the scheduled jump, 0.533, 0.284, 0.105 for the Poisson model and 0.691, 0.429, 0.114 for the Gaussian. At these thresholds the Gaussian has the most mass beyond the threshold and the Poisson model the least; the far-tail advantage of the double-exponential laws, which the kurtosis records, lies beyond the thresholds tabulated. The smiles cross: near the money the Poisson model is below the scheduled jump and the Gaussian above both, and twenty percent out of the money the order reverses, which is the reading of the wing clause of P7 given in Section 5.4.

B.20 Run times of the jump operator and the LCP solvers

Table 22 gives wall-clock times on one machine, taken in a single process after the pooled runs of the script had finished: the median of 50 evaluations of the jump operator and of 5 full pricings. One evaluation of the jump operator costs 35.78 milliseconds as a dense matrix-vector product at $M = 6400$, 2.45 by FFT and 0.28 by the recursion, with set-up costs of 5661, 172 and 127 milliseconds and a dense matrix of 313 megabytes. The FFT overtakes the dense form from $M = 800$ (0.28 against 0.71 milliseconds) and the recursion from $M = 400$ (0.07 against 0.10). A full American pricing on the (400, 200) grid takes 0.444 seconds with the splitting, 0.685 with Brennan-Schwartz and 4.346 with PSOR, the last because the sweeps are pure-Python loops. Panel C sets the run time against the error at $S_0 = 100$, measured against the time-converged exact-LCP reference of Appendix B.12: on the (1600, 800) grid the splitting

Table 22: Wall-clock timings (single process, numpy/scipy; the PSOR and Brennan–Schwartz sweeps are pure-Python loops).

<i>Panel A: one evaluation of the discrete jump operator $\mathcal{J}$ (milliseconds, median of 50; build time in parentheses)</i>					
M	dense	FFT	recursion	dense matrix (MB)	
200	0.02 (3.94)	0.11 (3.09)	0.06 (6.24)	0.3	
400	0.10 (13.26)	0.18 (3.53)	0.07 (6.24)	1.2	
800	0.71 (102.00)	0.28 (4.60)	0.09 (73.19)	4.9	
1600	2.55 (381.73)	0.62 (8.31)	0.11 (81.27)	19.6	
3200	6.57 (1468.74)	0.96 (66.34)	0.17 (88.73)	78.2	
6400	35.78 (5661.44)	2.45 (172.46)	0.28 (126.90)	312.6	
<i>Panel B: full American put pricing (Toivanen parameters, dense $\mathcal{J}$), seconds (median of 5)</i>					
M	N	LCP solver	time (s)	$V(100)$	PSOR sweeps
200	100	Ikonen–Toivanen	0.203	2.785214	–
200	100	Brennan–Schwartz	0.179	2.785147	–
200	100	PSOR	0.869	2.785147	16.3/19
400	200	Ikonen–Toivanen	0.444	2.802126	–
400	200	Brennan–Schwartz	0.685	2.802082	–
400	200	PSOR	4.346	2.802082	14.3/18
<i>Panel C: work and precision at $S_0 = 100$, error against the time-converged exact-LCP reference</i>					
M	N	IT time (s)		BS time (s)	
200	100	0.125	-2.27×10^{-2}	0.120	-2.27×10^{-2}
400	200	0.342	-5.75×10^{-3}	0.455	-5.80×10^{-3}
800	400	1.255	-1.46×10^{-3}	2.245	-1.49×10^{-3}
1600	800	4.877	-3.70×10^{-4}	8.480	-3.86×10^{-4}

Panel A: the dense evaluation is an $(M + 1)^2$ matrix–vector product, the FFT evaluation a zero-padded circular correlation of length $\geq 3M + 1$ with the reversed kernel, and the recursion two first-order recursive filters (Toivanen 2008); all three agree to 10^{-13} (Table 1). Panel B: timings include the operator set-up; the LCP solvers share the same tridiagonal factorization cost. The system-wide processor load during the timing pass, all processes included, was 100 percent of 12 logical processors.

takes 4.88 seconds for an error of 3.7×10^{-4} and Brennan–Schwartz 8.48 seconds for 3.9×10^{-4} , so in this implementation the splitting is both the faster and, on every grid of the panel, the more accurate of the two. The machine was not idle: the system-wide processor load over the timing pass, all processes included, was 100 percent of its twelve logical processors, so the times are indicative only, and the per-grid run times recorded alongside the ladders in the JSON were measured inside the process pool and are less reliable still. For the empirical study, where a few thousand solves of size (800, 400) or smaller are needed per model, the ordering of the columns matters more than their values.

B.21 Run times and software

The full verification takes 1662 seconds of wall clock for the script `verify_core.py` and 209 seconds for the script `verify_scheduled.py`, with a pool of four worker processes for the ladders on a Windows 11 machine with twelve logical processors and 15.8 GiB of memory that was running other work at the same time (Appendix B.20); the (12800, 6400) American put at $S_0 = 90$ with the recursion takes 31.8 seconds inside the pool. The code is pure `numpy` and `scipy`; the runs reported here used Python 3.14.3, numpy 2.4.2, scipy 1.17.1 and matplotlib 3.10.8. The results are deterministic up to floating-point rounding; the only randomness in the suite is the fixed-seed vector used to compare the three jump-operator forms. The timings are the one quantity that is not reproducible to the digit, for the reasons given in Appendix B.20.

B.22 Reproducing the results

From the paper folder, with the project interpreter,

```
python code/scripts/verify_core.py      [--fast] [--workers W]
python code/scripts/verify_scheduled.py [--fast] [--workers W]
python code/scripts/verif_env.py
python code/scripts/verif_macros.py
python -m pytest code/tests -q
```

regenerate the verification tables (`tables/verif_*.tex`) and figures (in `figures/`), the file `output/verification.json` (every number that appears in a table, merged from the two scripts) and the log `output/verification_log.txt`. The third command records the machine, the versions of the software and the run time of the two scripts with `--fast`, which it runs in a temporary copy of `code/` so that nothing in the paper folder is overwritten, in `output/verification_env.json`; the fourth reads these outputs and writes `tables/verif_macros.tex`, the macros through which the text of Section 5, this appendix, the abstract and the conclusion quotes the verification numbers. `verify_core.py` runs first, since the convergence and boundary figures combine data from both scripts and are rendered by whichever runs last; `--fast` shortens the ladders and grids to a check that took 501.6 and 319.8 seconds with the default pool of six workers, and `--workers 1` disables the process pool. The test suite carries the reference numbers of Section 5 as literals. It also contains a comparison of every solver path with an earlier, independent implementation of the scheme to a relative tolerance of 10^{-12} ; that implementation is included in `code/reference/` of the public repository. The scripts of the empirical part (Appendix D) call the same `price_fd` and `price_european_cos` with the defaults stated here: hat weights, dense jump operator, operator splitting, four Rannacher half steps, the half-width rule (54) with $c_E = 6$, and $x_c = \ln K$.

C Estimation: derivations, Monte Carlo studies and definitions

This appendix supports Section 6: it gives the derivations behind the estimators, the three Monte Carlo studies that Section 6 summarizes (the announcement-law estimators, the background jump parameters and the two-expiry decomposition under quote noise), the definitions of the predictive laws, the designs and the firm-level laws. Throughout it, $Z \sim \text{DE}(p, \eta_1, \eta_2)$ denotes an announcement return with the density f_E of (2); the subscript E of Section 2 is dropped when no confusion with the background jump law can arise.

C.1 Moments of the double-exponential law

For $-\eta_2 < u < \eta_1$ the moment generating function is

$$\mathbb{E}[e^{uZ}] = \frac{p\eta_1}{\eta_1 - u} + \frac{(1-p)\eta_2}{\eta_2 + u}, \quad (58)$$

so $\kappa_E = \ln \mathbb{E}[e^Z]$ of (5) is finite exactly when $\eta_1 > 1$, and the moments are $m_n = \mathbb{E}[Z^n] = n! (p\eta_1^{-n} + (1-p)(-1)^n\eta_2^{-n})$, as in Proposition 2.3(ii). The first two give (43) and

$$\text{Var } Z = \frac{2p}{\eta_1^2} + \frac{2(1-p)}{\eta_2^2} - \left(\frac{p}{\eta_1} - \frac{1-p}{\eta_2} \right)^2 = \frac{p(2-p)}{\eta_1^2} + \frac{(1-p)(1+p)}{\eta_2^2} + \frac{2p(1-p)}{\eta_1\eta_2}. \quad (59)$$

Conditional on its sign, Z is exponential: $\mathbb{P}(Z > z \mid Z > 0) = e^{-\eta_1 z}$ and $\mathbb{P}(Z < -z \mid Z < 0) = e^{-\eta_2 z}$ for $z \geq 0$, so $\mathbb{E}|Z| = p/\eta_1 + (1-p)/\eta_2$. In the symmetric case $p = \frac{1}{2}$, $\eta_1 = \eta_2 = \eta$, the law is Laplace with $\text{Var } Z = 2/\eta^2$, $\mathbb{E}|Z| = 1/\eta = \text{sd}(Z)/\sqrt{2}$ and excess kurtosis $m_4/m_2^2 - 3 = 3$; a Gaussian with the same standard deviation has $\mathbb{E}|Z| = \text{sd}(Z)\sqrt{2/\pi}$, which is larger.

Lemma C.1 (Location). *Let $c \in \mathbb{R}$ and $Z' = Z + c$. Then $\ln \mathbb{E}[e^{Z'}] = \kappa_E + c$, hence $Z' - \ln \mathbb{E}[e^{Z'}] = Z - \kappa_E$, and every price of the Kou-SEJ model is the same under Z' as under Z .*

Proof. $\mathbb{E}[e^{Z'}] = e^c \mathbb{E}[e^Z]$. The announcement enters the dynamics (4), the jump condition (13) and the characteristic function (9) only through the law of $Z - \kappa_E$. $\square$

Lemma C.1 is why the estimators of Section 6 match the variance of the announcement law rather than its second moment: only the centered law enters a price, and a constant in the announcement return is absorbed by the compensator.

C.2 The closed-form estimator and its small-sample law

For n independent draws $z_1, \dots, z_n$ of Z , with n_+ positive, $n_- = n - n_+$ negative, $S_+ = \sum_{z_i > 0} z_i$ and $S_- = \sum_{z_i < 0} |z_i|$, the log-likelihood is

$$\ell(p, \eta_1, \eta_2) = n_+ \ln p + n_- \ln(1 - p) + n_+ \ln \eta_1 - \eta_1 S_+ + n_- \ln \eta_2 - \eta_2 S_-. \quad (60)$$

It separates into three one-parameter problems, whose maximizers are (44); (n_+, S_+, S_-) is sufficient. Given $n_+ = k \geq 1$, S_+ is a sum of k independent exponential variables with rate η_1 , a gamma variable with shape k and rate η_1 , and $\mathbb{E}[S_+^{-j}] = \eta_1^j \Gamma(k - j) / \Gamma(k)$ for $j < k$. Hence, for the maximum-likelihood rate $\hat{\eta}_1 = k/S_+$ and the corrected rate $\tilde{\eta}_1 = (k - 1)/S_+$,

$$\begin{aligned} \mathbb{E}[\hat{\eta}_1 \mid k] &= \eta_1 \frac{k}{k-1} \quad (k \geq 2), & \text{Var}(\hat{\eta}_1 \mid k) &= \eta_1^2 \frac{k^2}{(k-1)^2(k-2)} \quad (k \geq 3), \\ \mathbb{E}[\tilde{\eta}_1 \mid k] &= \eta_1 \quad (k \geq 2), & \text{Var}(\tilde{\eta}_1 \mid k) &= \frac{\eta_1^2}{k-2} \quad (k \geq 3), \end{aligned} \quad (61)$$

and the same with η_2 , S_- and n_- . For $k = 1$, $\mathbb{E}[S_+^{-1}] = \int_0^\infty s^{-1} \eta_1 e^{-\eta_1 s} ds = \infty$, and for $k = 2$, $\mathbb{E}[S_+^{-2}] = \infty$. Since n_+ is binomial with $\mathbb{P}(n_+ = 1) = np(1-p)^{n-1} > 0$, neither rate estimator has a finite unconditional mean or variance at any n , Monte Carlo averages over all samples do not converge, and the relative bias in (61) is $1/(k-1)$ conditional on k . Table 23 therefore reports means conditional on $n_\pm \geq 2$ and standard deviations conditional on $n_\pm \geq 3$, together with the exact conditional values, which mix (61) over the binomial probabilities $b_k = \binom{n}{k} p^k (1-p)^{n-k}$; for example

$$\mathbb{E}\left[\frac{\hat{\eta}_1}{\eta_1} - 1 \mid n_+ \geq 2\right] = \frac{\sum_{k \geq 2} b_k / (k-1)}{\sum_{k \geq 2} b_k}, \quad \mathbb{E}\left[\frac{\hat{\eta}_1^2}{\eta_1^2} \mid n_+ \geq 3\right] = \frac{\sum_{k \geq 3} b_k k^2 / ((k-1)(k-2))}{\sum_{k \geq 3} b_k}, \quad (62)$$

which, with the analogous conditional first moment given $n_+ \geq 3$, give the exact bias +46.2% and the exact standard deviation of 1.10 times the truth at $n = 8$ quoted in Appendix C.3. The fourth moment of $\hat{\eta}_1$ is finite only for $k \geq 5$, so the Monte Carlo standard deviation is itself imprecise at $n = 8$.

Two interval estimators follow. The Wald interval $\hat{\eta}_1 \pm 1.645 \hat{\eta}_1 / \sqrt{n_+}$ uses the Fisher information n_+ / η_1^2 of (60). The exact interval uses the pivot $2\eta_1 S_+ \sim \chi_{2k}^2$:

$$\left[\frac{\chi_{2k}^2(0.05)}{2S_+}, \frac{\chi_{2k}^2(0.95)}{2S_+} \right] \quad (63)$$

covers η_1 with probability exactly 0.90 given $n_+ = k$, and hence unconditionally. For p the table uses the Clopper-Pearson interval from the quantiles of the beta law. When the observed move contains diffusion, $x = Z + \varepsilon$, the sum of the positive observed moves has a density that is positive near zero, and the same moments are infinite.

C.3 Accuracy of the announcement-law estimators

Table 23 simulates samples of $n = 8, 12, 20$ and 40 announcement returns from the fixed law $DE(0.45, 20, 16)$, whose mean moves are 5.00% up and 6.25% down and whose standard deviation is 8.00%, close to the pooled $\hat{\sigma}_J$ of Table 5; Figure 8 extends the bias and the coverage to a finer grid of n (design in Appendix C.15). Six findings matter for the pricing models.

First, the announcement variance, which is the input of DJ-hist and the scale of Kou-SEJ-P, is unbiased but imprecise: in Panel A the standard deviation of $\hat{\sigma}_J^2$ was 0.81 times the truth at $n = 8$, 0.50 times at $n = 20$ and 0.36 times at $n = 40$. At the history lengths of the sample, the announcement variance of a single firm is known only to within a large fraction of its own size.

Second, the rates of (44) have no finite mean and no finite variance at any n , because $n_+ = 1$ and $n_+ = 2$ have positive probability and, given $n_+ = k$, $\sum_{z_i > 0} z_i$ has a gamma law whose reciprocal has an infinite mean for $k = 1$ and an infinite variance for $k \leq 2$ (Appendix C.2). Monte Carlo means over all samples therefore change from seed to seed, and the table reports moments conditional on $n_{\pm} \geq 2$ (means) and $n_{\pm} \geq 3$ (standard deviations), with the exact conditional values. Conditional on $n_+ \geq 2$, the bias of $\hat{\eta}_{1,E}$ was +44.7% at $n = 8$ (exactly +46.2%), +26.8% at $n = 12$ and +14.4% at $n = 20$: a firm's mean up move is estimated without bias, its reciprocal is not. The median bias, which needs no conditioning, was smaller, +9.6% at $n = 8$. The corrected rate $(n_+ - 1) / \sum_{z_i > 0} z_i$ of Panel B removes the conditional bias (−0.6% at $n = 8$) at the price of a median bias of −21.9%.

Third, the rates are imprecise: conditional on $n_+ \geq 3$, the standard deviation of $\hat{\eta}_{1,E}$ was 1.04 times the truth at $n = 8$ (exactly 1.10), 0.72 at $n = 12$ and 0.45 at $n = 20$. Exact intervals from the gamma law covered the true rate in 89.6% to 90.7% of the samples at $n = 8$ and 12, and Wald intervals in 92.1% and 92.4%, so the stated uncertainty is accurate even where it is large.

Fourth, the standard deviation implied by the fitted law is nearly unbiased (relative bias −0.2% at $n = 8$), because it depends on the mean moves $1/\hat{\eta}_{1,E}$ and $1/\hat{\eta}_{2,E}$, which are sample means, rather than on the rates.

Fifth, diffusion noise matters for the scale and can be removed. In Panels C and D each move contains 2 trading days of diffusion at 2% per day, which adds 12% to the second moment. The mean squared move was then biased by +11.9% at $n = 8$ and +11.9% at $n = 20$, and the correction removed the bias (−0.3% and −0.3%). For the rates the noise partly offset the small-sample bias (conditional bias +20.7% at $n = 8$, against +44.7% without noise), and the correction, which rescales both rates, brought it back (+40.4%). The correction of Panel D matches the second moment of the fitted law to the mean squared move less the diffusion variance; the pricing estimator (45) matches the variance of the fitted law to s_z^2 and uses the diffusion of each move, a close analog.

Sixth, the shape is not identified for a single firm. When the truth was double exponential, the Akaike criterion picked the double exponential in 22.3% of the samples at $n = 8$, 31.7% at $n = 20$ and 43.8% at $n = 40$, and the Gaussian in 56.7% at $n = 8$; when the truth was Gaussian with the same standard deviation, it picked the Gaussian in 79.6% of the samples at $n = 8$ and the double exponential in 7.2% (6.4% at $n = 12$). The shares of picks across many firms are informative, but the samples of this table carry no diffusion, while the per-firm fits of Table 5, Panel C, are to raw moves. The row of that panel marked † therefore repeats the selection with each firm's own number of announcements and its own diffusion share (Appendix C.16).

Table 23: Monte Carlo behavior of the closed-form estimator of the double-exponential announcement law from n announcement returns, with diffusion contamination and AIC model selection.

n	relative bias (%)			sd/truth		90% coverage (%)			samples (%)	
	$\hat{\eta}_{1,E}$	$\hat{\eta}_{2,E}$	$\widehat{\text{sd}}(Z)$	$\hat{\eta}_{1,E}$	$\hat{\sigma}_J^2$	$\eta_{1,E}$ Wald	$\eta_{1,E}$ exact	p_E exact	$n_+ \leq 1$	$n_+ \leq 2$
<i>A. Announcement returns observed without diffusion noise: MLE</i>										
8	+44.7	+34.5	-0.2	1.04	0.81	92.1	89.6	97.6	6.4	21.3
12	+26.8	+21.4	-0.8	0.72	0.66	92.4	90.7	92.3	0.8	4.5
20	+14.4	+10.7	-0.4	0.45	0.50	91.7	90.4	96.3	0.0	0.1
40	+5.4	+5.3	+0.1	0.27	0.36	90.4	90.0	91.9	0.0	0.0
<i>B. As A, bias-corrected rates $\tilde{\eta} = (n_{\pm} - 1) / \sum z_i$ (undefined when $n_{\pm} \leq 1$)</i>										
8	-0.6	-0.6	+38.7	0.74	0.81	80.8	-	-	6.4	21.3
12	-0.4	+1.0	+21.0	0.56	0.66	85.1	-	-	0.8	4.5
20	+0.5	+0.2	+11.0	0.38	0.50	88.0	-	-	0.0	0.1
40	-0.7	+0.4	+5.4	0.25	0.36	88.6	-	-	0.0	0.0
<i>C. Observed move = Z + diffusion noise, $\sigma_d \sqrt{k} = 2.83\%$: naive MLE</i>										
8	+20.7	+20.2	+8.8	0.79	0.85	89.9	90.4	97.6	6.0	21.7
12	+10.8	+8.4	+8.7	0.63	0.68	90.2	90.3	92.9	0.8	3.7
20	+0.6	+0.5	+9.4	0.38	0.53	89.1	89.9	95.8	0.0	0.1
40	-6.4	-3.7	+9.8	0.22	0.37	85.9	87.2	92.6	0.0	0.0
<i>D. As C, second-moment correction: $\tilde{\eta}_j = \hat{\eta}_j \sqrt{m_2 / (m_2 - k\sigma_d^2)}$, $\tilde{\sigma}_J^2 = m_2 - k\sigma_d^2$</i>										
8	+40.4	+40.9	+0.5	1.09	0.85	90.9	-	-	6.0	21.7
12	+23.5	+21.7	+1.1	0.77	0.68	90.8	-	-	0.8	3.7
20	+9.1	+9.4	+2.4	0.44	0.53	90.8	-	-	0.0	0.1
40	+0.2	+3.2	+3.3	0.25	0.37	88.9	-	-	0.0	0.0
n	truth DE: AIC picks (%)			truth Gaussian: AIC picks (%)			$\hat{\sigma}_J^2$ bias (%)			
	DE	Gaussian	t_4	DE	Gaussian	t_4	naive (C)	corrected (D)		
8	22.3	56.7	20.9	7.2	79.6	13.3	+11.9	-0.3		
12	25.0	48.0	27.0	6.4	79.6	14.0	+10.8	-1.4		
20	31.7	34.7	33.6	5.1	80.8	14.1	+11.9	-0.3		
40	43.8	19.0	37.2	3.4	85.2	11.3	+12.2	0.0		

5,000 replications per row. Truth $Z \sim \text{DE}(p_E, \eta_{1,E}, \eta_{2,E}) = \text{DE}(0.45, 20, 16)$: mean up move $1/\eta_{1,E} = 5.0\%$, mean down move $1/\eta_{2,E} = 6.25\%$, $\text{sd}(Z) = 8.00\%$, $\mathbb{E}Z^2 = 0.00655$. Closed-form MLE $\hat{p}_E = n_+/n$, $\hat{\eta}_{1,E} = n_+ / \sum_{z_i > 0} z_i$, $\hat{\eta}_{2,E} = n_- / \sum_{z_i < 0} |z_i|$; $\widehat{\text{sd}}(Z)$ is the standard deviation implied by the fitted law and $\hat{\sigma}_J^2 = n^{-1} \sum z_i^2$ the mean squared announcement return (the DJ-hist estimator), whose relative bias is zero in A and B by construction. The rate estimators have no finite mean, because $n_{\pm} = 1$ has positive probability at every n and $\mathbb{E}[1 / \sum z_i^+ | n_+ = 1] = \infty$, and no finite variance, because $n_{\pm} = 2$ has positive probability; Monte Carlo means and standard deviations over all samples therefore change from seed to seed. The relative bias of a rate is the bias of its mean over the samples with $n_{\pm} \geq 2$, and its sd/truth is computed over the samples with $n_{\pm} \geq 3$ (n_+ for $\eta_{1,E}$, n_- for $\eta_{2,E}$; the last two columns give the shares of samples left out for $\eta_{1,E}$; $n_- \leq 1$ occurs in 1.7% of the samples at $n = 8$). Conditional on n_+ the relative bias of $\hat{\eta}_{1,E}$ is exactly $1/(n_+ - 1)$, and with the binomial law of n_+ the exact values for A at $n = 8, 12, 20, 40$ are a bias of +46.2, +27.7, +13.8, +6.1% and an sd/truth of 1.10, 0.81, 0.48, 0.27 (bias-corrected rate: bias 0, sd/truth 0.77, 0.61, 0.41, 0.26); at $n = 8$ the Monte Carlo sd is itself imprecise, because the fourth moment of the rate is finite only for $n_+ \geq 5$. The median relative bias of $\hat{\eta}_{1,E}$ over all samples, which needs no conditioning, is +9.6, +8.7, +5.3, +1.6% in A and -3.0, -4.6, -6.4, -10.0% in C. Wald intervals are $\hat{\eta} \pm 1.645 \hat{\eta} / \sqrt{n_{\pm}}$; exact intervals use $\eta_{1,E} \sum z_i^+ \sim \Gamma(n_+, 1)$ and, for p_E , Clopper-Pearson. In C each observed move is Z plus $\mathcal{N}(0, k\sigma_d^2)$ with $k = 2$ trading days and $\sigma_d = 2\%$, which adds 12% to the second moment; in D the correction subtracts $k\sigma_d^2$ from $m_2 = n^{-1} \sum x_i^2$ and rescales both rates ($m_2 - k\sigma_d^2$ is floored at $m_2/4$, which binds in at most 1.2% of the samples). AIC = $2k - 2\ell$ with $k = 3$ (DE), 2 (Gaussian), 2 (Student t with 4 degrees of freedom, location and scale); the Gaussian truth has the same standard deviation as Z . Written by `code/scripts/estimate_mc.py`; numbers in `output/estimation_mc.json`.




Figure 8: Accuracy of the double-exponential announcement-law estimators against the number n of announcement returns (the law of Table 23, without diffusion noise). (a) Relative bias: means of the rate estimators conditional on $n_{\pm} \geq 2$ (markers) with the exact conditional values (lines), the median of $\hat{\eta}_{1,E}$, the bias-corrected rate $\hat{\eta}_{1,E} = (n_+ - 1) / \sum_{z_i > 0} z_i$, the standard deviation implied by the fitted law and the mean squared return $\hat{\sigma}_J^2$. (b) Coverage of nominal 90% intervals: Wald and exact gamma intervals for the rates, Wald and Clopper-Pearson intervals for p_E .

C.4 The pricing estimator

Let a double-exponential law have up-move probability p and rates $\eta_1 = 1/(c\mu_+)$ and $\eta_2 = 1/(c\mu_-)$ for mean moves $\mu_{\pm} > 0$ and a scale $c > 0$. By (43) and (59),

$$\begin{aligned} \mathbb{E}Z &= c(p\mu_+ - (1-p)\mu_-), & \mathbb{E}Z^2 &= 2c^2(p\mu_+^2 + (1-p)\mu_-^2), \\ \text{Var } Z &= c^2(2p\mu_+^2 + 2(1-p)\mu_-^2 - (p\mu_+ - (1-p)\mu_-)^2). \end{aligned} \quad (64)$$

With (p, μ_+, μ_-) the sign split of the raw moves, $c = 1$ is the maximum-likelihood law of (44), and setting $\text{Var } Z = s_z^2$ gives the scale of (45): the estimator rescales both rates of the maximum-likelihood law by a common factor. Under (41), $\mathbb{E}[z_i^2] = \mathbb{E}[Z_i^2]$, so s_z^2 is an unbiased estimator of the announcement second moment for a firm whose announcement law does not change; matching the variance of the fitted law to it overstates the variance by $(\mathbb{E}Z)^2$, which is negligible when the mean announcement return is small relative to its standard deviation, as in Table 5. The floor $\eta_1 \geq 1.5$ keeps $\mathbb{E}[e^Z]$ finite, as (58) requires, and the floor on η_2 is a numerical guard. The second-moment correction of Table 23, Panel D, multiplies both maximum-likelihood rates by $(m_2/(m_2 - k\sigma_d^2))^{1/2}$, with m_2 the mean squared observed move and $k\sigma_d^2$ the diffusion variance of the design, which multiplies the second moment of the fitted law by $(m_2 - k\sigma_d^2)/m_2$. On the 7,079 events of the main study that Kou-SEJ-P priced, the fitted laws had a median $\hat{p}_E$ of 0.50, median rates of 26.9 and 26.6 and a median announcement standard deviation of 5.50%; 4 of them had a rate above 1,000, the near-degenerate case of Appendix C.7.

C.5 The double-exponential law with diffusion

Let $X = Z + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, v)$ independent of Z . Completing the square in $\int_0^\infty \eta_1 e^{-\eta_1 y} \varphi_v(x - y) dy$, with φ_v the $\mathcal{N}(0, v)$ density, and in its mirror image gives the density and the distribution

function

$$\begin{aligned} f_X(x) &= p \eta_1 e^{\eta_1^2 v/2 - \eta_1 x} \Phi\left(\frac{x - \eta_1 v}{\sqrt{v}}\right) + (1 - p) \eta_2 e^{\eta_2^2 v/2 + \eta_2 x} \Phi\left(\frac{-x - \eta_2 v}{\sqrt{v}}\right), \\ F_X(x) &= \Phi\left(\frac{x}{\sqrt{v}}\right) - p e^{\eta_1^2 v/2 - \eta_1 x} \Phi\left(\frac{x - \eta_1 v}{\sqrt{v}}\right) + (1 - p) e^{\eta_2^2 v/2 + \eta_2 x} \Phi\left(\frac{-x - \eta_2 v}{\sqrt{v}}\right), \end{aligned} \quad (65)$$

which reduce to the double-exponential density and distribution function as $v \rightarrow 0$. Equation (65) serves three purposes. First, it is the exact likelihood $\sum_i \ln f_X(x_i; p, \eta_1, \eta_2, v_i)$ of the moves (41), with each move’s own diffusion variance; maximized over the parameters, it gives the per-firm law “with diffusion” of Tables 5 and 29. Without bounds this likelihood has a degenerate ridge on samples that the diffusion alone explains ($p \rightarrow 1$ and $\eta_2 \rightarrow 0$: a down jump of vanishing probability and unbounded size), so the fit is confined to the region stated in the notes of Table 5 and computed by the adaptive simplex method (Nelder and Mead, 1965; Lagarias et al., 1998) from the pricing-estimator start; fits that end on the boundary of the region are flagged and left out of the medians. Second, it is the predictive density of the Kou-SEJ-P announcement law, convolved with the diffusion of the event being predicted. Third, its distribution function gives the probability integral transform behind the coverage columns of Table 24. The Akaike comparisons use $\text{AIC} = 2k - 2\ell$ with $k = 3$ for the double exponential (closed-form estimator), $k = 2$ for the Gaussian and $k = 2$ for the Student t_4 (location and scale).

C.6 Predictive laws

For a target event i , the prior set P_i holds the $n_i \geq 4$ earlier announcements of the same firm with a finite move whose post-event snapshot is not later than the pre-event snapshot of i ; an announcement whose spot fails the spot-consistency screen of Section 7.3 has no finite move, and the firm-level data screen of Section 6.1 flags none of the other moves. Let $s_i^2 = n_i^{-1} \sum_{j \in P_i} x_j^2$ be the mean squared prior move (the DJ-hist variance), $s_{z,i}^2 = n_i^{-1} \sum_{j \in P_i} z_j^2$ the corrected one, where a prior event without σ_{HV} takes the target’s σ_{HV} (known at d_{pre}), and $v_i = \sigma_{\text{HV},i}^2 d_i / 252$ the target’s own diffusion variance. The laws of Table 24 are, with location zero throughout:

- Gaussian: $\mathcal{N}(0, s_i^2)$. Student t_4 : four degrees of freedom and variance s_i^2 (scale $s_i/\sqrt{2}$). Laplace: variance s_i^2 (scale $s_i/\sqrt{2}$). Sign-split DE: the firm’s sign split rescaled by (64) to variance s_i^2 .
- Gaussian, corrected: $\mathcal{N}(0, s_{z,i}^2 + v_i)$. DE-P: the Kou-SEJ-P law (45) convolved with $\mathcal{N}(0, v_i)$ by (65); both require $s_{z,i}^2 > 0$.
- Pooled shapes. For the calendar quarter q of the target, the standardized moves $u_j = x_j/\sigma_{j,j}^{\text{hist}}$ of all matched events with at least 4 prior announcements whose post-event snapshot precedes the start of q (at least 250 of them) give a sign split $(\hat{p}_q, \hat{\mu}_{+,q}, \hat{\mu}_{-,q})$ and the maximum-likelihood degree of freedom $\hat{\nu}_q$ of a Student t law with location zero and free scale. $t_{\hat{\nu}}$: degree of freedom $\hat{\nu}_q$ and variance s_i^2 . DE, pooled shape: the pooled sign split rescaled to variance s_i^2 . DE-P, pooled shape: rescaled to $s_{z,i}^2$ and convolved with $\mathcal{N}(0, v_i)$. This is the rule by which Kou-SEJ-Q takes a pooled shape (Section 6.6).
- Option-implied: with $\hat{\sigma}_J^{2,\mathbb{Q}} > 0$ and $\hat{\sigma}_{\text{bg}}^2$ from (47) at the target’s d_{pre} , the Gaussian $\mathcal{N}(0, \hat{\sigma}_J^{2,\mathbb{Q}} + \hat{\sigma}_{\text{bg}}^2 d_i / 252)$, and the firm’s sign split rescaled to variance $\hat{\sigma}_J^{2,\mathbb{Q}}$ and convolved with $\mathcal{N}(0, \hat{\sigma}_{\text{bg}}^2 d_i / 252)$.

Each group of laws is compared on the events where all of its laws are defined, which is why the counts of Table 24 differ across groups. The score and the differences are (46); the coverage is the share of events with $\hat{F}_{L,i}(x_i) \in [0.05, 0.95]$. A sign-split rate above 1,000, a mean jump below a tenth of a percent, marks the near-degenerate fits discussed in Appendix C.7.

C.7 Out-of-sample prediction: the scores (exploratory)

Table 24 reports, for the laws defined above, LS_L of (46), the mean of $\Delta_{L,i}$ with a 90% event-bootstrap interval, its median, the share of events with $\Delta_{L,i} > 0$ and the coverage of the central 90% predictive interval. Like Section 6.4, this is a post-freeze exploratory analysis: the pre-registration names no predictive comparison of announcement laws, the comparison uses the holdout, and its results were computed after the specification freeze; no priced model depends on it.

Fat tails helped. The Student t_4 law improved on the Gaussian by +0.070 nats per event in the development sample and +0.059 in the holdout, and the Laplace law by +0.067 and +0.055, with intervals that exclude zero. The sign-split double exponential with the firm's own shape did worse than the Gaussian: its mean difference was -0.389 and -0.886 , its median difference -0.021 and -0.108 . The means are dominated by a few near-degenerate fits. When a firm's history has a single small move on one side, the fitted mean jump on that side is close to zero: in the holdout, 4 events had a rate above 1,000 (the largest 3.4×10^5), the worst of them contributed -0.740 to the mean difference, and without these events the mean difference was -0.136 . Even then the firm's own asymmetry cost more than the exponential tails gained, as the homogeneity of the up share in Section 6.1 leads one to expect.

The announcement law of Kou-SEJ-P scored -0.102 and -0.096 against the Gaussian, and -0.041 $[-0.087, +0.004]$ and -0.007 $[-0.041, +0.030]$ against the Gaussian with the same corrected variance: with the firm's own shape, the double exponential added nothing distinguishable from zero to a Gaussian. Both corrected laws lost to the uncorrected Gaussian on average (mean differences of -0.061 and -0.089 for the corrected Gaussian), although the corrected Gaussian scored higher for most events (69.7% and 60.1%). The correction shrinks the predictive law of firms with small past moves, sometimes to a fraction of its uncorrected width (Appendix C.17), and the log score punishes a law that is too narrow for the event at hand far more than one that is too wide; the 90% intervals of the Kou-SEJ-P law covered only 82.7% and 78.9% of the realized moves. With a pooled shape the same law improved on its firm-shape version by +0.128 $[+0.098, +0.156]$ and +0.121 $[+0.109, +0.134]$ nats and on the corrected Gaussian by +0.096 and +0.114. The pooled Student degree of freedom ranged from 3.1 to 4.6 over the 17 quarters from 2021Q2. The double-exponential shape helps when its asymmetry and tail ratio are estimated from the cross-section rather than from one firm. The pricing study keeps the frozen firm-shape estimator: these comparisons include the holdout, and changing the pricing law in response to them would tune the model on the holdout.

The option-implied Gaussian improved on the historical Gaussian by +0.080 nats in the development sample and +0.038 in the holdout, where the interval $[-0.009, +0.075]$ contains zero, and on the corrected Gaussian by +0.139 and +0.129, with intervals that exclude zero; its 90% intervals covered 92.1% and 89.9% of the realized moves, in the holdout the closest to nominal of all laws. The option-implied variance exceeded the realized one on average (Section 6.7), yet it measured the size of the coming move better than the firm's corrected history, and a law that is somewhat too wide costs little log score. The option-implied double exponential with the firm's shape scored below the option-implied Gaussian (-0.091 and -0.014 against the historical Gaussian), as the firm-shape laws did among the historical ones.

C.8 Background jump parameters: the threshold rule and maximum likelihood

Section 6.5 asks whether the data could identify a background jump intensity and its shape. Two estimators answer it; the designs are in Appendix C.15.

The first is a threshold rule: a day is a jump when its absolute return exceeds the sample standard deviation; $\hat{\lambda}$ is the share of such days, $\hat{p}$, $\hat{\eta}_1$ and $\hat{\eta}_2$ come from the flagged returns, and $\hat{\sigma}$ from all returns (Table 25). On Gaussian returns without jumps, the rule flagged a

Table 24: Out-of-sample predictive comparison of announcement laws: log scores of the realized announcement move (post-freeze, exploratory).

Law	n	score	Δ vs Gaussian [90% CI]	median Δ	better (%)	coverage (%)
<i>A. Development 2020–2021</i>						
<i>laws fitted on the firm’s prior moves</i>						
Gaussian	2,622	1.293	0	0	–	90.6
Student t_4	2,622	1.363	+0.070 [+0.046, +0.100]	+0.118	61.2	88.1
Laplace	2,622	1.360	+0.067 [+0.043, +0.094]	+0.011	50.7	90.2
DE, sign split	2,622	0.904	–0.389 [–0.638, –0.213]	–0.021	48.6	83.8
Gaussian, corrected	2,622	1.232	–0.061 [–0.084, –0.039]	+0.023	69.7	87.2
DE-P (Kou-SEJ-P law)	2,622	1.191	–0.102 [–0.143, –0.062]	+0.120	58.1	82.7
<i>pooled shape, firm scale</i>						
$t_{\hat{\nu}}$, pooled ν	2,230	1.379	+0.069 [+0.038, +0.108]	+0.143	60.1	87.4
DE, pooled shape	2,230	1.379	+0.069 [+0.040, +0.102]	+0.020	51.0	90.7
DE-P, pooled shape	2,230	1.347	+0.037 [+0.009, +0.069]	+0.095	57.6	87.1
<i>option-implied (risk-neutral reference)</i>						
Option-implied Gaussian	2,407	1.347	+0.080 [+0.044, +0.120]	+0.010	52.6	92.1
Option-implied DE	2,407	1.176	–0.091 [–0.240, +0.012]	+0.041	53.6	86.5
<i>B. Holdout 2022–2025</i>						
<i>laws fitted on the firm’s prior moves</i>						
Gaussian	11,553	1.061	0	0	–	84.8
Student t_4	11,553	1.120	+0.059 [+0.041, +0.080]	+0.017	51.5	81.7
Laplace	11,553	1.116	+0.055 [+0.038, +0.073]	–0.095	42.0	84.4
DE, sign split	11,553	0.175	–0.886 [–2.363, –0.127]	–0.108	41.7	79.9
Gaussian, corrected	11,553	0.972	–0.089 [–0.119, –0.063]	+0.007	60.1	82.8
DE-P (Kou-SEJ-P law)	11,553	0.965	–0.096 [–0.115, –0.077]	–0.016	48.4	78.9
<i>pooled shape, firm scale</i>						
$t_{\hat{\nu}}$, pooled ν	11,553	1.117	+0.056 [+0.038, +0.076]	+0.016	51.3	81.5
DE, pooled shape	11,553	1.116	+0.055 [+0.038, +0.074]	–0.100	41.8	84.4
DE-P, pooled shape	11,553	1.086	+0.025 [+0.010, +0.041]	–0.029	46.6	82.1
<i>option-implied (risk-neutral reference)</i>						
Option-implied Gaussian	10,875	1.090	+0.038 [–0.009, +0.075]	–0.019	44.8	89.9
Option-implied DE	10,875	1.039	–0.014 [–0.038, +0.013]	–0.067	44.1	85.4

Post-freeze, exploratory: this comparison was not pre-registered, includes the holdout and was computed after the specification freeze; the pricing models do not use it. Score: mean log predictive density of the realized announcement move (nats per event) for primary-sample events with at least 4 prior announcements of the same firm; every law uses information available at the pre-event snapshot only. Δ : per-event log score minus that of the Gaussian law on the same events (positive: better than Gaussian), mean with the 5th and 95th percentiles of 1,000 event-bootstrap draws, median, and the share of events with a positive difference. Coverage: share of realized moves inside the central 90% predictive interval (nominal 90). Firm-history laws: s^2 is the mean squared prior move (the DJ-hist variance), shared by the Gaussian, Student t_4 , Laplace and sign-split DE laws; “corrected” laws use the diffusion-corrected prior second moment $\bar{z}^2$ for the announcement and add the event’s own diffusion $\sigma_{HV}^2 d/252$; DE-P is the announcement law of Kou-SEJ-P (pricing estimator, Table 5). Pooled shape: the degree of freedom $\hat{\nu}$, or $\hat{p}_E$ and the ratio of mean up to mean down move, estimated from the moves of events settled before the quarter, each standardized by its own historical σ_J (at least 250 moves). Option-implied: σ_J^2 from the two-expiry decomposition at the pre-event snapshot (events with $\sigma_J^2 > 0$) and the decomposition’s background variance for the diffusion. The mean difference of the sign-split DE law is not a typical loss: it is pulled down by a few near-degenerate fits, in which the prior moves on one side are few and small, so that the fitted mean jump on that side is close to zero: in A, 11 events have a sign-split rate above 1,000 (11 of them with a single prior move on that side; largest rate 1.3×10^4), the worst event (ASTE, 2021-08-04) scores 338 nats below the Gaussian and contributes -0.129 to the mean difference, and without the 11 events the mean difference is -0.218 against a median of -0.021 ; in B, 4 events have a sign-split rate above 1,000 (4 of them with a single prior move on that side; largest rate 3.4×10^5), the worst event (TNDM, 2023-08-03) scores 8,546 nats below the Gaussian and contributes -0.740 to the mean difference, and without the 4 events the mean difference is -0.136 against a median of -0.108 . Written by `code/scripts/estimate_firms.py`; numbers in `output/estimation_firms.json`.

Table 25: Monte Carlo behavior of the one-standard-deviation threshold rule for the Kou parameters.

n	$\hat{\lambda}$ (per day)	$\hat{\eta}_1$	$\hat{\eta}_2$	$\hat{\sigma}$	variance ratio
<i>(i) Gaussian returns, no jumps: $\sigma = 0.3175$ (2% per day), $\lambda = 0$;</i>					
<i>large-sample limits $\hat{\lambda} \rightarrow 0.317$, $\hat{\eta} \rightarrow 32.8$, ratio $\rightarrow 2.48$</i>					
24	0.340 [0.250, 0.458]	35.1 [24.4, 48.2]	34.9 [24.2, 47.9]	0.307 [0.233, 0.386]	2.49 [1.41, 3.83]
80	0.324 [0.263, 0.375]	33.4 [27.8, 39.8]	33.4 [27.9, 39.6]	0.314 [0.273, 0.356]	2.48 [1.86, 3.16]
250	0.320 [0.288, 0.352]	33.0 [29.8, 36.5]	33.0 [29.6, 36.5]	0.316 [0.293, 0.341]	2.48 [2.12, 2.88]
2520	0.318 [0.307, 0.328]	32.8 [31.7, 33.8]	32.8 [31.7, 33.9]	0.317 [0.310, 0.325]	2.48 [2.36, 2.60]
<i>(ii) Kou returns: $\sigma = 0.30$, $\lambda = 0.3$ per day (= 75.6 per year), $p = 0.4$, $\eta_1 = 25$, $\eta_2 = 20$; total volatility 0.6431</i>					
24	0.213 [0.083, 0.375]	19.4 [6.9, 34.9]	16.4 [5.3, 34.0]	0.591 [0.315, 0.991]	2.50 [0.61, 6.47]
80	0.170 [0.100, 0.263]	15.6 [8.2, 24.6]	12.7 [6.6, 21.6]	0.625 [0.434, 0.858]	2.34 [1.07, 4.29]
250	0.155 [0.112, 0.200]	14.3 [9.9, 19.3]	11.4 [8.1, 15.2]	0.640 [0.526, 0.776]	2.30 [1.53, 3.36]
2520	0.149 [0.136, 0.163]	13.8 [12.4, 15.3]	11.0 [10.0, 12.2]	0.643 [0.604, 0.683]	2.25 [1.99, 2.53]
<i>(iii) 72 Gaussian days (2%) plus 8 announcement days with a DE(0.5, 23.6, 23.6) jump (sd 0.06); no Poisson jumps</i>					
80	0.229 [0.113, 0.338]	23.2 [12.2, 32.4]	23.2 [11.3, 32.8]	0.430 [0.328, 0.579]	2.28 [1.38, 3.80]

Each cell is the Monte Carlo mean followed by the 5th and 95th percentiles over 2,000 replications. The rule flags a day as a jump when $|\varrho_i|$ exceeds the sample standard deviation; $\hat{\lambda}$ is the share of flagged days, $\hat{\eta}_1$ and $\hat{\eta}_2$ the reciprocals of the mean flagged positive and negative returns, $\hat{\sigma}$ the standard deviation of all returns times $\sqrt{252}$, and the variance ratio is the total daily variance the pricer is handed, $\hat{\sigma}^2/252 + \hat{\lambda}(2\hat{p}/\hat{\eta}_1^2 + 2(1-\hat{p})/\hat{\eta}_2^2)$, divided by the true mean daily variance of the design. In (i) the truth is $\lambda = 0$ and no η is defined; the limits are $\hat{\lambda} \rightarrow 2(1 - \Phi(1))$ and $\hat{\eta} \rightarrow 1/(\sigma_d \mathbb{E}[Z | Z > 1])$ with $\mathbb{E}[Z | Z > 1] = 1.525$. In (iii) the eight announcement days make up 29% of the flagged days on average. Written by `code/scripts/estimate_mc.py`; numbers in `output/estimation_mc.json`.

share of 0.318 of the days at $n = 2520$, the large-sample limit $2(1 - \Phi(1)) = 0.317$, one “jump” every 3.15 days, with a mean flagged size of 1.525 daily standard deviations whatever the data (Appendix C.14). Since $\hat{\sigma}$ already contains all the variance, the jump component is counted twice, and the pricer is handed 2.48 times the true variance. On Kou returns with $\lambda = 0.3$ per day the rule reported $\hat{\lambda} = 0.149$ and $\hat{\sigma} = 0.643$, which is the total volatility 0.643 and not the diffusion coefficient 0.30; over a grid of true intensities the reported intensity was not even monotone in the truth (Figure 10). With 8 announcement days among 72 ordinary days and no Poisson jumps, it reported $\hat{\lambda} = 0.229$ per day, and only 29% of the flagged days were announcement days. A threshold fixed at a multiple of the standard deviation measures the threshold, not the jumps; a threshold that separates jumps from the diffusion must shrink with the sampling interval (Mancini, 2009), and the tests of Aït-Sahalia (2004), Barndorff-Nielsen and Shephard (2006), Lee and Mykland (2008) and Aït-Sahalia and Jacod (2009) likewise rely on high-frequency sampling.

The second is maximum likelihood, the approach of Ramezani and Zeng (2007). The density of a daily Kou return is the FFT inversion of the characteristic function (9) without the announcement factor on a fixed grid of 16,384 points, and the likelihood is maximized over $(\sigma, \lambda, p, \eta_1, \eta_2)$ and a drift from four starting points (Table 26). With ten years of daily returns, $n = 2520$, the estimator worked: the median $\hat{\lambda}$ was 76.9 per year against a truth of 75.6, with a 5–95% range of 66.9–92.5. At the sample sizes of a trailing window it did not. At $n = 250$ the range was 46.8–179.1, and at $n = 80$ 47.5–1553.9, with a Monte Carlo standard deviation of 7.22 times the truth; $\hat{\sigma}$ sat at its lower bound in 12.0% of the samples of 80 returns, and the Hessian was invertible in only 84.5%. Panel B fits the Kou model to returns without jumps. The median $\hat{\lambda}$ was the upper bound of the search region, 2520 per year, at both sample sizes, and $\hat{\lambda}$ lay on that bound in 78% of the samples of 250 returns and 84% of those of 2,520. The total volatility was nevertheless recovered (median ratio to the truth 1.0054 and 0.9988), and the likelihood-ratio

Table 26: Monte Carlo behavior of the characteristic-function maximum-likelihood estimator of the Kou daily-return law, and its ridge behavior on jump-free data.

		$n = 80$		$n = 250$		$n = 1000$		$n = 2520$	
quantity	truth	median	5–95%	median	5–95%	median	5–95%	median	5–95%
<i>A. Kou truth</i>									
σ	0.300	0.260	0.020–0.356	0.293	0.243–0.336	0.298	0.278–0.318	0.301	0.286–0.313
λ (per year)	75.6	135	47–1554	82	47–179	75	58–101	77	67–93
p	0.40	0.51	0.08–0.97	0.43	0.16–0.72	0.40	0.27–0.55	0.41	0.33–0.48
η_1	25.0	35.5	12.1–193.3	26.2	13.3–62.1	25.2	18.9–35.6	25.7	22.2–31.0
η_2	20.0	22.4	10.6–88.6	20.7	12.3–32.9	19.9	16.5–25.7	20.2	17.6–22.2
jump share of variance	0.78	0.86	0.66–1.00	0.79	0.67–0.88	0.79	0.72–0.82	0.78	0.76–0.81
total volatility	0.643	0.626	0.423–0.898	0.641	0.537–0.767	0.643	0.590–0.706	0.643	0.610–0.676
median SE / truth of $\hat{\lambda}$		1.01		0.42		0.18		0.12	
MC sd / truth of $\hat{\lambda}$		7.22		1.18		0.19		0.11	
Hessian invertible (%)		84.5		99.5		100.0		100.0	
$\hat{\sigma}$ at its lower bound (%)		12.0		0.5		0.0		0.0	
$\hat{\lambda}$ at its bound (%)		3.0		0.0		0.0		0.0	
replications		200		200		200		50	
<i>B. Jump-free truth: Gaussian returns, $\sigma = 0.317$, $\lambda = 0$</i>									
$\hat{\lambda}$ (per year)	0			2520	144–2520			2520	37–2520
$\hat{\sigma}$	0.317			0.251	0.020–0.322			0.296	0.266–0.322
jump share of variance	0			0.39	0.04–1.00			0.12	0.02–0.31
total volatility / truth	1			1.005	0.924–1.068			0.999	0.977–1.022
$2(\ell_{\text{Kou}} - \ell_{\text{Gauss}})$	0			0.86	−0.01–5.87			0.41	−0.15–4.26
$\hat{\lambda}$ at its bound (%)				78				84	
$\hat{\lambda} > 252$ (%)				92				88	
replications				100				50	

Median and 5th–95th percentile range of the estimates across replications (n daily returns per sample; panel B is run for $n = 250$ and $n = 2520$ only). A: daily log returns simulated from the Kou law with $\sigma = 0.30$, $\lambda = 75.6$ per year ($= 0.3$ per day), $p = 0.4$, $\eta_1 = 25$, $\eta_2 = 20$ and a drift that sets $\mathbb{E}g = 0$. The likelihood uses the density of the daily return obtained by FFT inversion of the characteristic function (9) (without the announcement factor, drift free) on a fixed grid of 2^{14} points with spacing 3.8×10^{-4} , the log density interpolated linearly; parameters ($\ln \sigma, \ln \lambda, \log p, \ln(\eta_1 - 1), \ln \eta_2, \mu_d$); adaptive Nelder–Mead followed by BFGS from four starts, best end point kept. “Median SE/truth” is the median asymptotic standard error (inverse numerical Hessian) divided by the true value; “MC sd/truth” the Monte Carlo standard deviation of the estimates divided by the true value. The jump share of variance is $\lambda \mathbb{E}[Y^2] / (\sigma^2 + \lambda \mathbb{E}[Y^2])$. The search is confined to $\sigma \geq 0.02$, $\lambda \leq 2520$ per year and $\eta \leq 5000$; “at its bound” means within 1% of the bound. The Hessian is not computed when an estimate lies on a bound of the search region (an end point of the difference stencil would leave it); of the samples without an invertible Hessian, 30 of the 31 at $n = 80$ (24 with $\hat{\sigma}$ at its lower bound) and the one at $n = 250$ ($\hat{\sigma}$ at its lower bound) have an estimate on a bound. Profile likelihood on jump-free samples ($n = 2520$, 10 samples): with λ fixed at 1, 5, 25, 75.6, 250, 750, 2520 per year and the other parameters re-optimized, the median of $-\ell(\lambda) + \ell_{\text{Gauss}}$ is +0.05, −0.01, −0.20, −0.40, −0.56, −0.64, −0.66; the median range of the profile over this grid is 0.47 log-likelihood units. Written by `code/scripts/estimate_mc.py`; numbers in `output/estimation_mc.json`.

statistic against the Gaussian had a median of 0.86 and 0.41. The profile likelihood in λ was flat, with a median range of 0.47 log-likelihood units over the grid of intensities in the table notes. In these experiments the data identify the total variance of ordinary returns; they do not identify its split into a diffusion and many small jumps. This flat ridge is a finding of the Monte Carlo study above. It is distinct from the problem that Honoré (1998) analyzes, that the likelihood of a jump-diffusion can be unbounded, so that unrestricted maximum likelihood is inconsistent. For the Kou law the likelihood diverges as $\sigma \rightarrow 0$ with the drift placed at one of the returns, since a return without a jump then has a density without bound, and the search region therefore bounds σ from below (table notes); $\hat{\sigma}$ ended on that bound in 12.0% of the samples of 80 returns.

Table 27: Two-expiry decomposition of the announcement variance under quote noise.

σ_J	spread	IV ₁	IV ₂	one quote per expiration (ATM call)			ATM call and put		
				bias (%)	SE (%)	$\hat{\sigma}_J^2 < 0$ (%)	bias $\hat{\sigma}_J$ (pts)	SE (%)	$\hat{\sigma}_J^2 < 0$ (%)
0.04	5%	0.3447	0.3231	−0.1	33.1	0.0	−0.06	22.4	0.0
	10%			−0.6	65.5	7.1	−0.31	44.4	0.9
	20%			+0.7	134.1	23.7	−0.44	89.1	13.6
0.06	5%	0.3934	0.3499	+0.2	18.4	0.0	−0.02	12.6	0.0
	10%			+0.8	36.4	0.0	−0.09	25.1	0.0
	20%			+0.6	73.7	9.9	−0.49	49.6	2.1
0.08	5%	0.4530	0.3842	+0.1	13.2	0.0	−0.01	9.1	0.0
	10%			+0.1	26.3	0.0	−0.07	17.8	0.0
	20%			+1.8	52.1	1.7	−0.29	35.7	0.0
0.10	5%	0.5196	0.4243	+0.1	10.7	0.0	−0.01	7.4	0.0
	10%			+0.5	21.7	0.0	−0.03	14.8	0.0
	20%			+0.7	43.1	0.0	−0.24	29.5	0.0

5,000 replications per row. Two at-the-money expirations $T_1 = 14/252$ and $T_2 = 28/252$, background $\sigma = 0.30$, $r = 0.04$, $q = 0$, $S_0 = K = 100$; true implied variances $IV_i^2 = \sigma^2 + \sigma_J^2/T_i$. Each observed mid is the true price times $(1 + \epsilon)$ with ϵ uniform on $\pm$ half the relative spread, independently across expirations; the implied volatility is inverted from the noisy mid and $\hat{\sigma}_J^2 = T_1 IV_1^2 - T_1(T_2 IV_2^2 - T_1 IV_1^2)/(T_2 - T_1)$, which for $T_2 = 2T_1$ equals $2T_1(IV_1^2 - IV_2^2)$. Bias and standard error (Monte Carlo standard deviation) of $\hat{\sigma}_J^2$ are in percent of the true σ_J^2 ; the bias of $\hat{\sigma}_J = \sqrt{\max(\hat{\sigma}_J^2, 0)}$ is in volatility points (percent per $\sqrt{\text{event}}$). “2 quotes” averages the implied volatilities of the ATM call and the ATM put, each with its own mid error. Without noise the decomposition recovers σ_J^2 with a largest absolute error of 4.3×10^{-18} . Written by `code/scripts/estimate_mc.py`; numbers in `output/estimation_mc.json`.

C.9 The two-expiry decomposition under quote noise

Table 27 measures the effect of quote noise on (47), with two at-the-money expirations of 14 and 28 trading days, a background volatility of 0.30 and a mid error uniform on plus or minus half the relative spread. Without noise the decomposition recovered σ_J^2 to rounding error. With noise it remained unbiased (relative bias between −0.6% and 1.8%), but its standard error ranged from 11% of the truth for the largest announcement volatility and the tightest spread to 134% for the smallest announcement volatility and the widest spread, where 23.7% of the estimates were negative. The delta method of Appendix C.13 explains the pattern: the standard error is proportional to the spread and to $T_1 T_2 / (T_2 - T_1)$ and, relative to σ_J^2 , roughly inversely proportional to the announcement variance; it reproduced every Monte Carlo standard error of the table to within 1.3% with one quote per expiration and 1.4% with two. Averaging the call and the put about halves the variance of the quote error (33% against 22% in the first row). The option-implied announcement variance of a single event is therefore noisy for firms with small announcements and wide spreads, and precise for firms with large announcements and tight spreads; averages over many events, as in Section 6.7, are unbiased either way.

C.10 The earnings variance premium in detail

Figure 9 shows the premium ratio (49) of Table 6 by VIX tercile and calendar year (panel b) and the two announcement volatilities by decile of the implied variance (panel a). Across deciles the realized variance was above the implied one in the lowest decile shown and well below it in the highest. Part of this pattern is mechanical: sorting on a noisy estimate exaggerates its spread across the sorted groups, and the decomposition is noisy for small announcements (Table 27). By calendar year the ratio was largest in 2020 (1.88) and smallest in 2023 (1.05 [0.97, 1.13], an interval that contains one). In the holdout it fell with the firm’s past move, from 1.39 in the lowest tercile to 1.04 in the highest, and was 0.98 in the highest price tercile (Panels C and D). In Panel G of Table 6 the expected absolute move of the Kou-SEJ-P law is lower than the other




Figure 9: The earnings variance premium. (a) Option-implied against realized announcement volatility by decile of $\hat{\sigma}_j^{2,Q}$ (deciles of the pooled sample; within each decile and sample, the square roots of the mean implied variance and of the mean z^2 ; the lowest decile, whose mean implied variance is negative, is not shown on the log axes). (b) The ratio (49) with 90% event-bootstrap intervals by VIX tercile at d_{pre} , for the development and the holdout samples (Table 6, Panel B), and by calendar year (Panel E).

measures for two reasons: the corrected variance excludes the diffusion, and at a given standard deviation the symmetric double exponential expects an absolute move of $1/\sqrt{2}$ times it, against $\sqrt{2/\pi}$ times it for the Gaussian (Appendix C.1). The at-the-money straddle, at 7.97% of the spot, also prices the diffusion to expiration.

C.11 Clock and term structure of the premium ratio (exploratory)

Table 28 collects the checks of the ratio (49) that Section 6.7 summarizes. They are post-freeze exploratory analyses: none was pre-registered, and the priced models use the calendar-clock decomposition of Table 6.

Clock. The decomposition (47) can be written $\hat{\sigma}_j^{2,Q} = w_1 - (w_2 - w_1) T_1/(T_2 - T_1)$ with the total implied variances $w_j = IV_j^2 \text{DTE}_j/365$, which do not depend on the clock, so the clock enters only through the ratio $T_1/(T_2 - T_1)$. On the calendar clock $T_j = \text{DTE}_j/365$; on the trading clock T_j is the number of NYSE trading days from d_{pre} to expiration j over 252, the trading-year maturity of the pricing (Section 7.4). The life of the shorter expiration starts on a weekday and usually holds proportionally fewer weekend days than the interval between the two expirations, which is 28 calendar but 19 trading days for the median pair, so the trading clock raised $T_1/(T_2 - T_1)$ for 89.2% of the pairs; when $w_2 > w_1$ that assigns more of the shorter expiration's variance to the background. The trading-clock variance was the smaller one for 87.2% of the events. The denominator of (49) does not change: its diffusion term $\sigma_{\text{HV}}^2 d/252$ is on the trading clock already, and d is one trading day for 82.4% of the events of the primary sample. Which clock the market's implied variances follow is not observed in these data, so the two ratios bracket R rather than correct it.

Table 28: The premium ratio on two clocks and two placebo decompositions (post-freeze, exploratory).

	Development 2020–2021			Holdout 2022–2025			All		
	<i>n</i>	ratio	[90% CI]	<i>n</i>	ratio	[90% CI]	<i>n</i>	ratio	[90% CI]
<i>A. The ratio $\sigma_J^{2,Q}/\bar{z}^2$ on two clocks</i>									
calendar clock (as priced)	8,226	1.70	[1.60, 1.82]	13,009	1.14	[1.09, 1.19]	21,235	1.31	[1.27, 1.36]
week clusters (<i>n</i> : weeks)	99	1.70	[1.50, 1.94]	169	1.14	[1.08, 1.22]	268	1.31	[1.23, 1.42]
trading clock	8,226	1.52	[1.42, 1.63]	13,009	1.06	[1.01, 1.11]	21,235	1.20	[1.15, 1.24]
week clusters (<i>n</i> : weeks)	99	1.52	[1.34, 1.72]	169	1.06	[0.99, 1.13]	268	1.20	[1.12, 1.29]
<i>trading clock by VIX tercile</i>									
low	2,747	1.04	[0.92, 1.20]	5,347	1.06	[0.99, 1.13]	8,094	1.05	[0.99, 1.12]
middle	2,748	1.57	[1.43, 1.72]	5,107	1.04	[0.98, 1.13]	7,855	1.17	[1.10, 1.24]
high	2,731	1.92	[1.73, 2.13]	2,555	1.08	[0.97, 1.20]	5,286	1.44	[1.33, 1.56]
VIX top minus bottom, calendar	8,226	1.02	[0.77, 1.29]	13,009	0.07	[−0.06, 0.22]	21,235	0.48	[0.35, 0.62]
trading clock	8,226	0.88	[0.64, 1.13]	13,009	0.02	[−0.10, 0.17]	21,235	0.39	[0.27, 0.52]
$\sigma_J^{2,Q} < 0$ (%), cal. / trad.	8,226	9.4 / 12.8		13,009	6.1 / 7.8		21,235	7.4 / 9.7	
<i>B. Placebo 1: the same two expirations at the post-event snapshot</i>									
calendar clock	7,882	0.48	[0.43, 0.54]	12,140	0.37	[0.34, 0.40]	20,022	0.41	[0.38, 0.43]
trading clock	7,882	0.38	[0.33, 0.43]	12,140	0.32	[0.29, 0.35]	20,022	0.34	[0.31, 0.36]
positive (%), cal. / trad.	7,882	59.5 / 52.8		12,140	73.4 / 65.4		20,022	68.0 / 60.4	
<i>R</i> , same events, cal. / trad.	7,882	1.73 / 1.54		12,140	1.14 / 1.05		20,022	1.32 / 1.20	
<i>by VIX tercile, calendar clock</i>									
low	2,610	0.26	[0.20, 0.33]	4,986	0.33	[0.29, 0.37]	7,596	0.31	[0.28, 0.35]
middle	2,626	0.36	[0.28, 0.45]	4,722	0.36	[0.32, 0.40]	7,348	0.36	[0.32, 0.40]
high	2,646	0.77	[0.66, 0.89]	2,432	0.46	[0.40, 0.54]	5,078	0.60	[0.54, 0.66]
top minus bottom	7,882	0.51	[0.38, 0.64]	12,140	0.13	[0.05, 0.21]	20,022	0.28	[0.21, 0.36]
<i>by VIX tercile, trading clock</i>									
low	2,610	0.20	[0.14, 0.27]	4,986	0.29	[0.26, 0.34]	7,596	0.27	[0.24, 0.30]
middle	2,626	0.25	[0.16, 0.33]	4,722	0.31	[0.26, 0.35]	7,348	0.29	[0.26, 0.33]
high	2,646	0.63	[0.53, 0.75]	2,432	0.38	[0.32, 0.46]	5,078	0.49	[0.43, 0.56]
top minus bottom	7,882	0.43	[0.30, 0.56]	12,140	0.09	[0.01, 0.16]	20,022	0.22	[0.15, 0.29]
<i>C. Placebo 2: matched expirations 20 trading days later</i>									
calendar clock	6,568	0.41	[0.34, 0.49]	9,304	0.29	[0.27, 0.32]	15,872	0.33	[0.30, 0.36]
trading clock	6,568	0.30	[0.23, 0.38]	9,304	0.24	[0.21, 0.26]	15,872	0.26	[0.23, 0.29]
positive (%), cal. / trad.	6,568	53.8 / 44.1		9,304	63.0 / 51.5		15,872	59.2 / 48.4	
<i>R</i> , same events, cal. / trad.	6,568	1.69 / 1.50		9,304	1.14 / 1.04		15,872	1.32 / 1.19	

Post-freeze, exploratory analyses, not pre-registered; the priced models use the calendar-clock decomposition of Table 6. *R* is the ratio (49); cal. and trad. are the calendar and the trading clock; intervals are from the event bootstrap unless the row says week clusters. Trading clock: (47) with the same expirations and total implied variances $IV_j^2 DTE_j/365$ and the maturities T_j counted in NYSE trading days from the snapshot over 252; the denominator $\bar{z}^2$ is the same on both clocks. Event bootstrap as in Table 6; announcement-week clusters: 1,000 draws of multinomial weights over the announcement weeks (Monday-dated calendar week of the announcement date, the clusters of the pre-registered tests), each event weighted by its week (seed 20261014). Top minus bottom VIX tercile: joint event bootstrap over the sample. Placebos: (47) applied to two expirations that contain no announcement, with at-the-money implied volatilities between 0.05 and 3 as in the decomposition; the rows of Panels B and C divide the mean placebo variance by the mean of z^2 of the same events (the rows *R* give the ratio itself on those events), and under a flat background term structure both placebos have mean zero. Placebo 1 uses the event’s own two expirations at the post-event snapshot, after the announcement; it also contains the elevated volatility that follows an announcement. Placebo 2 uses the first option snapshot at least 20 trading days after the post-event snapshot (median lag 20, at most 23 trading days) and, among the expirations that end before the firm’s next announcement, the two closest in calendar maturity to those of the event’s pair (within 7 days each); events without such a pair, or without a next announcement in the calendar, are left out. Neither placebo measures the background slope before the announcement, so neither is subtracted from the ratio. Written by `code/scripts/estimate_firms.py`; numbers in `output/estimation_firms.json`.

Announcement weeks. Announcements of the same week share market news, and the event bootstrap of Appendix C.13 treats them as independent. The week-clustered intervals of Panel A resample the announcement weeks, the clusters of the pre-registered tests of Section 8 (99 weeks in the development sample and 169 in the holdout). They are wider than the event-bootstrap intervals, and on the trading clock the holdout interval, $[0.99, 1.13]$, contains one.

Placebos. The decomposition is exact only when the background implied variance is the same at both expirations; any difference is counted as announcement variance. The placebos apply (47) to pairs of expirations that contain no announcement, where under a flat background it would average zero. Placebo 1 takes the event’s own two expirations at the post-event snapshot (7,882 development and 12,140 holdout events with both implied volatilities); it also contains the elevated volatility that follows an announcement, which the shorter expiration prices before the announcement but the one-day z^2 does not. Placebo 2 takes, at the first option snapshot at least 20 trading days after the post-event snapshot, the two expirations closest in calendar maturity to those of the event’s pair (within 7 days each) among those that end before the firm’s next announcement (6,568 and 9,304 events). Relative to the mean z^2 of the same events, placebo 1 averaged 0.48 $[0.43, 0.54]$ and 0.37 $[0.34, 0.40]$ on the calendar clock and 0.38 and 0.32 on the trading clock, and placebo 2 0.41 and 0.29 (0.30 and 0.24); placebo 1 was positive for 73.4% of the holdout events. Neither placebo measures the background slope before the announcement: the first adds the post-announcement volatility, the second describes a later stretch of the firm’s options. I therefore do not subtract either from R . On the events of each placebo, $R - 1$ exceeded the placebo on both clocks in the development sample (for placebo 1, R was 1.73 on the calendar and 1.54 on the trading clock) and fell short of it in the holdout (1.14 and 1.05). In the development sample placebo 1 rose across the VIX terciles, from 0.26 to 0.36 and 0.77, a top-minus-bottom difference of 0.51 $[0.38, 0.64]$ against 1.02 $[0.77, 1.29]$ for R itself. The slope that the decomposition attributes to an announcement is therefore larger when the VIX is high even without one, and part of the rise of the premium with the VIX can come from that slope.

C.12 The pre-registered variant of Kou-SEJ-Q

Kou-SEJ-Q-prereg (Section 6.6) takes $V = \hat{\sigma}_J^{2,\mathbb{Q}}$ and $\hat{\sigma}_{\text{bg}}$ from (47) and a double-exponential law with mean zero, which leaves one free parameter:

$$\begin{aligned} \frac{p_E}{\eta_{1,E}} = \frac{1 - p_E}{\eta_{2,E}} = m, \quad V = \frac{2m^2}{p_E(1 - p_E)}, \\ \text{so that} \quad m = \sqrt{V p_E(1 - p_E)/2}, \quad \eta_{1,E} = \frac{p_E}{m}, \quad \eta_{2,E} = \frac{1 - p_E}{m}. \end{aligned} \tag{66}$$

A value $p_E > \frac{1}{2}$ makes up moves more frequent and smaller and down moves rarer and larger, a negative skew. The parameter p_E is chosen so that the model’s slope of implied volatility in standardized moneyness $k = \ln(K/F)/(\sigma_{\text{ATM}}\sqrt{T})$, fitted by least squares on k and k^2 (on k alone when there are few strikes) over the out-of-the-money kept quotes of the nearest expiration, equals the market’s. The range is $\max\{0.05, c/(1 + c)\} \leq p_E \leq 0.95$ with $c = 1.5^2 V/2$, which keeps $\eta_{1,E} \geq 1.5$. When no p_E in the range matches the market’s slope, the closest is used and the fit is flagged as clipped. That happened for 47% of the development events and 41% of the holdout events, and in 90% of the clipped fits the market’s skew was steeper than that of the closest member of the family: the pre-event skew contains more than the asymmetry of a zero-mean announcement jump on a flat diffusion, presumably the skew of the background itself. The median fitted p_E was 0.71, against a median historical share of up moves of 0.50 in the Kou-SEJ-P laws.

C.13 Standard errors and intervals

Maximum likelihood of the Kou law. The asymptotic standard errors of Table 26 are the square roots of the diagonal of the inverse numerical Hessian of the negative log-likelihood in the natural parameters; the Hessian is not computed when an estimate lies on the boundary of the search region. In small samples they understate the sampling variability badly: at $n = 80$ the median asymptotic standard error of $\hat{\lambda}$ was 1.01 times the truth, while the Monte Carlo standard deviation was 7.22 times the truth; at $n = 2520$ the two were 0.12 and 0.11.

Two-expiry decomposition. Write (47) as $\hat{\sigma}_j^{2,Q} = A(\text{IV}_1^2 - \text{IV}_2^2)$ with $A = T_1 T_2 / (T_2 - T_1)$. A relative error ϵ_j in the mid of expiration j moves its implied volatility by $\delta \text{IV}_j \approx (P_j / \mathcal{V}_j) \epsilon_j$, with P_j the option price and $\mathcal{V}_j$ its Black–Scholes vega; at the money $P_j / \mathcal{V}_j \approx \text{IV}_j$. With ϵ_1 and ϵ_2 independent and uniform on $[-s/2, s/2]$, so that $\text{Var } \epsilon_j = s^2/12$,

$$\text{SE}(\hat{\sigma}_j^{2,Q}) \approx 2A \frac{s}{\sqrt{12}} \left(\sum_{j=1}^2 \text{IV}_j^2 \frac{P_j^2}{\mathcal{V}_j^2} \right)^{1/2} \approx 2A \frac{s}{\sqrt{12}} (\text{IV}_1^4 + \text{IV}_2^4)^{1/2}. \quad (67)$$

When the implied volatilities of the call and the put are averaged, each with its own error, P_j^2 is replaced by $(C_j^2 + P_j^2)/4$, which roughly halves the variance. In the first row of Table 27, (67) gives 33.1% of the truth against 33.1% in the Monte Carlo study with one quote, and 22.5% against 22% with two; over the twelve cells of the table the largest relative discrepancy was 1.3% with one quote and 1.4% with two. The normal approximation $\mathbb{P}(\hat{\sigma}_j^{2,Q} < 0) \approx \Phi(-\sigma_j^2/\text{SE})$ also tracks the share of negative estimates (8.6% against 9.9% in the cell where the two differ most).

Event bootstrap. For a ratio of pooled means $R = \sum_i a_i / \sum_i b_i$, such as (49), draw $b = 1, \dots, 1000$ assigns the events, ordered by their identifier, the multiplicities $w^{(b)}$ of a multinomial draw with N trials and equal probabilities (generator seed 20261014), and $R^{(b)} = \sum_i w_i^{(b)} a_i / \sum_i w_i^{(b)} b_i$; the 90% interval is formed by the 5th and 95th percentiles of the draws. Every subsample (period, tercile, year) is bootstrapped on its own. The intervals for the mean score differences of Table 24 are formed the same way from the paired per-event differences. The draws treat events as independent; announcements of the same week share market news, and the clustered tests of Section 8 allow for that dependence. Section 8.1, the Week-cl. columns of Tables 8, 9 and 12–13 and the week-clustered rows of Table 11 give intervals that resample announcement weeks, those of the predictions a median of 1.4 times as wide in the holdout. For the ratios of Table 28, Panel A, the draws resample announcement weeks instead: the weeks, sorted by date, receive the multinomial multiplicities of a draw with G trials and equal probabilities over the G weeks (same seed), and every event the multiplicity of its week.

C.14 The threshold rule

For Gaussian daily returns $r \sim \mathcal{N}(0, \sigma_d^2)$ without jumps, the rule of Table 25 flags $|r| > \hat{\sigma}_d$, and $\hat{\sigma}_d \rightarrow \sigma_d$. With X standard normal, the share of flagged days, the mean flagged size and the variance handed to the pricer converge to

$$\begin{aligned} \hat{\lambda} &\rightarrow 2(1 - \Phi(1)), & \frac{1}{\hat{\eta}} &\rightarrow \sigma_d \mathbb{E}[|X| \mid |X| > 1] = \sigma_d \frac{\varphi(1)}{1 - \Phi(1)}, \\ \frac{\hat{\sigma}_d^2 + 2\hat{\lambda}/\hat{\eta}^2}{\sigma_d^2} &\rightarrow 1 + 4(1 - \Phi(1)) \left(\frac{\varphi(1)}{1 - \Phi(1)} \right)^2, \end{aligned} \quad (68)$$

with φ the standard normal density and $\hat{p} \rightarrow \frac{1}{2}$: the limits are $\hat{\lambda} \rightarrow 0.317$ per day (80.0 per year), a mean flagged size of 1.525 daily standard deviations, $\hat{\eta} \rightarrow 32.8$ for a daily standard deviation of 2%, and a variance ratio of 2.48 (a volatility ratio of 1.57). None of these depends on the




Figure 10: The threshold rule of Table 25 on daily returns simulated from the Kou law of design (ii) with the jump intensity varied over a grid, for $n = 80$ and $n = 2520$ returns. (a) Reported intensity (mean and 5–95% band) against the true intensity; the dotted line is the truth. (b) Total daily variance handed to the pricer over the true variance.

data beyond σ_d : the rule returns a jump component for every continuous law. Figure 10 shows its output when the data do contain Kou jumps.

C.15 Monte Carlo designs

The four Monte Carlo studies of this appendix were run with fixed seeds by the script `code/scripts/estimate_mc.py`, which records the seeds with every result in `output/estimation_mc.json`; the designs were fixed before any simulation was run.

Threshold rule (Table 25, Figure 10): 2,000 replications of $n = 24, 80, 250$ and 2520 daily returns from (i) a Gaussian law with a daily standard deviation of 2%; (ii) the Kou law with $\sigma = 0.30$, $\lambda = 0.3$ per day (75.6 per year), $p = 0.4$, $\eta_1 = 25$ and $\eta_2 = 20$; and (iii) 72 Gaussian days with the standard deviation of (i) plus 8 announcement days carrying an additional symmetric double-exponential jump with a standard deviation of 6%, without Poisson jumps.

Maximum likelihood of the Kou law (Table 26): 200, 200, 200 and 50 replications of $n = 80, 250, 1000$ and 2520 daily returns from design (ii), with a drift that sets the mean log return to zero. The density of a daily return is the FFT inversion of the characteristic function (9) without the announcement factor, with a free daily drift, on a fixed grid of 16,384 points with spacing 3.8×10^{-4} , the log density being interpolated linearly at the observations; at seven test points the FFT density agreed with direct numerical inversion to a relative error of 3.8×10^{-5} . The parameters are $(\ln \sigma, \ln \lambda, \text{logit } p, \ln(\eta_1 - 1), \ln \eta_2)$ and the drift; each fit runs the adaptive simplex method (at most 3,000 iterations) followed by BFGS from four starting points (the threshold estimate, few large jumps, moderate jumps, many small jumps) and keeps the best end point, inside the search region stated in the notes of the table. Panel B repeats the fit on Gaussian returns without jumps, with the volatility of (i), and the profile likelihood fixes λ on the grid of the table notes and re-optimizes the other parameters on ten jump-free samples of 2,520 returns.

Announcement law (Table 23, Figure 8): 5,000 replications of $n = 8, 12, 20$ and 40 an-

nouncement returns (a finer grid for the figure) from $DE(0.45, 20, 16)$. In Panels C and D each observed move adds a $\mathcal{N}(0, k\sigma_d^2)$ diffusion with $k = 2$ trading days and $\sigma_d = 2\%$ per day. The model-selection panel fits the double exponential (closed form), the Gaussian (location and scale) and the Student t_4 (location and scale) to the same samples, and to samples from a Gaussian truth with the standard deviation of the double exponential.

Two-expiry decomposition (Table 27): 5,000 replications per cell. Two at-the-money expirations of 14 and 28 trading days, a background volatility of 0.30 and an interest rate of 0.04; the true implied variances are $\sigma_{bg}^2 + \sigma_J^2/T_j$, the observed mid of each expiration is the true price times $1 + \epsilon_j$ with ϵ_j uniform on plus or minus half the relative spread, independently across expirations and, in the two-quote columns, across the call and the put, and the implied volatility is inverted from the noisy mid.

C.16 The pooled shape and the per-firm laws of the announcement moves

Panel B of Table 5 pools the shape across firms by dividing each move by the root mean square of the other moves of the same firm in the same sample. The standardized moves had fat tails: 1.78% of them exceeded three in absolute value, several times the Gaussian share shown in the panel, and a Student t law fitted by maximum likelihood had 5.3 degrees of freedom. Part of this tail weight is the noise of the scale. A Gaussian move divided by the root mean square of m other Gaussian moves has a Student t law with m degrees of freedom, and every firm of the panel has at least eight events, so the noise of the scale alone would produce at least seven degrees of freedom. The fitted value lies below that bound: the moves have fat tails of their own, or the announcement variance of a firm changes over time, which has the same effect on the pooled law. Against a Gaussian with the same mean and variance, the Laplace law gained 0.027 nats per move, the Student t law 0.040, the double exponential with the pooled sign split 0.027 and the double exponential with diffusion 0.045 (Figure 11, panel a). The pooled law was slightly negatively skewed (skewness -0.34), with mean up moves 0.965 times as large as mean down moves.

Panel C fits one law per firm, for the firms with at least twelve announcements in the primary sample (eight in the development sample, which has at most eight per firm). The Akaike criterion picked the Gaussian for 66% of the 1,017 firms, the double exponential for 11% and the Student t_4 for 23%. For one firm the pick says little about the true shape: when the truth is double exponential, the criterion picks it in only 31.7% of the samples of twenty moves (Appendix C.3). Across firms the shares can be compared with those a known law would produce. The row of Panel C marked † draws, for every firm of the panel, 100 samples of the firm’s own number of moves from the double-exponential law of Table 23, or from a Gaussian law with mean zero and the same standard deviation, each move with Gaussian diffusion noise that gives the firm’s own diffusion share of the mean squared move (median 23.2% in the development and 7.6% in the holdout sample), applies the same criterion and reports the share of double-exponential picks; this simulation was added after the specification freeze. In the development sample a double-exponential law would give 12.7% double-exponential, 70.0% Gaussian and 17.2% t_4 picks and a Gaussian law 7.2%, 79.8% and 13.0%, against 11.7%, 70.8% and 17.4% observed, close to the double-exponential values. In the holdout the two laws would give 18.0%, 55.6% and 26.5% and 6.3%, 79.9% and 13.8%, against 4.1%, 77.0% and 18.9% observed: with twelve or more announcements per firm, the picks resemble those of a Gaussian law more than those of a double-exponential one. The share of positive moves varied across firms by little more than binomial sampling noise around a common share of 0.512: the ratio of the cross-firm variance of the share to its binomial variance was 1.06 (p -value 0.10 for homogeneity), and 1.18 in the development sample (p -value 0.02; Figure 11, panel b). The maximum-likelihood fit of the double exponential with diffusion ended on a parameter bound for 12.9% of the development firms and 0.3% of the holdout firms.




Figure 11: Announcement moves, 2020–2025 (primary-sample events). (a) Density of the standardized moves $u = x/\text{RMS}_{f,-i}$ of Table 5, Panel B (all events), on a log scale, with the Gaussian, Laplace, Student t and sign-split double-exponential laws fitted to them. (b) Share of positive moves of the firms with at least twelve events, against its distribution under a common probability of an up move (binomial simulation with the firms’ event counts and the pooled share).

C.17 Firm-level laws

Table 29 reports the announcement laws of the twenty firms with the most announcements in the primary sample. Across all firms with the maximum-likelihood fit, the standard deviations implied by the pricing estimator and by the fit with diffusion had a correlation of 0.97, so the two estimators agree on the scale. In the table they differ most for FITB, whose moves are small relative to the diffusion between the snapshots: its corrected announcement volatility is well below the root mean square of its raw moves, and the two estimators treat the diffusion differently (a variance correction against an explicit convolution). The last column is the firm’s own premium ratio (49). For firms with small announcement moves the mean of z_i^2 is close to zero and the ratio is unstable, which is one reason Section 6.7 reports pooled ratios.

Table 29: The 20 firms with the most announcements in the primary sample: announcement moves and double-exponential laws.

Firm	n	$\overline{ x }$ (%)	σ_J (%)		pricing estimator			MLE with diffusion			σ_J^Q (%)	ratio
			RMS	corr.	$\hat{p}_E$	$\hat{\eta}_{1,E}$	$\hat{\eta}_{2,E}$	$\hat{p}_E$	$\hat{\eta}_{1,E}$	$\hat{\eta}_{2,E}$		
TRV	22	3.8	4.8	4.6	0.68	31.8	27.8	0.71	29.0	23.7	4.0	0.92
FITB	22	2.4	3.1	1.4	0.36	86.5	115.3	0.44	68.5	51.0	4.0	10.79
VMI	22	7.1	8.8	8.6	0.68	17.4	14.2	0.68	15.9	13.1	6.2	0.60
MEDP	22	9.4	11.9	11.7	0.41	9.9	15.0	0.41	8.8	13.4	13.3	1.33
AON	21	4.0	4.9	4.6	0.52	29.3	31.9	0.51	24.9	28.2	3.6	0.80
CB	21	2.5	3.4	3.0	0.52	42.5	52.9	0.58	41.4	48.8	3.2	1.32
PLD	21	2.6	3.5	2.9	0.48	47.3	49.3	0.36	35.1	53.3	2.7	0.77
CBRE	21	4.5	5.9	5.4	0.76	22.2	30.0	0.72	21.6	39.6	5.4	1.05
TFC	21	3.3	4.1	3.5	0.33	41.0	38.9	0.44	40.8	30.8	3.0	4.89
HON	21	3.9	4.4	4.1	0.24	25.9	37.0	0.24	23.4	31.5	4.3	1.24
ORLY	21	4.6	5.5	5.3	0.52	29.9	24.3	0.51	28.0	21.5	5.5	1.40
EFX	21	5.5	7.2	7.0	0.33	17.2	22.0	0.27	13.0	22.2	6.7	1.21
CMC	21	4.5	5.7	5.3	0.57	24.1	30.7	0.64	25.1	25.2	6.2	1.97
IQV	21	4.8	5.9	5.6	0.62	20.3	38.8	0.61	18.8	41.3	5.6	1.02
BKR	21	3.9	4.5	3.7	0.57	35.1	43.0	0.57	31.2	44.5	4.7	2.90
DECK	21	7.8	9.5	9.1	0.71	14.4	15.7	0.76	13.8	12.3	9.5	1.17
DG	21	7.5	11.8	11.7	0.29	13.7	10.6	0.27	16.4	13.0	8.1	0.60
CINF	21	5.7	7.0	6.8	0.62	31.0	15.1	0.62	27.1	12.5	5.0	0.59
ODFL	21	5.1	6.6	6.2	0.52	20.7	25.8	0.57	20.4	20.5	5.4	0.73
MOH	21	5.7	7.2	6.9	0.48	15.0	35.1	0.54	14.3	28.2	6.5	1.37

x is the realized announcement log move between the pre- and post-event snapshots (parity-implied spots, split and dividend adjusted); $z^2 = x^2 - \sigma_{HV}^2 d/252$ removes the diffusion of the d non-announcement trading days between the snapshots, with σ_{HV} the trailing background volatility. Firms ranked by the number of announcements in the primary sample 2020–2025, ties broken by the tighter median pre-event relative spread; latest ticker shown. RMS: root mean square of x ; corr.: $(z^2)^{1/2}$. A missing σ_{HV} (1 of the 424 moves in the table) takes the median σ_{HV} of the firm’s other events (descriptive full-sample fits). The pricing estimator takes the sign split of the raw moves ($\hat{p}_E$, mean up and mean down move) and rescales both rates so that the variance of the law equals $\overline{z^2}$, with $\eta_{1,E} \geq 1.5$; the MLE maximizes the exact likelihood of $Z + \mathcal{N}(0, \sigma_{HV}^2 d/252)$, $Z \sim \text{DE}(p_E, \eta_{1,E}, \eta_{2,E})$. σ_J^Q : median two-expiry implied announcement standard deviation at the pre-event snapshot; ratio: the firm’s $\overline{\sigma_J^{2,Q}}/\overline{z^2}$ (Table 6). Written by `code/scripts/estimate_firms.py`; numbers in `output/estimation_firms.json`.

D Data details

This appendix documents the rules and counts behind Section 7, in the order in which the scripts apply them. Appendix D.7 names the scripts.

D.1 Snapshot calendar

The database’s `date` column is a snapshot label, not necessarily a trading day. I compared every label with an NYSE trading-day calendar (the exchange’s published holidays, including Good Fridays, Juneteenth from 2022 and the closure for the national day of mourning in January 2025) and mapped each of the 90 labels that are not trading days to its effective date, the previous trading day. The Saturday labels of 2019 carry the Friday chain, including the contracts that expired that Friday. A label whose effective date already had a snapshot is redundant: the 32 redundant labels, all of them holidays or weekend copies, were dropped from every derived table, and the other 58 labels were kept under their effective date. Two fingerprint tests support the rule. For the 20 most-quoted symbols, 25 labels were bit-identical to the previous snapshot, all of them redundant labels; for the whole snapshot, 1 label was. The label of the closure in January 2025 was a stale copy of the previous trading day for the first two expirations but with a reselected third expiration, so the bit-identity test missed it and the trading-day calendar

Table 30: Snapshot calendar of the option-quote database: cadence regimes, the treatment of labels that are not trading days, and the anomalies found. A snapshot is identified by its label (the database’s `date` column); a label that is not an NYSE trading day carries the closing quotes of the previous trading day (its effective date). Bit identity is tested on a fingerprint of all quote fields of the 20 most-quoted symbols. Counts are snapshot labels unless stated otherwise.

Item	<i>n</i>	Dates and rule
<i>A. Cadence regimes (snapshot labels)</i>		
Weekly, Saturday-dated	48	2019-02-09 to 2019-12-28
Monday, Wednesday, Friday	717	2020-01-04 to 2024-08-30
Every trading day	162	2024-09-02 to 2025-04-22
<i>B. Labels that are not NYSE trading days</i>		
All such labels	90	effective date = previous NYSE trading day
dropped as redundant	32	the effective date already has a snapshot: 2019-05-11, 2020-01-20, 2020-05-25, 2020-09-07, 2021-01-18, 2021-02-15, 2021-05-31, 2021-07-05, 2021-09-06, 2022-01-17, 2022-06-20, 2022-07-04, 2022-09-05, 2022-12-26, 2023-01-02, 2023-02-20, 2023-05-29, 2023-06-19, 2023-09-04, 2023-12-25, 2024-01-01, 2024-01-15, 2024-02-19, 2024-05-27, 2024-09-02, 2024-11-28, 2024-12-25, 2025-01-01, 2025-01-09, 2025-01-20, 2025-02-17, 2025-04-18
kept, Saturday-dated	46	the Friday chain (2019 regime and 2020-01-04)
kept, other	12	2019-06-30 → 2019-06-28; 2020-04-10 → 2020-04-09; 2020-07-03 → 2020-07-02; 2020-12-25 → 2020-12-24; 2021-01-01 → 2020-12-31; 2021-04-02 → 2021-04-01; 2021-12-24 → 2021-12-23; 2022-02-21 → 2022-02-18; 2022-04-15 → 2022-04-14; 2023-04-07 → 2023-04-06; 2024-03-29 → 2024-03-28; 2024-06-19 → 2024-06-18
<i>C. Other anomalies</i>		
Bit-identical to the previous snapshot (20 most-quoted symbols)	25	25 of them are among the redundant labels
Partial snapshots	10	symbols < 0.6× trailing median of the previous 10: 2020-06-08 (175 symbols, A-BEN); 2021-05-19 (796 symbols, A-MPC); 2021-05-26 (888 symbols, A-MPC); 2021-05-27 (624 symbols, MPW-ZYXI); 2021-06-14 (889 symbols, A-MPC); 2021-06-15 (625 symbols, MPC-ZYXI); 2023-07-03 (397 symbols, AA-ZI); 2023-11-01 (398 symbols, AA-ZI); 2023-11-08 (714 symbols, A-JBGS); 2024-12-26 (679 symbols, DOCS-ZWS)
Rows with expiration before the label	120,209	on 10 non-trading-day or copied labels; excluded (DTE ≥ 1 from the effective date)
Rows expiring on the label date	1,212,298	excluded (DTE ≥ 1 from the effective date)
Longest gap between effective snapshots	10	trading days, 2024-07-29 to 2024-08-12

caught it. Table 30 lists the regimes, the labels and the other anomalies: 10 partial snapshots, 120,209 rows that expired before their label (on 10 labels) and 1,212,298 rows expiring on the label date. The parity and implied-volatility calculations use only contracts with at least one calendar day to expiration from the effective date.

Some of the partial snapshots are alphabetical halves of one chain collected over two days, which is why the matching of Section 7.2 works symbol by symbol. Table 31 summarizes the structure of the database by year, and Figure 12 shows the cadence regimes and the days to expiration at which monthly and weekly expirations were quoted. The coverage of trading days rose from 59% in the Monday–Wednesday–Friday regime to 97% in the daily regime. From 2020 on, three or more expirations were quoted on 35% of the symbol-dates, and the median days to expiration of a weekly series at its last quote was 11. An expiration is monthly if it falls on the third Friday of its month, or on the Thursday before it when that Friday is an exchange holiday;

Table 31: Structure of the option-quote database by snapshot year. Snapshots are end-of-day chains; ‘effective’ counts distinct trading days after mapping weekend- and holiday-dated snapshots to the preceding trading day and dropping bit-identical copies. DTE is calendar days from the snapshot label to expiration; shares are of (date, symbol, expiration) cells. Weeklies are expirations that are not the standard third-Friday monthly. Spreads are (ask – bid)/mid over quotes with bid > 0; ‘core’ restricts to $0.15 \leq |\Delta| \leq 0.85$ and mid ≥ 0.20 (vendor delta). The 2019 snapshots are weekly (Saturday-dated); 2025 ends on 22 April.

	2019	2020	2021	2022	2023	2024	2025	All
<i>Snapshots</i>								
Snapshot dates	48	155	156	151	155	183	79	927
Effective trading days	47	152	151	146	149	176	74	895
Holiday/weekend copies dropped	1	3	5	5	6	7	5	32
Trading days covered (%)	18.7	60.1	59.5	58.2	59.6	69.8	98.7	
Symbols per snapshot (median)	1,053	1,635	1,517	1,521	1,520	1,525	1,525	
Quotes per snapshot (median, 000s)	36.8	89.3	92.6	90.6	90.3	93.2	71.3	
<i>Expirations per (symbol, date)</i>								
Mean	2.30	2.35	2.35	2.34	2.36	2.35	2.37	
Share with 3 or more (%)	39.5	34.9	34.7	33.7	36.1	34.9	37.3	
Nearest expiration: median DTE	13	14	14	14	14	14	15	
<i>DTE of quoted expirations (% of cells)</i>								
0–7	6.6	10.1	10.3	10.1	10.3	9.1	8.9	
8–14	19.7	13.8	13.2	13.4	13.4	12.5	12.0	
15–30	27.8	30.8	30.0	29.6	29.5	30.4	30.2	
31–45	14.2	16.7	16.8	18.1	17.2	18.5	21.9	
46–60	22.4	23.2	24.0	22.4	23.1	22.1	20.4	
61–90	7.0	5.4	5.6	6.2	6.4	7.4	6.6	
Weeklies, share of cells (%)	27.7	17.3	16.7	15.9	15.8	15.9	15.9	
Weeklies quoted inside 11 DTE (%)	0.0	7.0	5.1	0.2	2.7	1.5	0.1	
Monthlies quoted inside 11 DTE (%)	68.4	76.6	77.5	80.9	78.9	78.8	54.4	
<i>Quotes</i>								
Median relative spread, all (%)	15.7	22.2	19.4	18.6	16.1	22.4	24.5	20.0
Median relative spread, core (%)	12.7	22.6	19.0	18.0	15.6	22.5	25.0	
Zero-bid quotes (%)	12.4	17.5	19.5	16.6	19.4	15.1	10.1	16.9
Vendor IV missing (%)	0.3	1.2	0.7	0.8	0.7	0.4	0.7	0.7
<i>Parity-implied spot (per symbol-date)</i>								
Valid call–put pairs, nearest expiry (median)	5	5	5	5	5	6	6	5
Pass filter (≥ 2 pairs, dispersion $\leq 1\%$) (%)	91.5	67.1	72.2	67.4	68.0	65.5	65.7	68.5

every other expiration is a weekly. Weeklies made up 16–17% of the (date, symbol, expiration) cells in each year from 2020 on.

D.2 Earnings calendar, firm identifiers and matching

Table 32 follows the earnings calendar to the matched announcements. I kept the 38,363 rows of the 2,088 symbols that appear in the option database and dropped the 14 of them dated on a weekend. The superseded projections described in Section 7.2 are visible in the whole table: 3,030 rows follow the previous row of the same symbol by less than 30 days, 98% of them without a timing flag; merging them produced 57 announcements in the window that were merged from more than one calendar row. Of the announcements with a missing timing flag, treated as made after the close, 215 are in the final sample, where they carry a flag. Of the 2,183 symbols of the option database, 95 have no calendar row at all, 38 of them flagged as




Figure 12: Structure of the option-quote database. (a) Snapshots per month, by cadence regime: weekly Saturday-dated snapshots in 2019, Monday–Wednesday–Friday snapshots from 4 January 2020 to 30 August 2024, daily snapshots from 2 September 2024. (b) Quoted (date, symbol, expiration) cells by calendar days to expiration at the snapshot, per year, in 7-day bins, split into monthly (third-Friday) and weekly expirations; the dotted line marks 11 days. Weekly expirations are almost never quoted inside that line.

exchange-traded funds. The exchange-traded-fund flag of the DoltHub symbol table is current rather than point-in-time, and 16 symbols that it flags announced earnings in the sample window because an exchange-traded fund later reused the ticker of a delisted or renamed company, so I did not use the flag; the operational rule is the presence of calendar rows.

Firm identifiers matter for the firm’s past announcement moves and for its trailing return series. I linked a ticker change into one firm history only when both tickers are in the database, the old ticker stops being quoted before the new one starts (or the new one is a reused ticker split at a known date), the old holders kept their shares one for one, and the surviving company is predominantly the old one. The link date is the day after the old ticker’s last snapshot, because the vendor switched tickers at month ends. Table 33 lists the 38 linked changes, the 20 changes left unlinked by design, the 1 rejected candidate and the 8 reused tickers whose earlier rows belong to a different company. The 36,543 announcements in the window belong to 2,060 firms.

The matching ran symbol by symbol, so a partial snapshot served the symbols it contains: 121 events of the primary sample used a partial snapshot at d_{pre} and 96 at d_{post} . Most unmatched announcements belong to symbols that the vendor did not quote around the announcement: the

Table 32: From the earnings calendar to matched announcements. Panel A cleans the calendar (rows of one symbol less than 30 days apart are one announcement; the last date is kept, preferring a row with a non-missing timing flag). Panel B matches the pre-event snapshot d_{pre} (latest snapshot on or before T_{pre}) and the post-event snapshot d_{post} (earliest on or after T_{post}), where T_{pre} and T_{post} are the last close before and the first close after the announcement. Panel C gives the distribution of the lags and of the number of ordinary trading-day returns between d_{pre} and d_{post} besides the announcement return. Development: announcements in 2020–2021; holdout: 2022 to 22 April 2025.

	Dev.	Holdout	All				
<i>A. Earnings calendar</i>							
Calendar rows, all symbols			95,112				
of symbols in the options database			38,363				
dated on a weekend (dropped)			14				
in the window, before merging rolled dates			36,801				
Announcements in the window	13,831	22,712	36,543				
merged from more than one calendar row	24	33	57				
after market close	7,839	12,787	20,626				
before market open	5,992	9,925	15,917				
timing missing, treated as after close	995	1,595	2,590				
date not a trading day	14	41	55				
<i>B. Snapshot matching (lag at most 3 trading days)</i>							
Matched to d_{pre} and d_{post}	11,282	17,845	29,127				
Unmatched: no snapshot within the lag on either side	2,517	4,037	6,554				
symbol not quoted within 30 calendar days	2,476	3,847	6,323				
Unmatched: no pre-event snapshot only	14	294	308				
Unmatched: no post-event snapshot only	18	536	554				
Announcements 2024-07-26 to 2024-08-12 (snapshot gap)	0	1,032	1,032				
of which unmatched	0	1,005	1,005				
<i>C. Events with a primary expiration (21,604), by lag in trading days</i>							
Lag	0	1	2	3	4	5	6
d_{pre} before T_{pre}	13,292	8,034	202	76			
d_{post} after T_{post}	10,851	10,350	329	74			
Ordinary trading days between the snapshots	2,919	17,806	560	311	4	3	1

vendor’s list of symbols changed over time, and a row in the calendar did not imply quotes. Of the remaining 1,093 unmatched announcements, 841 fell into the one long interruption of the database, between the snapshots of 29 July 2024 and 12 August 2024. In the primary sample, d_{pre} fell at most one trading day before T_{pre} for 98.7% of the events and d_{post} at most one trading day after T_{post} for 98.1%. 20 events of the primary sample have an announcement date that is an exchange holiday; the timing rule of Section 7.2 handles them without change. The number of ordinary trading-day returns between the two snapshots besides the announcement return (Panel C of Table 32) is the diffusion term that the realized announcement variance removes.

Announcements without a timing flag. Section 7.2 compares the 917 announcements with a primary expiration whose timing was imputed with the 20,687 whose timing was flagged. The smaller fall of the at-the-money implied volatility of the nearest expiration from d_{pre} to d_{post} is not a matter of price level or of the flag that the others carry: by development-sample tercile of the spot its medians were 2.6–4.6 volatility points with an imputed timing against 9.6–11.4 with a flag, and the median was 11.2 points for the reports flagged after the close and 9.0 for those flagged before the open. Two other statistics point the same way. The two-expiry decomposition

Table 33: Ticker changes linked into one firm history (38). The earnings calendar and the option database key each announcement by the ticker in force at the time, so events match under their own symbol; the link only joins the firm’s past announcements and its trailing return series. A change is linked when the old holders kept their shares one for one and the surviving company is predominantly the old one; the link date is the day after the old ticker’s last snapshot. Mergers of near-equals and spin-offs are not linked.

Old	New	Linked from	Old	New	Linked from	Old	New	Linked from
BHGE	BKR	2019-10-12	WLTW	WTW	2022-01-08	RCII	UPBD	2023-02-25
HCP	PEAK	2019-11-02	CRY	AORT	2022-01-22	PKI	RVTY	2023-05-16
SYMC	NLOK	2019-11-02	VIAC	PARA	2022-02-17	FISV	FI	2023-06-03
CTRP	TCOM	2019-11-02	HFC	DINO	2022-03-15	INT	WKC	2023-06-15
JEC	J	2019-12-07	ADS	BFH	2022-04-02	RE	EG	2023-07-08
WTR	WTRG	2020-02-01	SGMS	LNW	2022-04-28	ABC	COR	2023-08-29
CTL	LUMN	2020-09-17	BLL	BALL	2022-05-10	CDAY	DAY	2024-02-01
IRET	CSR	2020-12-19	FB	META	2022-06-09	PEAK	DOC	2024-03-02
AAXN	AXON	2021-01-26	ANTM	ELV	2022-06-28	FLT	CPAY	2024-03-23
WYND	TNL	2021-02-13	NLOK	GEN	2022-11-08	PNM	TXNM	2024-07-30
ETH	ETD	2021-08-14	TBK	TFIN	2022-12-01	GPS	GAP	2024-08-22
CREE	WOLF	2021-10-02	FBHS	FBIN	2022-12-15	NYCB	FLG	2024-10-26
MLHR	MLKN	2021-10-30	BCOR	AVTA	2023-01-26			

Not linked (20): LB→BBWI (spin-off of Victoria’s Secret on the renaming day); UTX→RTX (spin-offs of Carrier and Otis plus merger with Raytheon); RTN→RTX (merger); DISCA→WBD (merger, Discovery holders 29%); MYL→VTRS (merger with Upjohn, Mylan holders 43%); BCEI→CIVI (merger of near-equals); BBT→TFC (merger of near-equals); COG→CTRA (merger of near-equals); IR→TT (spin-off; the IR ticker passed to Gardner Denver); ARNC→HWM (spin-off of the new Arconic); AAN→PRG (spin-off of the new Aaron’s); DISH→SATS (merger); WRK→SW (merger); VIAB→VIAC (merger); INFO→SPGI (merger); RBC→RRX (merger, Regal holders 61%); TOT→TTE (new ticker not in the database); RDS.A→SHEL (new ticker not in the database); SNE→SONY (new ticker not in the database); SQ→XYZ (neither ticker in the database). *Rejected* (1): ROLL→RBC (ticker not in the options database).

Reused tickers (8): the rows of the symbol before the date belong to a different firm: COR before 2023-08-29; DOC before 2024-03-02; IR before 2020-03-02; ARNC before 2020-04-01; AAN before 2020-12-01; RBC before 2021-11-01; CHK before 2021-02-10; VAL before 2021-05-01.

of Section 6.6 was negative for 16.2% of the imputed-timing announcements, against 7.1%, and the root-mean-square realized move was 6.8%, against 8.3%, as expected when for some of them both snapshots follow the release. The post-freeze check of Section 7.2 repeated the evaluation of Section 8 with the frozen prices on the common sample without its 31 development and 86 holdout events with an imputed timing. The vega-scaled event-average RMSE of BS-HV became 11.21 and 13.68 (11.29 and 13.71 with them), that of DJ-hist 8.57 and 8.34 (8.62 and 8.37), and that of Kou-SEJ-Q 3.78 and 4.24 (3.81 and 4.24). DJ-hist reduced the RMSE of BS-HV by 23.6% and 39.1% (a component of P1; 23.7% and 38.9% with them), and Kou-SEJ-Q that of Kou-SEJ-P by 55.0% and 51.4% (P5; 55.0% and 51.6%). Under the pre-registered metric the largest move of a component of P1 or P5 was 0.4 percentage points, and none changed its verdict. Without the 380 development and 502 holdout announcements with an imputed timing and a decomposition, the premium ratio was 1.67 [1.56, 1.79] over 7,846 announcements in the development period and 1.13 [1.08, 1.18] over 12,507 in the holdout. `sample_timing.py` computes these numbers with the functions of the main analysis, after checking that on the full sample they reproduce its stored results.

D.3 Parity-implied spot and the tolerance rule

The parity calculation treats the quotes as European with $q = 0$, uses the 1-month rate of the effective date and measures time in calendar days ($T = \text{DTE}/365$); a call–put pair is valid when both bids are positive and both asks exceed the bids. The spot of a symbol-date uses the nearest

Table 34: The parity tolerance rule against the plain 1% rule. At each snapshot the spot is the median of $C_{\text{mid}} - P_{\text{mid}} + Ke^{-rT}$ over the three call-put pairs nearest the money of the nearest expiration with at least two valid pairs; the dispersion is the range of the three values over their median. The plain rule requires at least two pairs and dispersion $\leq 1\%$; the tolerance rule requires at least two pairs and dispersion $\times S \leq \max(1\% S, 2 \times \text{ATM half-spread})$, the ATM half-spread being the mean half-spread of the call and put quotes of the three pairs. Upper panel: events with a primary expiration (tolerance rule and spot-consistency screen at both dates), by development-sample price tercile; half-spreads and dispersions at d_{pre} ; ‘fail’ counts events that fail the plain rule at d_{pre} or d_{post} .

	Price tercile at d_{pre}			All
	1 (low)	2	3 (high)	
Spot range (USD)	0.55–36.61	36.61–89.14	89.15–5,020.81	
Events with a primary expiration	6,480	7,245	7,879	21,604
median ATM half-spread (USD)	0.22	0.45	0.79	0.46
median ATM half-spread / spot (%)	1.23	0.76	0.41	0.66
median dispersion of the parity spots (%)	0.79	0.56	0.26	0.44
fail the plain 1% rule	3,736	3,314	1,173	8,223
fail the plain 1% rule (%)	57.7	45.7	14.9	38.1
Events with ≥ 1 kept quote	1,892	3,052	5,052	9,996
fail the plain 1% rule	271	271	142	684
<i>Matched announcements (29,127)</i>				
Parity spot on both dates				27,274
Tolerance rule on both dates				23,852
and the spot-consistency screen on both dates				23,851
Plain 1% rule on both dates				15,479
Pass the tolerance rule, fail the plain rule				8,373
Pass the plain rule, fail the tolerance rule				0

expiration with at least one day to expiration and at least 2 valid pairs, and among its pairs the 3 with the smallest $|C_{\text{mid}} - P_{\text{mid}}|$. Over all 1,248,308 symbol-dates, the median number of valid pairs was 5, the median dispersion 0.44% and its 90th percentile 3.0%, and 68.5% of the symbol-dates passed the plain rule of at least 2 pairs and dispersion at most 1%.

For the event dates I read the exact bid and ask of the parity pairs from the database and applied the tolerance rule of Section 7.3: at least 2 valid pairs and

$$\text{dispersion} \times S \leq \max(1\% \cdot S, 2h_{\text{ATM}}),$$

where h_{ATM} is the mean half-spread of the call and put quotes of the pairs used. At d_{pre} the median h_{ATM} of the matched announcements was \$0.40 and the median tolerance 1.32% of the spot. Because the tolerance is never below 1% of the spot, the rule accepts every spot that the plain rule accepts (0 exceptions). Of the 27,274 matched announcements with a parity spot on both dates, 23,852 passed the tolerance rule on both and 15,479 the plain rule; 8,373 passed only the tolerance rule. Table 34 shows that the difference comes from the low-priced stocks, whose at-the-money spreads are wide relative to the spot, so that mids that differ by a fraction of the spread give parity values that differ by more than 1%: among the events of the primary sample the plain rule would have rejected 58% of those in the lowest price tercile and 15% of those in the highest, and among the events with a kept quote it would have removed 684 (7%).

Spot-consistency screen. Let S_{t-1} and S_{t+1} be the parity spots of the same ticker, within one firm identity, at the previous and at the next snapshot at which it has one. The screen

Table 35: Parity-implied spot versus the daily close (DoltHub OHLCV) for the 60 most-quoted symbols. The deviation is $|S_{\text{parity}}/\text{Close} - 1|$ on common trading dates; ‘filtered’ keeps the (date, symbol) pairs that pass the plain parity filter (≥ 2 pairs, dispersion $\leq 1\%$).

Year	All pairs				Filtered		
	n	Median (%)	P95 (%)	$\leq 1\%$ (%)	n	Median (%)	P95 (%)
2019	2,709	0.03	0.38	100.0	2,707	0.03	0.38
2020	9,077	0.07	0.50	99.3	8,915	0.07	0.48
2021	8,791	0.05	0.36	100.0	8,729	0.05	0.35
2022	8,708	0.05	0.39	99.9	8,681	0.05	0.38
2023	8,772	0.04	0.38	99.9	8,709	0.04	0.37
2024	10,505	0.06	0.40	99.9	10,388	0.06	0.39
2025	4,438	0.08	0.96	95.2	4,346	0.08	0.93
All	53,000	0.05	0.43	99.4	52,475	0.05	0.42

statistic of the parity spot S_t is

$$c_t = \begin{cases} \min(|\ln(S_t/S_{t-1})|, |\ln(S_t/S_{t+1})|) & \text{if the two log ratios have the same sign,} \\ 0 & \text{otherwise,} \end{cases}$$

and the spot fails when $c_t > \ln 3 = 1.10$. The statistic is large only for a jump that reverts at the next snapshot: an announcement move that persists leaves the post-event spot close to the next one, and c_t close to zero. The first and the last spot of a series have one neighbor and are not tested. I set the threshold from the distribution of c_t over the 1,243,993 of the 1,248,307 parity spots that have two neighbors, before the sample was rebuilt and without reference to any pricing error. Its 99.99th percentile is 0.27. The largest value that passes is 0.93 (GNW on 20 September 2021, a spot from two pairs with a dispersion of 31%), and the 3 spots that fail have values from 1.33 to 1.80; 2 of them equal the parity spot of SPY on the same day. Every threshold between 0.93 and 1.33, that is every factor between 2.53 and 3.80, rejects the same spots. A wrong chain that persists for several snapshots would escape the statistic; no spot departs from its ticker’s level by more than the factor and returns within two to 5 snapshots (0 such runs).

The screen changed the inputs in three places. The announcement of SR on 4 February 2021 left the primary sample: its post-event spot gave it a log move of 1.836, which had entered the past moves, the historical announcement variance and the past-move tercile of the firm’s 15 later announcements; none of its quotes had passed the quote filters. The other announcement with a failing spot had already failed the tolerance rule. And 2 failing spots left the firms’ return series, which lowered σ_{HV} of 2 announcements of SSD from 326–349% to 20–25%. Section 7.3 and deviation (viii) of Section 8.1 state when the screen was added and how the main study was priced again; the tercile cut points of Table 39 are those of the corrected development sample.

Table 35 and Figure 13 compare the parity spot with the daily close of the 60 most-quoted symbols (97,862 closes downloaded one trading day per request, 1,632 requests). Over the 53,000 common symbol-dates the median absolute deviation was 0.05% and its 95th percentile 0.43%, 99.4% of the deviations were below 1%, and 0 looked like an unadjusted split. The weakest year was 2025, in which 95.2% of the deviations were below 1%; the monthly 95th percentile peaked at 2.5% in March 2025 (Figure 13, panel b). The check covers liquid names only.

D.4 Background volatility and event-level inputs

The background volatility σ_{HV} of an event comes from the firm’s own series of parity-implied spots, so that it measures the same object as the spot used for pricing. The series contains




Figure 13: Parity-implied spot against the daily close (DoltHub OHLCV) for the 60 most-quoted symbols. (a) A random draw of the common symbol-dates, log scales, with the 45-degree line. (b) Median and 95th percentile of $|S_{\text{parity}}/\text{close} - 1|$ by month.

every non-redundant snapshot at which the firm’s spot passed the tolerance rule (1,037,728 firm-dates), less the 2 spots that fail the spot-consistency screen of Appendix D.3; because the exact half-spread was read from the database only for the event dates, the series uses an approximate h_{ATM} , half the median relative spread of the expiration times the mean of the at-the-money call and put mids. A return $\varrho_i = \ln((S_i + D_i)/S_{i-1})$ runs between consecutive snapshots of the firm, with D_i the dividends going ex in the interval and a split adjustment where a split falls in it; g_i is the number of trading days it spans. The window of an event contains the returns that end on or before d_{pre} and start at most 120 trading days before it, minus three groups: returns spanning more than 10 trading days; returns that overlap the interval from 1 trading day before T_{pre} to 1 trading day after T_{post} of any announcement of the firm in the cleaned calendar, matched or not (the current one included); and returns that overlap the corresponding interval around announcement dates projected backward from the firm’s first calendar announcement in steps of 91 days, which stand in for the announcements before 22 January 2020, where the calendar has no rows (the timing of a projected date is unknown, so the interval covers both timings). The estimator is

$$\sigma_{\text{HV}}^2 = 252 \frac{\sum_i \varrho_i^2}{\sum_i g_i},$$

a zero-mean variance per trading day, annualized, that weights each return by the time it spans; it requires at least ten returns covering at least twenty trading days. It was available for 28,213 of the 29,127 matched announcements, with a median of 33% and a median of 54 returns in the window; 2,768 windows reached back before the start of the calendar, and 2,000 windows excluded returns around a projected date. For the first months of 2020 the window reaches into the weekly regime of 2019, whose returns span about a week each. Table 36 reports these and the other event-level inputs for the events with a kept quote.

The at-the-money implied volatilities that enter the two-expiry decomposition and the at-the-money column of Table 38 were computed per snapshot, symbol and expiration at the strike nearest the parity spot, as the mean of the call and the put implied volatilities, with $q = 0$

Table 36: Event-level inputs at d_{pre} for the announcements with at least one kept quote. σ_{HV} is the background volatility from the firm’s parity-implied spot returns over the trailing 120 trading days, excluding returns that overlap a past announcement (Appendix D.4). The two-expiry decomposition uses the ATM implied volatilities of the two shortest expirations that contain the announcement. Past announcements are the firm’s earlier matched announcements whose post-event snapshot precedes d_{pre} and whose realized move is defined (spot within tolerance on both dates); the historical announcement variance needs at least 4.

	Dev.	Holdout	All
Events	3,861	6,135	9,996
Firms	1,003	1,052	1,281
After market close (%)	51.8	50.5	51.0
Timing missing, treated as after close (%)	3.0	1.6	2.2
$d_{\text{pre}} = T_{\text{pre}}$ (%)	57.1	64.0	61.3
$d_{\text{post}} = T_{\text{post}}$ (%)	45.1	54.5	50.9
Median spot S_{pre} (USD)	83.07	96.25	90.45
Median 1-month rate at d_{pre} (%)	0.07	4.67	2.20
Median VIX at d_{pre}	21.17	18.84	19.15
Ex-dividend date in a kept option’s life (%)	25.0	29.3	27.6
σ_{HV} available (%)	99.4	99.5	99.5
median σ_{HV} (%)	36.3	30.4	32.1
median returns in the window	60	64	63
median returns excluded around past announcements	4	5	4
Two-expiry decomposition available (%)	99.6	100.0	99.8
negative implied announcement variance (% of available)	5.0	1.6	2.9
At least 4 past announcements (%)	42.6	94.1	74.2
median number of past announcements	3	12	8

and $T = \text{DTE}/365$, the convention of the parity calculation. The quote-level implied volatilities of Section 7.4 use trading time and the dividend yield of the option’s life instead. For a given price the two conventions give nearly the same total implied variance $\sigma^2 T$ (they differ through the discount and dividend terms), so the annualized volatilities differ mainly by the ratio of the two time measures. The two-expiry decomposition, however, depends on the ratio of the two maturities, which weekends and holidays change; it is computed in calendar time.

D.5 Quote filters

The implied volatilities of the quotes were inverted by bisection to machine precision. Table 37 complements the cumulative funnel of Table 7. Panel A counts the quotes that fail each filter on its own: the spread filter and the delta filter are the binding ones. The strikes rule acts on whole (event, expiration) pairs; applied last, it removed comparatively few quotes but many events (Section 7.4). Panel B concerns the quotes whose mid lies at or below the European lower bound and that therefore have no European implied volatility. They were deep in the money (median $\ln(K/S)$ of 0.21 for the puts and -0.30 for the calls) and close to the bound (median shortfall 0.22 and 0.12 of the half-spread). Most of them had an ask at or above the bound (90% of the puts and 99% of the calls), so the mid fell below the bound by a fraction of the spread; 61% of the puts had an ex-dividend date in their life; and at the vendor’s delta only 5% of the puts and 1% of the calls would have passed the delta filter, so the bound filter removed few quotes that the delta filter would have kept. 23 quotes had a mid at or above the upper bound. Panel C links the kept quotes to d_{post} ; the 138,928 post-event quotes on contracts that were not quoted at d_{pre} are not used. Panel D is the secondary sample. Its settlement spot is, in order of preference, the parity spot of the expiration day that passes the tolerance rule, the same one trading day earlier, and the expiration-day parity spot that fails the rule (flagged as




Figure 14: The kept primary sample. (a) Events with at least one kept quote per announcement quarter, development (2020–2021) and holdout (from 2022). The third quarter of 2024 is thin because the database has no snapshot between 29 July 2024 and 12 August 2024, and the last quarter ends on 22 April 2025. (b) Kept quotes by the pre-registered bins of $|\Delta_{BS}|$ (rows) and of days to expiration at d_{pre} (columns). (c) Distribution of $\ln(K/S_{pre})$ of the kept calls and puts.

not validated). 141 (event, weekly expiration) pairs have no settlement spot; 103 of them had no snapshot on the expiration day or the trading day before, 100 of those in 2022.

On the kept quotes, inverted again with time measured in calendar days, my implied volatilities differ from the vendor’s by a median of 0.62 volatility points, against 1.48 in trading time, so most of the gap reflects the vendor’s calendar-time convention. On all pre-event quotes with an implied volatility, the correlation between my implied volatilities and the vendor’s was only 0.58, against 0.979 on the kept quotes; the difference comes mostly from quotes that the filters remove, whose implied volatility is poorly determined (far from the money, wide spreads, mids near a bound). As a check of the quote pipeline, 3 randomly drawn events were rebuilt from the raw rows with an independent scalar implementation (Brent’s method for the implied volatility, closed-form Greeks, row-by-row filter logic): the contracts, the filter flags and the post-event links agreed exactly, and the largest implied-volatility difference was 2×10^{-13} .

Table 38 describes the kept primary sample by announcement year, and Figure 14 by announcement quarter, by the pre-registered bins of $|\Delta_{BS}|$ and days to expiration, and by log-moneyness; 3,866 events had two or more kept expirations.

D.6 Tercile cut points

The pre-registered cuts by price level, past announcement moves, VIX and relative spread use terciles whose cut points were computed on the development events with a primary expiration and applied unchanged to the holdout (Table 39). The price variable is the parity spot at d_{pre} , a proxy for size and liquidity because the data contain no market capitalization; the past-move variable is the mean absolute realized move of the firm’s earlier announcements in the sample and needs at least one of them; the spread variable is the median, over the event’s primary expirations, of the median relative spread of the expiration’s quotes at d_{pre} . The remaining pre-registered cut, by the ratio of option-implied to historical announcement variance, follows the same rule and is defined only for the events with at least 4 past announcements. The moneyness and maturity bins are fixed intervals of $|\Delta_{BS}|$ and of days to expiration and need no cut points. The filters of Section 7.4 kept at least one quote for 50%, 46% and 41% of the events with a primary expiration by VIX tercile, and for 45%, 49% and 47% by tercile of the firm’s past

Table 37: Quote filters in detail. Panel A counts the pre-event quotes of the primary expirations that fail each filter whatever the others do (the cumulative counts are in Table 7). Panel B describes the quotes whose mid lies at or below the European lower bound (discounted intrinsic value with the dividend yield), for which no European implied volatility exists. Panel C links each kept quote to the same contract at d_{post} . Panel D is the secondary sample; the settlement spot uses the tolerance rule of Table 34. Settlement spot by source, in (event, weekly expiration) pairs: parity spot on the expiration day 1,594; parity spot one trading day earlier 23; expiration-day spot failing the tolerance rule 79; none 141.

	Dev.	Holdout	All
<i>A. Pre-event quotes failing each filter (each filter on its own)</i>			
Pulled quotes (all strikes, calls and puts)	299,894	492,376	792,270
bid > 0	43,262	62,675	105,937
ask > bid	56	593	649
relative spread $\leq 20\%$	176,713	280,407	457,120
mid $\geq \$0.20$	21,729	35,865	57,594
mid inside the European no-arbitrage bounds	18,116	25,919	44,035
implied volatility in (0.05, 3)	19,909	26,036	45,945
$0.15 \leq \Delta_{BS} \leq 0.85$ at own IV	150,823	257,258	408,081
≥ 2 strikes per side and expiration	129,258	198,108	327,366
<i>B. European bounds, pulled quotes</i>			
Mid at or below the European lower bound, puts	11,269	14,571	25,840
Mid at or below the European lower bound, calls	6,770	11,117	17,887
Mid at or above the upper bound	23	0	23
after the first four filters (puts / calls):			
quotes		14,873 / 11,681	
median $\ln(K/S)$		0.21 / -0.30	
median shortfall below the bound / half-spread		0.22 / 0.12	
ask at or above the bound (%)		89.5 / 99.5	
ex-dividend date in the option's life (%)		61.0 / 4.9	
inside the delta range at the vendor's delta (%)		5.1 / 1.1	
<i>C. Kept pre-event quotes at d_{post}</i>			
Kept quotes	54,840	97,129	151,969
same contract quoted at d_{post}	42,817	77,464	120,281
post-event quote passes the quote filters	22,234	39,274	61,508
<i>D. Earnings-week weeklies</i>			
Quotes pulled	8,170	30,318	38,488
kept	390	1,887	2,277
kept, relative spread $\leq 10\%$	145	826	971
kept, validated settlement spot	379	1,768	2,147

announcement moves.

D.7 Scripts

The data pipeline consists of scripts in `code/scripts/`, run in this order. `data_coverage.py` builds the snapshot calendar and the coverage index of every (date, symbol, expiration) cell; `data_parity_spot.py` computes the parity spots and at-the-money implied volatilities; `data_rates.py`, `data_earnings.py` and `data_ohlc_subset.py` download the rates and VIX, the earnings calendar, symbol table, dividends and splits, and the closing prices, with every request logged and every file recorded with its retrieval time and checksum; `data_ohlc_subset.py` also writes Table 35, and `data_tables.py` writes Table 31 and Figures 12 and 13. `sample_events.py` cleans the calendar, links firm identifiers, matches the snapshots, applies the tolerance rule and the

Table 38: The kept primary sample by announcement year. Events and firms with at least one kept pre-event quote. DTE (calendar days from d_{pre} to expiration) and the relative spread (ask – bid)/mid are medians over kept quotes. ATM IV and the straddle-implied move (ATM call mid plus put mid over spot) are medians over events, taken from the shortest kept expiration at d_{pre} . The realized move is $|\ln(S_{\text{post}}/S_{\text{pre}})|$ with parity-implied spots, adjusted for dividends and splits. 2025 ends on 22 April.

Year	Events	Firms	Kept quotes	Median				
				DTE	spread (%)	ATM IV (%)	straddle (%)	move (%)
2020	1,883	837	27,251	25	10.8	50.4	9.8	4.1
2021	1,978	812	27,589	25	10.3	42.0	8.3	3.4
2022	1,783	733	29,856	28	9.5	48.3	9.5	4.7
2023	2,029	760	31,473	25	8.7	41.0	7.7	4.4
2024	1,761	788	26,290	23	9.4	42.2	7.9	4.8
2025	562	527	9,510	28	10.1	46.4	8.8	4.5
Development	3,861	1,003	54,840	25	10.5	46.2	8.9	3.8
Holdout	6,135	1,052	97,129	25	9.2	43.8	8.3	4.6
All	9,996	1,281	151,969	25	9.7	44.6	8.5	4.3

Table 39: Tercile cut points of the pre-registered cuts, fixed on the development sample (events with a primary expiration announced in 2020–2021; n is the number with the variable available) and applied unchanged to the holdout. A value at or below $q_{1/3}$ is tercile 1 (T1), above $q_{2/3}$ tercile 3 (T3). The last three columns give, for each tercile, the share of events with a primary expiration (both samples) that keep at least one quote after the filters. Variables: the parity spot at d_{pre} ; the mean absolute realized move of the firm’s earlier announcements in the sample (at least one needed); the VIX close at d_{pre} ; the median over the event’s primary expirations of the median relative spread of the expiration’s quotes at d_{pre} .

Variable	Development cut points			Kept (%)		
	$q_{1/3}$	$q_{2/3}$	n	T1	T2	T3
Spot S_{pre} (USD)	36.61	89.14	8,407	29	42	64
Mean past move (%)	3.70	6.51	6,930	45	49	47
VIX at d_{pre}	17.97	24.52	8,407	50	46	41
Median relative spread (%)	17.2	31.8	8,407	90	49	6

spot-consistency screen and computes the event-level inputs; `sample_quotes.py` pulls the quotes of the sample with one key-table join per year, computes the implied volatilities and the filter flags and writes Tables 7 and 38 and Figure 14; `data_sample_tables.py` writes the other tables of this appendix and the numbers quoted in Section 7. As a check, `data_sample_tables.py` also recomputes 25 of the counts and cut points reported here from the derived files by a second route (the funnel counts, the kept sample, the tercile cut points, the plain-rule failures, the unmatched announcements, and the spots and announcements that fail the consistency screen, from an independent reimplement of the screen); all 25 agreed with the values written by the earlier scripts. `sample_timing.py`, run after the analysis of Section 8, writes the post-freeze check of the announcements without a timing flag (Appendix D.2). The recent announcements of Section 8.10 come from `data_recent.py`. The pricing scripts read the quote and event files described here; Appendix B states the solver settings they use.

E Ex-ante model prices for announcements after 12 October 2026

Table 40 records, before the events, the model prices for the announcements of the firms of Section 8.10 scheduled between 13 October and 30 November 2026. The prices use only the end-of-day chain of 12 October 2026 and information available at that close, with the models and the calibration of Section 8.10: Kou-SEJ-Q is fitted to the two shortest expirations that contain the announcement, and every column of the table refers to the first of them. Each of the 24 scheduled announcements has two such expirations and is priced (24 rows). The straddle and 25-delta columns use only quotes that pass the filters of Section 8.10, the straddle at the strike nearest the spot whose call and put both pass them; where no strike does, the straddle columns are left empty.

The table was first written on 15 October 2026, 13:01 UTC, before the first announcement it covers (20 October 2026). The SHA-256 hash of the table file begins `c50765812debd1c0`. A later run of the script, on 15 October 2026, 13:15 UTC, reproduced the identical table. The full hashes of the table and of the file of all priced quotes are in `output/exante_hash.txt` of the code repository and of the replication archive, and every recorded write of the table, with its time and hash, is in `output/exante_stamp.json`.

A first version of the table, first written on 13 October 2026, 23:26 UTC (hash beginning `e318e749763a2544`), took the calibration pair and every column from the nearest quoted expiration. For 20 of its 22 priced events that expiration ends before the announcement, so its move and smile columns did not forecast the announcement, and it left 2 events unpriced whose second quoted expiration preceded the event. I corrected the error before the first announcement of the table, as a recorded amendment of the frozen code in `output/spec_freeze.json`; the first version is archived in `output/exante_superseded/e318e749763a/exante_forecasts.tex` with its hash.

The table makes three forecasts that can be checked after the announcements: the size of the move (the market’s straddle-implied move, the option-implied announcement standard deviation σ_J^Q , and the historical σ_J^P and Kou-SEJ-P move); the shape of the pre-event smile (the Kou-SEJ-Q minus market implied volatility at the 25-delta put and call); and, through the stored quotes, the prices of the longer expirations. A later version of this paper will report the realized moves and the post-event chains against these numbers. No claim about them is made here.

F Additional empirical results

This appendix collects the tables and figures of Section 8 whose content the main text summarizes: the errors in three other samples, the errors by the pre-registered cuts with four typical events, the tests of equal accuracy, the Model Confidence Set and the regressions, the post-event detail, the post-freeze exploratory American calibration, and the recent announcements by firm. The models, samples, metrics and conventions are those of Section 8.1; nothing here changes a price or a verdict.

F.1 Pre-event errors in other samples

Table 41 repeats the comparison of Table 8 in three other samples. In the broad sample, which includes 2020, the option-implied models again came first (Kou-SEJ-Q 4.69 and 4.65, DJ-implied 4.67 and 4.58), and so they did on the core quotes with a relative spread of at most 10%. On the out-of-calibration quotes, where DJ-implied and Kou-SEJ-Q extrapolate from the pair, Kou-SEJ-Q had the lower RMSE (4.38 against 4.47 in the holdout), but over 50 events only, and Mkt-ATM, which uses the quote’s own expiration, was lower still (3.46).

Table 40: Ex-ante model prices for announcements scheduled after the 12 October 2026 close, computed from that snapshot only. Event: date and timing (A after close, B before open, A* timing not yet announced, treated as after close). Exp./DTE: nearest expiration containing the event, the shorter of the two calibration expirations (the two shortest that contain the event, as in Section 8.10), and trading days to it; every other column refers to that expiration. ATM IV: the at-the-money implied volatility the calibration uses (strike nearest the spot); Mkt move: straddle-implied move, mid of the ATM straddle / spot at the strike nearest the spot whose call and put both pass the quote filters of Section 8.10 (– if no strike does); σ_J^Q : two-expiry implied announcement s.d.; σ_J^P : diffusion-corrected historical announcement s.d.; P move: the same straddle / spot under Kou-SEJ-P (– without a historical law); last two columns: Kou-SEJ-Q minus market IV at the 25-delta put and call among the quotes that pass the filters (vol. points). All in %.

Firm	Event	Spot	Exp./DTE	ATM IV	Mkt move	σ_J^Q	σ_J^P	P move	25 Δ P	25 Δ C
TSLA	10-21 A	378.25	10-26/10	46.7	7.4	6.0	9.5	9.7	0.5	0.5
NVDA	11-18 A*	226.67	11-20/29	34.3	9.3	4.4	7.4	9.0	–2.3	2.1
META	11-04 A*	742.38	11-06/19	44.3	9.7	8.1	12.4	14.1	0.1	0.3
NFLX	10-20 A	70.63	10-23/9	54.7	8.3	8.4	12.0	10.1	0.9	1.7
AMD	11-03 A	594.73	11-06/19	52.0	11.4	7.4	7.8	12.6	–1.0	0.5
AMZN	10-29 A*	258.99	11-06/19	39.2	8.6	6.7	6.3	7.2	0.4	0.2
GOOGL	10-28 A	352.61	11-06/19	36.2	8.0	6.1	5.1	6.5	0.5	0.4
MSFT	10-28 A	540.58	11-06/19	30.6	6.7	4.5	4.3	6.0	–0.9	0.9
AAPL	11-02 A	340.97	11-06/19	25.9	5.7	3.9	3.8	5.5	–1.3	1.7
PLTR	11-02 A*	208.99	11-06/19	55.3	12.1	8.3	15.3	15.1	–1.9	2.1
COIN	10-29 A*	180.16	11-06/19	65.7	–	10.2	12.2	–	–10.4	–
SMCI	11-03 A*	39.95	11-06/19	78.2	17.1	6.8	13.9	16.9	1.5	–2.8
INTC	10-29 A	101.36	11-06/19	60.4	13.2	8.9	9.3	16.3	0.3	–1.9
UBER	11-04 B	71.45	11-06/19	41.5	9.1	8.2	5.1	6.9	0.9	0.1
HOOD	10-27 A	109.26	11-06/19	62.9	13.7	9.1	9.4	15.8	0.8	–0.5
DIS	11-12 A*	108.57	11-20/29	31.5	8.7	7.6	6.5	8.2	4.3	2.0
BA	10-27 B	189.65	11-06/19	34.5	7.6	3.8	3.5	7.5	–0.6	–0.4
PYPL	10-27 B	56.09	11-06/19	42.9	10.2	7.6	10.4	10.3	0.8	–1.9
ROKU	10-29 A*	152.74	11-06/19	19.0	–	3.5	–	–	–	–
CVNA	10-28 A	59.56	11-06/19	77.0	16.9	12.9	–	–	–3.6	3.0
MRNA	11-05 A*	213.78	11-06/19	105.4	23.0	20.0	9.7	19.4	6.1	–5.1
ABNB	11-05 A	163.64	11-06/19	44.7	–	8.9	7.7	–	0.2	1.9
NOW	10-28 A	140.22	11-06/19	58.5	12.8	10.1	6.1	10.9	1.2	–0.8
TTD	11-05 A*	11.95	11-06/19	58.7	–	11.9	23.9	–	–	2.3

Variance against shape (post-freeze, exploratory). DJ-hist and Kou-SEJ-P differ in two ways: DJ-hist takes the mean squared past move as its announcement variance and Kou-SEJ-P the diffusion-corrected second moment var_P of (45), and Kou-SEJ-P adds the double-exponential shape. Black–Scholes with total variance $\sigma_{\text{HV}}^2 T + \text{var}_P$, a Gaussian announcement law with the variance of Kou-SEJ-P computed from the stored inputs and not one of the frozen models, separates the two. On the common sample its event-average RMSE in the holdout was 9.31 [9.02, 9.62], between 8.91 for DJ-hist and 9.40 for Kou-SEJ-P: the diffusion correction of the variance raised the RMSE by 0.40 [0.37, 0.43] points and the shape by a further 0.10 [0.07, 0.13]. In the development sample the two parts were 0.14 [0.06, 0.22] and –0.13 [–0.21, –0.06], so there the shape lowered the RMSE. With the pre-registered vega-scaled metric the parts were 0.27 and 0.11 in the holdout and –0.03 and –0.12 in the development sample. Most of the holdout gap between DJ-hist and Kou-SEJ-P is therefore the variance, not the shape, and the sign of the shape’s contribution differs between the two samples.

F.2 Errors by the pre-registered cuts

Tables 42 and 43 report the RMSE by the pre-registered cuts, and Figure 16 the mispricing captured in the same cells; Table 44 gives the RMSE by the same cuts with the two samples

Table 41: Pre-event RMSE in implied-volatility points (event averages) in three other samples; the second header row gives the announcement years each column covers. Broad: every model except DJ-hist and Kou-SEJ-P (which need four prior announcements) defined, so 2020 is included. Out-of-calibration: common-sample quotes whose expiration is not one of the two expirations of the decomposition pair, to which DJ-implied and Kou-SEJ-Q are fitted at the money; Mkt-ATM uses the ATM IV of each quote’s own expiration, out-of-calibration expirations included. Core: common-sample quotes with a relative spread of at most 10%. Out-of-calibration and core quotes are subsets of the common sample, whose development part has no 2020 event.

Model	Broad sample			Out-of-calibration quotes			Core quotes (spread $\leq 10\%$)		
	Dev. 2020–2021	Holdout 2022–2025	All 2020–2025	Dev. 2021	Holdout 2022–2025	All 2021–2025	Dev. 2021	Holdout 2022–2025	All 2021–2025
Events	3,623	5,974	9,597	4	50	54	1,240	5,044	6,284
Quotes	52,068	95,210	147,278	93	624	717	9,357	48,383	57,740
BS-HV	16.07	16.00	16.02	15.31	15.38	15.37	12.36	15.28	14.71
CRR-HV	15.43	14.85	15.07	15.22	14.82	14.85	11.62	14.16	13.65
BS-IVlong	8.00	8.89	8.55	7.06	5.92	6.01	6.45	8.32	7.95
Mkt-ATM	4.67	4.57	4.61	3.77	3.46	3.48	3.53	3.91	3.83
DJ-hist	–	–	–	5.61	8.07	7.89	8.70	8.49	8.53
DJ-implied	4.67	4.58	4.61	3.76	4.47	4.42	3.53	3.91	3.84
Kou-SEJ-P	–	–	–	8.11	8.84	8.78	8.69	9.05	8.98
Kou-SEJ-Q	4.69	4.65	4.67	3.62	4.38	4.32	3.77	4.09	4.03
Kou-SEJ-Q-prereg	5.34	5.51	5.45	3.78	4.92	4.84	3.95	4.98	4.78

pooled.

Moneyness. The option-implied models fitted to the at-the-money level (Mkt-ATM, DJ-implied and Kou-SEJ-Q) had their smallest errors for $0.35 \leq |\Delta| \leq 0.65$ (Kou-SEJ-Q 1.51 and 1.76), as they should. Kou-SEJ-Q-prereg, which takes the decomposition variance without refitting the level, had its smallest error in the wing bin (2.37 against 2.47 at the money in the development sample, 3.12 against 3.46 in the holdout). For $0.15 \leq |\Delta| < 0.35$ the double-exponential smile of Kou-SEJ-Q was worse than the flat volatility of DJ-implied (RMSE 2.42 against 2.01 in the development sample, 2.97 against 2.50 in the holdout), although its mean error there was smaller (-0.44 against -0.95 , and -0.26 against -0.87). Figure 15 shows four typical events, the two per sample whose Kou-SEJ-Q RMSE was closest to the sample median (CI 2021-11-04, ALGT 2021-07-28, MPWR 2023-10-30, LVS 2024-01-24); in the second and third panels the asymmetry taken from the firm’s history tilts the Kou-SEJ-Q smile upward to the right while the market smile falls.

Delta bins at an ex-ante volatility (post-freeze, exploratory). The pre-registered $|\Delta|$ bins use the delta at the quote’s own implied volatility, the volatility whose error is measured, so a quote with a high market volatility can be placed in the in-the-money bin by that volatility itself. With the delta taken instead at the at-the-money implied volatility of the quote’s expiration, the volatility of Mkt-ATM, 8.7% of the holdout quotes change bin and 7.1% fall outside $0.15 \leq |\Delta| \leq 0.85$ (development 8.5% and 7.0%). The in-the-money RMSE of Kou-SEJ-Q falls from 6.36 to 4.52 in the holdout and that of DJ-implied from 6.37 to 4.19, but the in-the-money bin keeps the largest errors (Kou-SEJ-Q 2.87 in the wing and 1.76 at the money; in-the-money minus at-the-money 2.77 [2.70, 2.83]), and so it does in the development sample (3.99 against 1.48 at the money). Part of the in-the-money excess of the pre-registered bins thus comes from the binning, and most of it remains.

Maturity and past moves. BS-HV’s error was largest for 5–15 days to expiration (27.48 and 28.45) and smallest for 31–45 days (10.23 and 11.50): the announcement is a larger share of the




Figure 15: Pre-event implied-volatility smiles of the nearest expiration for four typical events: per sample, the two events of distinct firms whose Kou-SEJ-Q RMSE is closest to the sample median, among events with an exact Kou-SEJ-Q calibration and at least twelve kept quotes (three calls and three puts) in that expiration. Market mids with bid–ask ranges (circles: calls; triangles: puts), the Kou-SEJ-Q, Kou-SEJ-P and Kou-SEJ-Q-prereg smiles, and the flat DJ-implied and BS-HV levels. Panel titles: announcement date and timing, days to expiration, realized move.

variance when little time remains. It also rose from the bottom to the top tercile of the firm’s past announcement moves (9.68 to 16.60, and 10.30 to 23.15), and so did the share captured by Kou-SEJ-Q (0.72 to 0.84, and 0.67 to 0.85): the variance that BS-HV omits was largest where past moves had been large. Where the option-implied announcement variance was low relative to the historical one (bottom tercile of their ratio), DJ-hist captured only 0.00 and 0.10 of BS-HV’s error: the firm’s history overstated what the market priced.

VIX, spreads, price level, timing and year. The errors of Kou-SEJ-Q were largest in the top VIX tercile (5.99 and 5.58, against 4.15 and 4.11 in the bottom one); with cut points from the whole development period, including the high-VIX months of 2020, the top tercile holds only 53 events of 2021. The top relative-spread tercile is small too (45 and 270 events), because the filters removed most quotes of wide-spread names (Section 7.4), and in it Kou-SEJ-Q’s error was below that of the middle tercile (4.55 against 4.89, and 4.80 against 4.94). In the holdout the lowest price tercile, the pre-registered size proxy, had the largest errors of every model. After-close announcements had larger BS-HV errors than before-open ones (13.69 against

Table 42: Pre-event RMSE in implied-volatility points (event averages) by the pre-registered cuts, development sample (announcements of 2021 only: no 2020 event has the four prior announcements that DJ-hist and Kou-SEJ-P need), common sample. Quote-level cuts ($|\Delta|$ at the quote's own IV, days to expiration) average the per-event RMSE over the event's quotes in the bin; event-level cuts use terciles fixed on the development sample (implied/historical variance: $\sigma_J^{2,Q}/\sigma_J^{2,\text{hist}}$). Events: events with quotes in the cell.

Cut	Events	BS-HV	CRR-HV	BS-IVlong	Mkt-ATM	DJ-hist	DJ-impl.	Kou-SEJ-P	Kou-SEJ-Q	Q-prereg
<i>Year</i>										
2021	1,386	12.84	12.16	6.98	4.04	8.99	4.04	9.00	4.20	4.44
<i>Timing</i>										
AMC	733	13.69	13.22	7.13	3.85	9.26	3.85	9.42	4.12	4.31
BMO	653	11.89	10.97	6.81	4.26	8.68	4.26	8.52	4.29	4.59
<i>Δ bin</i>										
0.15–0.35	959	11.47	11.46	5.37	2.01	7.72	2.01	7.63	2.42	2.37
0.35–0.65	1,347	10.80	10.74	4.60	1.37	7.57	1.37	7.29	1.51	2.47
0.65–0.85	1,378	14.25	13.14	8.51	5.57	9.88	5.57	10.12	5.71	5.77
<i>DTE bin (days)</i>										
5–15	72	27.48	27.27	17.48	5.23	14.21	5.23	13.85	5.94	6.25
16–30	1,119	13.14	12.44	7.36	3.95	9.27	3.95	9.23	4.14	4.42
31–45	393	10.23	9.69	4.01	3.84	7.19	3.84	7.46	3.84	3.98
<i>Past move tercile</i>										
T1	327	9.68	8.47	5.71	4.15	7.25	4.15	8.00	4.26	4.32
T2	619	11.84	11.21	6.41	3.98	7.95	3.98	8.34	4.05	4.25
T3	440	16.60	16.23	8.72	4.05	11.73	4.05	10.66	4.38	4.81
<i>Implied/hist. variance tercile</i>										
T1	425	10.77	10.17	5.69	3.76	10.56	3.76	8.85	3.87	4.06
T2	596	12.50	11.80	7.00	4.04	6.95	4.04	7.24	4.14	4.40
T3	365	15.81	15.06	8.44	4.38	10.47	4.38	12.04	4.69	4.95
<i>VIX tercile</i>										
T1	822	12.39	11.67	6.86	3.98	8.74	3.98	8.70	4.15	4.34
T2	511	12.55	11.98	6.86	3.99	8.89	3.99	8.83	4.11	4.43
T3	53	22.65	21.41	9.95	5.51	13.82	5.51	15.20	5.99	6.18
<i>Spread tercile</i>										
T1	921	11.98	11.31	6.60	3.66	8.68	3.66	8.60	3.87	4.00
T2	420	14.43	13.76	7.73	4.81	9.58	4.81	9.75	4.89	5.31
T3	45	15.62	14.55	7.87	4.74	9.78	4.74	10.08	4.55	5.46
<i>Price tercile</i>										
T1	181	14.68	14.25	8.00	4.39	11.38	4.38	10.83	4.55	4.82
T2	405	14.55	13.29	8.36	5.16	10.48	5.16	10.32	5.15	5.69
T3	800	11.56	11.11	6.05	3.40	7.69	3.40	7.91	3.65	3.73

11.89, and 17.10 against 14.09), while for the option-implied models the difference was small and changed sign. The option-implied models had their largest holdout errors in 2022.

Table 43: Pre-event RMSE in implied-volatility points (event averages) by the pre-registered cuts, holdout sample (2022–2025), common sample. Quote-level cuts ($|\Delta|$ at the quote’s own IV, days to expiration) average the per-event RMSE over the event’s quotes in the bin; event-level cuts use terciles fixed on the development sample (implied/historical variance: $\sigma_J^{2,Q}/\sigma_J^{2,\text{hist}}$). Events: events with quotes in the cell.

Cut	Events	BS-HV	CRR-HV	BS-IVlong	Mkt-ATM	DJ-hist	DJ-impl.	Kou-SEJ-P	Kou-SEJ-Q	Q-prereg
<i>Year</i>										
2022	1,582	15.06	14.09	8.59	5.39	10.44	5.39	10.91	5.44	6.26
2023	1,880	13.54	12.23	8.39	4.64	8.31	4.64	8.35	4.68	5.48
2024	1,561	17.79	16.67	9.47	3.64	8.08	3.65	9.04	3.83	4.69
2025	484	18.23	17.09	9.07	4.41	8.93	4.42	9.75	4.61	5.36
<i>Timing</i>										
AMC	2,745	17.10	15.95	9.41	4.61	9.79	4.61	10.37	4.77	5.65
BMO	2,762	14.09	12.96	8.23	4.49	8.03	4.50	8.45	4.53	5.29
<i>Δ bin</i>										
0.15–0.35	4,241	14.50	14.47	7.20	2.50	7.53	2.50	8.04	2.97	3.12
0.35–0.65	5,359	13.43	13.28	6.30	1.59	7.09	1.59	7.52	1.76	3.46
0.65–0.85	5,406	17.36	15.26	10.63	6.37	10.22	6.37	10.80	6.36	7.06
<i>DTE bin (days)</i>										
5–15	676	28.45	27.21	19.27	4.25	11.54	4.25	13.06	4.83	6.12
16–30	4,264	15.58	14.41	8.80	4.60	9.05	4.60	9.49	4.70	5.57
31–45	1,640	11.50	10.60	4.26	4.11	7.45	4.14	7.73	4.06	4.55
<i>Past move tercile</i>										
T1	1,389	10.30	8.70	6.75	4.94	7.82	4.94	8.38	4.84	5.26
T2	2,579	13.93	12.94	7.94	4.33	8.18	4.33	8.74	4.34	5.07
T3	1,539	23.15	22.18	12.14	4.58	11.13	4.59	11.44	4.99	6.32
<i>Implied/hist. variance tercile</i>										
T1	704	12.10	10.75	7.20	4.91	10.16	4.91	8.63	4.92	5.37
T2	2,685	14.27	13.23	8.04	4.24	6.87	4.24	7.28	4.29	5.05
T3	2,118	18.43	17.22	10.33	4.83	11.08	4.84	12.36	5.01	6.04
<i>VIX tercile</i>										
T1	2,398	15.83	14.63	8.78	4.01	7.94	4.02	8.39	4.11	4.91
T2	2,121	15.18	14.01	8.80	4.72	9.45	4.73	9.88	4.82	5.61
T3	988	15.89	14.95	8.92	5.50	10.09	5.50	10.84	5.58	6.51
<i>Spread tercile</i>										
T1	3,637	15.17	14.01	8.64	4.42	8.97	4.43	9.39	4.51	5.36
T2	1,600	16.24	15.12	9.23	4.81	8.69	4.82	9.29	4.94	5.66
T3	270	17.41	16.39	8.64	4.73	9.44	4.77	10.30	4.80	5.78
<i>Price tercile</i>										
T1	874	20.87	19.20	11.74	6.05	12.33	6.06	13.00	6.15	7.51
T2	1,607	16.17	14.99	9.20	5.00	9.59	5.00	10.05	5.00	5.86
T3	3,026	13.76	12.79	7.76	3.88	7.56	3.88	8.02	4.03	4.67

Table 44: Pre-event RMSE in implied-volatility points (event averages) by the pre-registered cuts, both samples (2021–2025), common sample. Quote-level cuts ($|\Delta|$ at the quote’s own IV, days to expiration) average the per-event RMSE over the event’s quotes in the bin; event-level cuts use terciles fixed on the development sample (implied/historical variance: $\sigma_J^{2,Q}/\sigma_J^{2,\text{hist}}$). Events: events with quotes in the cell.

Cut	Events	BS-HV	CRR-HV	BS-IVlong	Mkt-ATM	DJ-hist	DJ-impl.	Kou-SEJ-P	Kou-SEJ-Q	Q-prereg
<i>Year</i>										
2021	1,386	12.84	12.16	6.98	4.04	8.99	4.04	9.00	4.20	4.44
2022	1,582	15.06	14.09	8.59	5.39	10.44	5.39	10.91	5.44	6.26
2023	1,880	13.54	12.23	8.39	4.64	8.31	4.64	8.35	4.68	5.48
2024	1,561	17.79	16.67	9.47	3.64	8.08	3.65	9.04	3.83	4.69
2025	484	18.23	17.09	9.07	4.41	8.93	4.42	9.75	4.61	5.36
<i>Timing</i>										
AMC	3,478	16.38	15.37	8.93	4.45	9.68	4.45	10.17	4.64	5.37
BMO	3,415	13.67	12.58	7.96	4.45	8.16	4.46	8.46	4.48	5.16
<i>Δ bin</i>										
0.15–0.35	5,200	13.94	13.92	6.86	2.41	7.57	2.41	7.97	2.87	2.98
0.35–0.65	6,706	12.90	12.77	5.95	1.54	7.19	1.55	7.47	1.71	3.27
0.65–0.85	6,784	16.73	14.83	10.20	6.21	10.15	6.21	10.66	6.23	6.79
<i>DTE bin (days)</i>										
5–15	748	28.35	27.22	19.10	4.34	11.80	4.34	13.14	4.94	6.13
16–30	5,383	15.07	14.00	8.50	4.46	9.10	4.46	9.44	4.59	5.33
31–45	2,033	11.26	10.42	4.21	4.06	7.40	4.08	7.68	4.02	4.44
<i>Past move tercile</i>										
T1	1,716	10.18	8.66	6.55	4.79	7.71	4.79	8.31	4.73	5.08
T2	3,198	13.53	12.60	7.65	4.26	8.13	4.26	8.67	4.28	4.91
T3	1,979	21.69	20.85	11.38	4.46	11.26	4.47	11.26	4.86	5.98
<i>Implied/hist. variance tercile</i>										
T1	1,129	11.59	10.54	6.63	4.48	10.31	4.48	8.71	4.52	4.88
T2	3,281	13.95	12.97	7.85	4.20	6.89	4.20	7.27	4.27	4.93
T3	2,483	18.04	16.91	10.06	4.77	10.99	4.77	12.31	4.96	5.88
<i>VIX tercile</i>										
T1	3,220	14.95	13.87	8.29	4.00	8.15	4.01	8.47	4.12	4.77
T2	2,632	14.67	13.62	8.42	4.58	9.34	4.58	9.67	4.69	5.38
T3	1,041	16.23	15.28	8.97	5.50	10.28	5.50	11.06	5.60	6.49
<i>Spread tercile</i>										
T1	4,558	14.53	13.46	8.23	4.27	8.91	4.27	9.23	4.38	5.09
T2	2,020	15.86	14.84	8.92	4.81	8.88	4.82	9.39	4.93	5.59
T3	315	17.16	16.13	8.53	4.73	9.49	4.77	10.27	4.77	5.73
<i>Price tercile</i>										
T1	1,055	19.81	18.35	11.10	5.76	12.17	5.77	12.63	5.88	7.05
T2	2,012	15.84	14.64	9.03	5.03	9.77	5.03	10.11	5.03	5.83
T3	3,826	13.30	12.44	7.40	3.78	7.59	3.78	8.00	3.95	4.47




Figure 16: Mispricing captured relative to BS-HV by pre-registered cut, with 90% event-bootstrap intervals, for the four event-jump models; common sample, (a) development (announcements of 2021), (b) holdout (2022–2025). Quote-level cuts ($|\Delta|$, days to expiration) use the quotes in the bin; event-level cuts use terciles fixed on the development sample.

F.3 Tests of equal accuracy, the Model Confidence Set and the regressions

Table 45 reports the tests of equal accuracy discussed in Section 8.5: in Panel A the Diebold–Mariano and Giacomini–White tests with the implied-volatility loss, in Panel B the Mincer–Zarnowitz regressions, and in Panel C the tests with the pre-registered vega-scaled loss.

Table 46 reports the 90% Model Confidence Set of Hansen et al. (2011) (10% level). With the implied-volatility loss it was {Mkt-ATM, DJ-implied, Kou-SEJ-Q, Kou-SEJ-Q-prereg} in the development sample and {Mkt-ATM, DJ-implied, Kou-SEJ-Q} in the holdout, and blocks of one, two and eight weeks instead of 4 gave the same sets. With the vega-scaled loss it shrank to {Mkt-ATM, DJ-implied} and {Mkt-ATM}: Kou-SEJ-Q left the set, as the Diebold–Mariano tests of Section 8.5 anticipate. Under either loss the set held only models that use the at-the-money quotes, and always Mkt-ATM.

Table 47 regresses the per-event loss differentials on the pre-registered characteristics, with standard errors clustered by firm and calendar quarter (Cameron et al., 2011); the loss is the squared implied-volatility difference, not the pre-registered vega-scaled metric. The development sample has only 4 quarter clusters and the holdout 14, and the VIX terciles vary mostly between quarters, so I use the development regressions for no inference and check the holdout ones in Table 48, a post-freeze sensitivity check that was not pre-registered: the key coefficients with the per-event RMSE and with the vega-scaled loss, with firm-only clustered errors and with quarter fixed effects, and a wild cluster restricted bootstrap by quarter. In the holdout the gain of Kou-SEJ-Q over BS-HV was larger in the top past-move tercile (-692.24 with standard error 120.94) and in the top tercile of the ratio of option-implied to historical announcement variance, where the variance that BS-HV omits was largest, with wild bootstrap p -values below 0.01 under every loss. The disadvantage of Kou-SEJ-Q relative to DJ-implied was concentrated in the same top past-move tercile under every loss and every kind of inference (two-way t from 4.97 to 5.56 , wild bootstrap p -value < 0.001): the shape taken from history hurt most for the firms with large past moves. A top-VIX-tercile coefficient appeared only with the squared implied-volatility loss (3.88 , two-way $t = 3.30$ and firm-only $t = 2.60$, against the 5% critical value 2.16 of the t distribution with one degree of freedom fewer than the 14 quarters; wild bootstrap p -value 0.029). It did not appear with the per-event RMSE ($t = 0.12$) or with the vega-scaled loss ($t = 1.30$ squared, -1.63 as an RMSE), nor with quarter fixed effects ($t = 1.41$), so I do not read it as a finding. R^2 was at most 0.052 .

Table 45: Tests of equal predictive accuracy on the common sample of Table 8. The loss of an event is the mean squared IV error of its quotes; $\bar{d} < 0$ favors the first model. Diebold–Mariano t (Diebold and Mariano, 1995): mean differential over its standard error clustered by announcement week; † marks $|t| > 3$, the hurdle of Harvey et al. (2016). Giacomini–White (Giacomini and White, 2006): conditional test with instruments (1, VIX at d_{pre} , mean differential of the previous announcement week), χ^2_3 p -value. ‡ marks a Holm-adjusted p -value below 0.05, the adjustment of the plan over the comparisons of a panel within a split (adjusted values in `output/results_summary.json`). Panel C repeats Panel A with the pre-registered vega-scaled loss. The Mincer–Zarnowitz slopes of Mkt-ATM, DJ-implied and Kou-SEJ-Q are close to one largely by construction, since these models are fitted to the at-the-money level of the expirations of most quotes. Common sample and conventions as in Table 8.

Comparison	Development 2021			Holdout 2022–2025			All 2021–2025		
	$\bar{d} / a$	t_{DM} / b	p_{GW} / R^2	$\bar{d} / a$	t_{DM} / b	p_{GW} / R^2	$\bar{d} / a$	t_{DM} / b	p_{GW} / R^2
<i>A. Loss differential $d = L_{\text{model}} - L_{\text{bench}}$ (vol. points²), Diebold–Mariano t (announcement-week clusters) and Giacomini–White p</i>									
CRR-HV vs BS-HV	−27.06	−5.94 ^{†‡}	0.005	−48.00	−17.99 ^{†‡}	0.000	−43.79	−17.78 ^{†‡}	0.000
BS-IVlong vs BS-HV	−192.45	−5.76 ^{†‡}	0.001	−358.25	−4.66 ^{†‡}	0.000	−324.91	−5.21 ^{†‡}	0.000
Mkt-ATM vs BS-HV	−236.29	−6.37 ^{†‡}	0.000	−439.63	−5.69 ^{†‡}	0.000	−398.74	−6.32 ^{†‡}	0.000
DJ-hist vs BS-HV	−102.07	−3.96 ^{†‡}	0.012	−230.09	−11.27 ^{†‡}	0.000	−204.35	−11.60 ^{†‡}	0.000
DJ-impl. vs BS-HV	−236.29	−6.37 ^{†‡}	0.000	−439.56	−5.69 ^{†‡}	0.000	−398.69	−6.32 ^{†‡}	0.000
Kou-SEJ-P vs BS-HV	−106.93	−5.13 ^{†‡}	0.004	−217.58	−11.84 ^{†‡}	0.000	−195.33	−12.45 ^{†‡}	0.000
Kou-SEJ-Q vs BS-HV	−235.56	−6.42 ^{†‡}	0.000	−440.15	−5.70 ^{†‡}	0.000	−399.01	−6.33 ^{†‡}	0.000
Q-prereg vs BS-HV	−233.69	−6.46 ^{†‡}	0.000	−428.31	−5.55 ^{†‡}	0.000	−389.18	−6.18 ^{†‡}	0.000
Kou-SEJ-Q vs DJ-impl.	0.73	0.99	0.350	−0.59	−0.97	0.247	−0.33	−0.64	0.105
Kou-SEJ-Q vs Mkt-ATM	0.73	0.98	0.351	−0.52	−0.84	0.228	−0.27	−0.53	0.092
Kou-SEJ-Q vs Kou-SEJ-P	−128.63	−4.89 ^{†‡}	0.000	−222.57	−2.80 [‡]	0.004	−203.68	−3.18 ^{†‡}	0.000
DJ-impl. vs DJ-hist	−134.23	−4.70 ^{†‡}	0.000	−209.47	−2.61 [‡]	0.009	−194.34	−3.01 ^{†‡}	0.000
Kou-SEJ-Q vs BS-IVlong	−43.11	−9.50 ^{†‡}	0.000	−81.90	−11.23 ^{†‡}	0.000	−74.10	−12.19 ^{†‡}	0.000
Q-prereg vs DJ-impl.	2.60	1.36	0.543	11.25	13.97 ^{†‡}	0.000	9.51	11.43 ^{†‡}	0.000
Q-prereg vs Kou-SEJ-Q	1.87	1.15	0.462	11.84	12.17 ^{†‡}	0.000	9.83	10.69 ^{†‡}	0.000
Kou-SEJ-Q vs DJ-impl. ($ \Delta < 0.35$)	3.91	4.47 ^{†‡}	0.007	4.68	8.68 ^{†‡}	0.000	4.54	9.65 ^{†‡}	0.000
Kou-SEJ-Q vs DJ-impl. (out of calib.)	−1.68	–	–	−1.16	−1.14	0.500	−1.20	−1.26	0.426
<i>B. Mincer–Zarnowitz regression of the market IV on the model IV (quote level, vol. points): intercept, slope, R^2</i>									
BS-HV	24.65	0.607	0.536	35.16	0.406	0.215	33.21	0.444	0.261
CRR-HV	23.86	0.621	0.565	33.92	0.429	0.241	32.02	0.466	0.289
BS-IVlong	2.36	1.063	0.854	3.00	1.075	0.804	2.85	1.074	0.813
Mkt-ATM	1.73	0.996	0.914	2.35	0.982	0.909	2.24	0.985	0.910
DJ-hist	18.88	0.574	0.665	26.34	0.489	0.384	24.90	0.506	0.429
DJ-implied	1.75	0.995	0.914	2.35	0.982	0.909	2.24	0.984	0.910
Kou-SEJ-P	20.05	0.584	0.646	27.54	0.486	0.366	26.07	0.506	0.411
Kou-SEJ-Q	2.28	0.975	0.907	2.85	0.960	0.907	2.76	0.962	0.907
Kou-SEJ-Q-prereg	2.80	0.999	0.917	1.46	1.052	0.908	1.71	1.043	0.910
<i>C. Panel A with the pre-registered vega-scaled loss (squared price error over the BS vega at the quote’s own IV): $\bar{d}$, t_{DM}, p_{GW}</i>									
CRR-HV vs BS-HV	−10.03	−5.09 ^{†‡}	0.005	−16.52	−19.63 ^{†‡}	0.000	−15.21	−18.45 ^{†‡}	0.000
BS-IVlong vs BS-HV	−152.52	−5.34 ^{†‡}	0.001	−320.65	−3.62 ^{†‡}	0.000	−286.85	−4.01 ^{†‡}	0.000
Mkt-ATM vs BS-HV	−187.09	−5.96 ^{†‡}	0.000	−386.27	−4.38 ^{†‡}	0.000	−346.22	−4.83 ^{†‡}	0.000
DJ-hist vs BS-HV	−46.77	−2.17	0.054	−157.71	−10.39 ^{†‡}	0.000	−135.40	−10.13 ^{†‡}	0.000
DJ-impl. vs BS-HV	−187.10	−5.96 ^{†‡}	0.000	−386.21	−4.37 ^{†‡}	0.000	−346.17	−4.83 ^{†‡}	0.000
Kou-SEJ-P vs BS-HV	−58.53	−3.58 ^{†‡}	0.024	−149.37	−11.34 ^{†‡}	0.000	−131.10	−11.48 ^{†‡}	0.000
Kou-SEJ-Q vs BS-HV	−184.54	−5.92 ^{†‡}	0.000	−384.06	−4.35 ^{†‡}	0.000	−343.94	−4.80 ^{†‡}	0.000
Q-prereg vs BS-HV	−183.24	−5.92 ^{†‡}	0.000	−376.22	−4.26 ^{†‡}	0.000	−337.42	−4.72 ^{†‡}	0.000
Kou-SEJ-Q vs DJ-impl.	2.56	4.07 ^{†‡}	0.003	2.14	6.08 ^{†‡}	0.000	2.23	7.26 ^{†‡}	0.000
Kou-SEJ-Q vs Mkt-ATM	2.55	4.07 ^{†‡}	0.004	2.21	6.01 ^{†‡}	0.000	2.28	7.17 ^{†‡}	0.000
Kou-SEJ-Q vs Kou-SEJ-P	−126.01	−4.76 ^{†‡}	0.000	−234.69	−2.57 [‡]	0.017	−212.84	−2.90 [‡]	0.001
DJ-impl. vs DJ-hist	−140.33	−4.77 ^{†‡}	0.001	−228.50	−2.49 [‡]	0.023	−210.77	−2.85 [‡]	0.001
Kou-SEJ-Q vs BS-IVlong	−32.02	−8.77 ^{†‡}	0.000	−63.41	−10.56 ^{†‡}	0.000	−57.10	−11.43 ^{†‡}	0.000
Q-prereg vs DJ-impl.	3.86	5.27 ^{†‡}	0.004	9.99	15.48 ^{†‡}	0.000	8.75	15.12 ^{†‡}	0.000
Q-prereg vs Kou-SEJ-Q	1.30	1.37	0.287	7.84	12.27 ^{†‡}	0.000	6.53	11.32 ^{†‡}	0.000

Table 46: The 90% Model Confidence Set (10% level) (Hansen et al., 2011). Loss: event average of the per-event mean squared IV error (vol. points²). $T_{\max}$ statistic, elimination of the model with the largest standardized excess loss, MCS p -values as running maxima, 1,000 moving-block bootstrap draws of four consecutive announcement weeks. In: the model belongs to the 90% MCS; In_v: the same with the pre-registered vega-scaled loss (squared price error over the BS vega at the quote’s own IV). Common sample and conventions as in Table 8.

Model	Development 2021				Holdout 2022–2025				All 2021–2025			
	Loss	p_{MCS}	In	In _v	Loss	p_{MCS}	In	In _v	Loss	p_{MCS}	In	In _v
BS-HV	273.6	0.003			482.3	0.000			440.3	0.000		
CRR-HV	246.5	0.007			434.3	0.000			396.5	0.000		
BS-IVlong	81.1	0.016			124.0	0.033			115.4	0.012		
Mkt-ATM	37.3	0.382	✓	✓	42.6	0.412	✓	✓	41.6	0.623	✓	✓
DJ-hist	171.5	0.016			252.2	0.033			235.9	0.012		
DJ-implied	37.3	1.000	✓	✓	42.7	0.293	✓		41.6	0.483	✓	
Kou-SEJ-P	166.7	0.012			264.7	0.023			245.0	0.008		
Kou-SEJ-Q	38.0	0.382	✓		42.1	1.000	✓		41.3	1.000	✓	
Kou-SEJ-Q-prereg	39.9	0.273	✓		53.9	0.033			51.1	0.012		
Events / weeks	1,386 / 48				5,507 / 167				6,893 / 215			

Table 47: Regressions of per-event loss differentials (mean squared IV error, IV points²; not the pre-registered vega-scaled metric) on the pre-registered characteristics, common sample. Q – BS-HV: Kou-SEJ-Q minus BS-HV; Q – DJ-impl.: Kou-SEJ-Q minus DJ-implied; negative coefficients favor Kou-SEJ-Q. Tercile dummies (T1 omitted) use the development-sample cut points; wing and deep shares are the event’s shares of quotes with $|\Delta|$ in 0.15–0.35 and 0.65–0.85. Standard errors in parentheses, two-way clustered by firm and calendar quarter (Cameron et al., 2011). The development sample has only 4 quarter clusters (the holdout 14), so its standard errors are not used for inference; Table 48 gives, for the key coefficients, wild-cluster bootstrap p -values and the same regressions with other losses.

Regressor	Development 2021		Holdout 2022–2025		All 2021–2025	
	Q – BS-HV	Q – DJ-impl.	Q – BS-HV	Q – DJ-impl.	Q – BS-HV	Q – DJ-impl.
Intercept	–369.45 (102.72)	4.60 (6.58)	–144.61 (378.15)	–6.27 (3.57)	–198.58 (297.33)	–3.38 (3.41)
BMO	81.71 (32.68)	–4.09 (1.27)	255.11 (175.97)	–0.06 (1.39)	220.31 (141.41)	–0.89 (1.24)
Past move T2	–120.86 (22.61)	–1.29 (1.60)	–350.99 (220.35)	2.53 (1.73)	–314.86 (181.99)	1.76 (1.52)
Past move T3	–376.23 (64.66)	4.38 (1.76)	–692.24 (120.94)	9.80 (1.97)	–634.57 (105.79)	8.69 (1.81)
Var. ratio T2	–128.57 (50.93)	–0.21 (1.97)	–192.82 (80.83)	2.47 (2.29)	–193.93 (65.60)	1.23 (1.65)
Var. ratio T3	–255.72 (58.23)	1.17 (2.88)	–522.42 (182.95)	4.46 (1.98)	–499.36 (155.68)	3.13 (1.47)
VIX T2	32.29 (38.34)	1.11 (0.75)	–195.52 (236.08)	2.32 (0.57)	–149.66 (186.54)	1.79 (0.40)
VIX T3	–251.61 (212.28)	13.24 (3.87)	–87.56 (110.83)	3.88 (1.18)	–98.06 (99.33)	3.68 (0.94)
Spread T2	–9.90 (42.75)	–1.96 (3.42)	102.72 (111.74)	2.14 (1.46)	83.69 (96.32)	1.32 (1.38)
Spread T3	–7.14 (71.56)	–8.59 (6.35)	54.76 (98.82)	3.34 (2.66)	34.43 (84.24)	1.24 (2.75)
Price T2	16.32 (66.40)	–3.31 (2.60)	78.61 (126.43)	0.98 (3.93)	71.05 (107.11)	0.20 (3.28)
Price T3	47.12 (85.28)	0.32 (2.54)	90.57 (165.70)	5.37 (4.21)	89.55 (128.21)	4.63 (3.51)
Wing share	110.96 (305.86)	7.03 (15.42)	–674.71 (645.62)	–2.43 (4.27)	–546.09 (523.41)	–0.87 (4.62)
Deep share	113.43 (246.19)	–7.00 (9.63)	290.99 (174.16)	–5.91 (5.76)	255.88 (146.89)	–5.86 (5.16)
Mean DTE (10 days)	110.08 (39.22)	–0.07 (1.55)	83.10 (64.16)	–1.51 (0.63)	97.58 (50.59)	–1.25 (0.65)
Mean of d	–235.56	0.73	–440.15	–0.59	–399.01	–0.33
Events	1,386	1,386	5,507	5,507	6,893	6,893
Firms / quarters	636 / 4	636 / 4	940 / 14	940 / 14	1016 / 18	1016 / 18
R^2	0.052	0.021	0.005	0.012	0.005	0.011

Table 48: Post-freeze sensitivity check of Table 47 (not pre-registered): the key coefficients of the loss-differential regressions with four per-event losses (IV difference or the pre-registered vega-scaled error, as the mean squared error or as the RMSE of the event) and four kinds of inference. t_{2w} : two-way (firm, quarter) clustered, as in Table 47; t_{firm} : clustered by firm only; t_{FE} : with calendar-quarter fixed effects and standard errors clustered by firm, which absorb most of the variation of the VIX terciles and so test a within-quarter effect; p_{WCR} : wild cluster restricted bootstrap p -value of the two-way t (999 draws of Webb’s six-point weights by calendar quarter, the dimension with fewer clusters: 4 in the development sample, 14 in the holdout). With 4 clusters the development p -values are coarse, and no development standard error is used for inference.

Regressor	Loss	Development 2021		Holdout 2022–2025				
		Coef.	p_{WCR}	Coef.	t_{2w}	t_{firm}	t_{FE}	p_{WCR}
<i>A. Kou-SEJ-Q minus DJ-implied</i>								
Past move T3	IV, squared	4.38	0.050	9.80	4.97	4.80	4.66	<0.001
	IV, RMSE	0.41	0.040	0.58	5.05	6.13	5.89	<0.001
	vega, squared	3.53	0.047	7.64	5.56	5.57	5.40	<0.001
	vega, RMSE	0.40	0.038	0.60	5.37	6.51	6.25	<0.001
VIX T2	IV, squared	1.11	0.216	2.32	4.09	1.96	1.31	<0.001
	IV, RMSE	−0.06	0.454	0.02	0.51	0.50	0.58	0.629
	vega, squared	−0.29	0.708	0.65	0.94	1.02	0.66	0.349
	vega, RMSE	−0.08	0.401	−0.01	−0.27	−0.35	−0.14	0.790
VIX T3	IV, squared	13.24	0.163	3.88	3.30	2.60	1.41	0.029
	IV, RMSE	0.25	0.269	0.00	0.12	0.08	0.50	0.907
	vega, squared	12.24	0.272	1.13	1.30	1.12	0.58	0.263
	vega, RMSE	0.27	0.259	−0.07	−1.63	−1.30	−0.27	0.141
<i>B. Kou-SEJ-Q minus BS-HV</i>								
Past move T3	IV, squared	−376.23	0.042	−692.24	−5.72	−5.74	−5.19	<0.001
	IV, RMSE	−8.53	0.049	−13.53	−13.35	−22.70	−21.65	<0.001
	vega, squared	−295.04	0.039	−576.49	−3.96	−3.96	−3.66	<0.001
	vega, RMSE	−7.45	0.048	−11.67	−13.73	−20.22	−19.10	<0.001
Var. ratio T3	IV, squared	−255.72	0.053	−522.42	−2.86	−3.02	−2.26	0.001
	IV, RMSE	−6.12	0.059	−8.76	−10.20	−16.14	−12.12	<0.001
	vega, squared	−179.73	0.046	−447.77	−2.21	−2.37	−1.89	0.003
	vega, RMSE	−4.98	0.059	−7.40	−10.18	−15.14	−11.04	<0.001

F.4 After the announcement

Table 49 reports the post-event errors summarized in Section 8.6. Mkt-ATM uses the post-event at-the-money volatility and is a reference, not a forecast.

Table 50 reports the change in implied volatility across the announcement. At the money the market’s implied volatility fell by 6.69 points on average in the development sample and by 8.35 in the holdout. DJ-implied and Kou-SEJ-Q predicted larger falls (8.00 and 11.09; 8.65 and 11.93): their background, which is what they predict for the option after the announcement, was below the post-event market level. Kou-SEJ-Q-prereg came closest on average, but only because its pre-event at-the-money volatility was too low. Every model with an announcement component had a Mincer–Zarnowitz slope below one (Kou-SEJ-Q 0.61 and 0.63): the predicted crushes were more dispersed across events than the realized ones.

Table 49: Post-event pricing errors: event-average RMSE in implied-volatility points at d_{post} of the contracts quoted at d_{pre} (quotes passing the filters at d_{post}), each model priced with its background diffusion only, and the reduction relative to BS-HV with 90% event-bootstrap interval; Red. vega: the reduction with the pre-registered vega-scaled error (price error over the BS vega at the post-event quote’s own IV), the statistic of P1. DJ-hist and Kou-SEJ-P share the BS-HV background, so their post-event prices equal BS-HV’s. Mkt-ATM at d_{post} uses the post-event ATM IV (information of d_{post}): it is a reference, not a forecast. The development post-event sample, like the common sample, has no 2020 event.

Model	Development 2021			Holdout 2022–2025			All 2021–2025		
	RMSE	Red. vs BS-HV (%)	Red. vega (%)	RMSE	Red. vs BS-HV (%)	Red. vega (%)	RMSE	Red. vs BS-HV (%)	Red. vega (%)
Events	932			3,986			4,918		
Quotes	6,888			32,325			39,213		
BS-HV	7.14	–	–	7.77	–	–	7.65	–	–
CRR-HV	6.66	6.8 [5.3, 8.3]	4.3	6.96	10.4 [9.5, 11.4]	5.7	6.90	9.8 [9.0, 10.7]	5.4
BS-IVlong	6.29	12.0 [6.8, 17.0]	6.4	7.78	–0.2 [–5.2, 4.6]	–7.5	7.50	2.0 [–2.2, 6.2]	–5.0
Mkt-ATM	3.18	55.5 [52.5, 58.4]	57.2	3.95	49.2 [46.6, 51.8]	51.8	3.80	50.3 [48.1, 52.5]	52.8
DJ-hist	7.14	0.0 [0.0, 0.0]	0.0	7.77	0.0 [0.0, 0.0]	0.0	7.65	0.0 [0.0, 0.0]	0.0
DJ-implied	4.80	32.9 [29.2, 36.5]	34.1	6.02	22.4 [18.6, 26.1]	25.3	5.79	24.3 [21.0, 27.5]	26.9
Kou-SEJ-P	7.14	0.0 [0.0, 0.0]	0.0	7.77	0.0 [0.0, 0.0]	0.0	7.65	0.0 [0.0, 0.0]	0.0
Kou-SEJ-Q	5.07	29.0 [25.2, 32.7]	30.6	6.42	17.3 [13.3, 21.2]	20.6	6.17	19.4 [15.9, 22.8]	22.4
Kou-SEJ-Q-prereg	4.80	32.9 [29.2, 36.5]	34.1	6.02	22.4 [18.6, 26.1]	25.3	5.79	24.3 [21.0, 27.5]	26.9

Table 50: The pre-to-post change in implied volatility. Change: event-average RMSE of the predicted minus the market change in IV from d_{pre} to d_{post} , quotes passing the filters on both dates. Crush at the money: for each event and expiration the contract closest to the forward; market crush $IV_{\text{pre}} - IV_{\text{post}}$ (event average) against the crush each model predicts (its pre-event IV minus its post-event IV); RMSE: event average of the per-event RMSE of predicted minus market crush; MZ: slope of the regression of the market crush on the predicted crush (announcement-week clusters; – for the models without an announcement component, whose predicted crush is close to zero). All in vol. points. Mkt-ATM uses the post-event ATM IV and is a reference, not a forecast.

Model	Development 2021				Holdout 2022–2025				All 2021–2025			
	Change	Crush	RMSE	MZ	Change	Crush	RMSE	MZ	Change	Crush	RMSE	MZ
Market crush		6.69				8.35				8.04		
Events	932	932			3,986	3,986			4,918	4,918		
BS-HV	7.98	0.00	7.50	–	9.96	0.00	9.29	–	9.58	0.00	8.95	–
CRR-HV	7.77	0.00	7.40	–	9.52	–0.17	9.16	–	9.19	–0.14	8.83	–
BS-IVlong	8.42	–0.52	7.94	–	10.36	–0.48	9.70	–	9.99	–0.49	9.37	–
Mkt-ATM	3.03	7.41	2.20	0.95	3.60	9.36	2.95	0.89	3.49	8.99	2.81	0.90
DJ-hist	6.22	10.12	5.85	0.39	5.73	9.70	5.34	0.65	5.82	9.78	5.44	0.59
DJ-implied	4.27	8.00	3.71	0.66	5.03	11.09	4.70	0.68	4.89	10.51	4.51	0.68
Kou-SEJ-P	5.66	7.49	5.09	0.43	5.69	7.59	5.14	0.74	5.69	7.57	5.13	0.67
Kou-SEJ-Q	4.65	8.65	4.03	0.61	5.65	11.93	5.23	0.63	5.46	11.31	5.00	0.63
Kou-SEJ-Q-prereg	4.23	6.56	3.60	0.73	4.84	8.65	4.19	0.76	4.73	8.26	4.08	0.76

F.5 Post-freeze exploratory analysis: a consistent American calibration

This analysis was not pre-registered. I ran it after the holdout results of Section 8 had been computed, it changes no frozen price and no verdict, and I report it because it tests the explanation of the put result given in Section 8.7. On the events of the American subsample I recalibrated Kou-SEJ-Q so that its American prices, rather than its European ones, match the at-the-money mids of the two decomposition expirations (Kou-SEJ-Q-AmCal). The announcement shape is that of Kou-SEJ-Q; the background volatility and the announcement variance solve the same least-squares problem in implied-volatility units, with model and market prices inverted in the same way; and the American price is the European price plus the finite-difference premium on the grid of Kou-SEJ-Q-Am (400×200 ; on 8 events a 1600×800 grid moved no quote’s implied volatility by more than 0.0019 points). Every kept quote of these events was then priced American, and the comparison uses the quotes of Table 10. The fit was exact to 10^{-4} in 541 of the 545 holdout events; in the other 4, all fitted exactly by the frozen calibration, the least-squares solution set the background volatility to zero and still missed the mids by up to 0.90 points. As an alternative, Kou-SEJ-Q-deAm calibrates the European model to at-the-money implied volatilities de-Americanized with the binomial American Black–Scholes model (500 steps) and is then priced American in the same way. Table 51 reports the errors.

At the calibration strikes of expirations without a dividend the double count was real and small: in the holdout the frozen American model over-priced those puts by 0.30 [0.23, 0.37] points on average, and the consistent calibration brought the mean errors of the puts and calls there to 0.09 and -0.08 . Where a dividend falls in the option’s life, every calibration missed those strikes in opposite directions (Kou-SEJ-Q-AmCal: puts 1.49, calls -1.31): the continuous dividend yield of the models does not reproduce the split between call and put prices there, and the column ‘Calib.’ of Table 51 pools both cases. Over all holdout puts it lowered the mean error from 1.15 to 0.95 (difference -0.20 [-0.22 , -0.19]), which removes 25% [23, 28] of the rise from the European 0.35, and it lowered the put RMSE by 1.8% [1.4, 2.1]. The rise was largest in the money: the frozen premium raised the mean put error by 1.58 points for $0.65 < |\Delta| \leq 0.85$, against 0.42 at the money, and the recalibration took back 11% of the

Table 51: Post-freeze exploratory analysis, not pre-registered: a consistent American calibration. Errors on the common-sample quotes of the stratified 10% American subsample (142 development and 545 holdout events), by $|\Delta|$ bin (Black–Scholes delta at the quote’s own implied volatility), for puts and calls. Kou-SEJ-Q (European) and Kou-SEJ-Q-Am (frozen) are the frozen prices: Kou-SEJ-Q-Am adds the finite-difference early-exercise premium to the European model calibrated to European-inverted IVs of the American at-the-money quotes. Kou-SEJ-Q-AmCal has the same announcement shape, with background volatility and announcement variance chosen so that its American at-the-money prices match the market’s (both inverted to IVs in the same way). Kou-SEJ-Q-deAm is the European model calibrated to at-the-money IVs de-Americanized with the binomial American Black–Scholes model, priced American. Errors are model minus market implied volatility in volatility points (European inversions of model and market prices at the quote’s own S, r, q, T); per-event RMSE and mean error averaged over events. ‘Calib.’ holds the quotes at the at-the-money strike of each of the two calibration expirations, to whose mids the three calibrated models are fitted (where those quotes passed the filters); ‘All’ covers every $|\Delta|$ bin. ‘Quotes’ rows give the number of quotes.

$ \Delta $	Development 2021					Holdout 2022–2025				
	0.15–0.35	0.35–0.65	0.65–0.85	Calib.	All	0.15–0.35	0.35–0.65	0.65–0.85	Calib.	All
<i>A. RMSE (IV points, event average)</i>										
Puts: quotes	178	426	408	147	1,012	1,044	1,956	1,664	659	4,664
Kou-SEJ-Q (European)	2.13	1.39	3.98	0.69	3.11	2.86	1.74	4.91	0.88	3.56
Kou-SEJ-Q-Am (frozen)	2.13	1.39	3.98	0.69	3.10	2.79	1.91	5.10	0.98	3.65
Kou-SEJ-Q-AmCal	2.14	1.36	3.95	0.66	3.08	2.88	1.80	4.98	0.82	3.59
Kou-SEJ-Q-deAm	2.14	1.37	3.97	0.67	3.09	2.83	1.86	5.06	0.91	3.63
CRR-HV	10.69	9.30	10.58	8.93	10.31	15.55	13.17	11.84	13.25	13.26
Calls: quotes	152	403	419	152	974	790	1,888	1,904	677	4,582
Kou-SEJ-Q (European)	2.27	1.46	5.10	0.76	3.90	3.17	1.82	6.55	0.89	5.00
Kou-SEJ-Q-Am (frozen)	2.25	1.38	4.50	0.70	3.45	3.19	1.67	5.49	0.78	4.25
Kou-SEJ-Q-AmCal	2.27	1.41	4.52	0.73	3.47	3.06	1.71	5.61	0.84	4.30
Kou-SEJ-Q-deAm	2.26	1.39	4.51	0.71	3.46	3.14	1.68	5.55	0.79	4.27
CRR-HV	10.11	9.73	13.10	9.38	11.62	12.98	13.87	17.20	14.05	15.66
<i>B. Mean error (IV points, event average)</i>										
Puts: quotes	178	426	408	147	1,012	1,044	1,956	1,664	659	4,664
Kou-SEJ-Q (European)	−1.39	0.41	0.69	0.32	0.30	−1.87	0.65	1.41	0.37	0.35
Kou-SEJ-Q-Am (frozen)	−1.39	0.41	0.70	0.32	0.31	−1.71	1.07	2.98	0.71	1.15
Kou-SEJ-Q-AmCal	−1.41	0.38	0.66	0.28	0.27	−1.93	0.84	2.81	0.48	0.95
Kou-SEJ-Q-deAm	−1.40	0.39	0.68	0.30	0.29	−1.80	0.97	2.92	0.61	1.07
CRR-HV	−9.73	−7.69	−8.55	−7.53	−8.55	−14.85	−11.19	−9.31	−11.44	−11.10
Calls: quotes	152	403	419	152	974	790	1,888	1,904	677	4,582
Kou-SEJ-Q (European)	0.92	−0.24	−3.55	−0.32	−1.65	2.42	−0.32	−5.02	−0.31	−2.18
Kou-SEJ-Q-Am (frozen)	0.95	−0.15	−3.00	−0.24	−1.29	2.48	−0.16	−4.08	−0.20	−1.61
Kou-SEJ-Q-AmCal	0.92	−0.19	−3.02	−0.28	−1.32	2.25	−0.40	−4.27	−0.43	−1.82
Kou-SEJ-Q-deAm	0.93	−0.17	−3.01	−0.27	−1.31	2.39	−0.26	−4.16	−0.30	−1.70
CRR-HV	−8.19	−8.53	−12.08	−8.11	−10.18	−10.82	−12.19	−16.32	−12.47	−13.84

former and 54% of the latter. In the money the European Kou-SEJ-Q already over-priced the puts (mean error 1.41), and the premium, which an at-the-money recalibration hardly changes, adds to that. Why these puts are over-priced I leave open; the shape of the announcement law (Section 8.4) and the parity-implied spot, which applies European parity to American quotes (Section 7.3), are candidates; Section 9 argues for the former. The recalibration lowered the model’s level (median change of the background volatility -0.7%), which cost the calls a little (RMSE 4.25 for Kou-SEJ-Q-Am against 4.30, a reduction of -1.3% $[-1.6, -1.0]$), and over all quotes the RMSE did not change (-0.1% $[-0.3, 0.1]$). Kou-SEJ-Q-deAm came between the frozen and the consistent calibration (put mean error 1.07): the binomial premium it removes from the mids is smaller than the model’s own, so its European level stayed higher. In the development sample, with rates near zero, the three American variants differ by at most 0.03 points of RMSE.

The reading of Section 8.7 therefore stands, with a qualification. Calibrated consistently, the American price still did not help for puts (RMSE 3.59 against 3.56 for the European Kou-SEJ-Q, a reduction of -0.6% $[-2.8, 1.8]$, and of -3.7% $[-5.7, -1.8]$ with the pre-registered vega-scaled

metric), while for calls with a dividend it cut the RMSE of the European model by 29.5% [24.3, 34.1]. The double count at the money explains only part of the put result.

F.6 Recent announcements: universe, firms and predictions

Table 52 lists the firms of the recent sample of Section 8.10, and Figure 17 compares the straddle-implied move with the realized move. Table 53 gives the mispricing captured on the out-of-sample expiration, Table 54 the verdicts of the pre-registered predictions applied to that expiration, Figure 18 the smiles of the eight largest moves, Figure 19 the errors across moneyness, and Table 55 the errors by firm.

Table 56, a post-freeze analysis that was not pre-registered, splits the third expiration of Table 14 by the length of the chain. In the 82 events whose chain has three expirations the longest one is the third, so BS-IVlong uses the at-the-money volatility of the evaluation expiration itself and coincides with the in-sample reference Mkt-ATM; only in the 59 events with a fourth expiration is it a forecast from a longer expiration.




Figure 17: Straddle-implied move (ATM straddle mid of the nearest expiration at d_{pre} divided by the spot) against the absolute realized log move from d_{pre} to d_{post} , one marker per announcement, with the 45-degree line. The nearest expiration is about two weeks out, so the implied move also contains two weeks of ordinary diffusion.

Table 52: Recent announcements (2025-05-01 to 2026-10-10) with pre- and post-event chains: number of events, median straddle-implied move (ATM straddle mid of the nearest expiration at d_{pre} divided by the parity-implied spot, %), median absolute realized log move between the parity-implied spots at d_{pre} and d_{post} (%), and median ATM implied volatility of the nearest expiration (about two weeks to expiry) at d_{pre} and at d_{post} (% annualized, trading-day convention). The nearest expiration in the source rolls with the snapshot date, so the post-event figure is a constant-tenor quantity rather than the same contract.

Symbol	Events	Straddle move	Realized	ATM IV pre	ATM IV post
TSLA	5	8.6	3.3	54	45
NVDA	6	8.4	2.8	52	40
META	5	8.6	10.3	52	34
NFLX	5	8.6	7.8	52	31
AMD	6	11.8	7.1	67	60
AMZN	6	8.7	7.2	53	32
GOOGL	5	7.3	2.2	46	32
MSFT	5	6.3	5.0	40	29
AAPL	6	5.6	2.3	35	28
PLTR	6	12.8	8.5	85	53
COIN	5	12.9	13.9	78	69
AVGO	6	9.5	7.6	61	42
SMCI	6	17.4	12.6	98	74
MU	6	12.5	3.9	72	54
INTC	5	13.4	11.2	80	61
UBER	6	9.0	5.2	51	37
CRWD	6	10.3	4.5	59	43
HOOD	4	13.8	10.5	78	67
LULU	6	12.1	12.8	74	45
DIS	6	7.8	7.2	44	27
BA	5	6.7	4.5	37	34
ORCL	6	13.2	11.9	76	52
CRM	6	9.5	3.4	56	38
ADBE	6	10.3	3.7	62	37
PYPL	5	9.5	8.0	63	34
ROKU	1	4.4	1.3	26	17
CVNA	3	16.5	8.7	95	72
MRNA	6	14.5	4.0	85	69
ABNB	6	9.7	2.6	60	36
NOW	5	10.4	4.6	60	46
TTD	5	17.5	6.6	104	61
All	165	10.0	5.8	62	41

Table 53: Mispricing captured, $MC = 1 - \sum |V_{\text{model}} - \text{mid}| / \sum |V_{\text{bench}} - \text{mid}|$, ratio of sums over the common sample of Table 14, with 90% event block-bootstrap intervals (1,000 draws). Columns 4–7 use BS-HV as the benchmark within subsamples: wing quotes ($0.15 \leq |\Delta| < 0.35$), at-the-money quotes ($0.35 \leq |\Delta| \leq 0.65$), and announcement year.

Model	vs BS-HV	vs BS-IVlong	Wing	ATM	2025	2026
DJ-hist	0.25 [0.14, 0.35]	−3.73 [−4.54, −3.06]	0.27	0.23	0.39	0.10
DJ-implied	0.86 [0.83, 0.88]	0.11 [−0.06, 0.26]	0.85	0.90	0.87	0.85
Kou-SEJ-P	0.27 [0.15, 0.36]	−3.64 [−4.49, −2.95]	0.29	0.25	0.40	0.10
Kou-SEJ-Q	0.86 [0.83, 0.88]	0.11 [−0.07, 0.27]	0.85	0.90	0.87	0.85
Kou-SEJ-Q-Am	0.86 [0.84, 0.88]	0.14 [−0.02, 0.28]	0.85	0.90	0.87	0.85
Mkt-ATM	0.89 [0.87, 0.90]	0.29 [0.17, 0.41]	0.88	0.94	0.90	0.88




Figure 18: Pre-event implied-volatility smiles of the nearest expiration for the eight largest recent announcement moves (one per firm, at least twelve usable quotes): market mid with the bid-ask range, Kou-SEJ-P and Kou-SEJ-Q smiles, the flat BS-HV and DJ-implied levels, and the market smile of the nearest expiration after the announcement.

Table 54: Pre-registered predictions (output/preregistration.md, written before any pricing error was computed) evaluated on the out-of-sample expiration of the recent sample, with a simpler reading than Tables 12 and 13: no intervals, and a prediction holds only if every stated component holds. RMSE in implied-volatility points, event averages unless stated; reductions relative to BS-HV. P2a, P2b, P4 and P6 are judged on Kou-SEJ-Q only (P2 of the main study on four models); the P2b terciles are those of the historical announcement s.d. of the recent events. P4: Kou-SEJ-Q RMSE pooled over the quotes of each tercile, with VIX terciles of the recent events and relative-spread terciles of the recent quotes (the main study uses event averages and cut points of its development sample). P6: gain $1 - \text{RMSE}_{\text{Kou-SEJ-Q}}/\text{RMSE}_{\text{DJ-implied}}$.

	Prediction	Measured	Verdict
P1	Event-jump models cut pre-event RMSE vs BS-HV by $\geq 30\%$	DJ-hist 32%; DJ-implied 75%; Kou-SEJ-P 31%; Kou-SEJ-Q 75% (post-event leg not priced)	holds
P2a	Reduction larger in the wings than at the money	Kou-SEJ-Q wing 79%, ATM 81%	fails
P2b	Reduction larger for high past-move firms	top tercile 70%, bottom 77%	fails
P3	No model beats Mkt-ATM at the money	Mkt-ATM 1.20 vs best other 1.61	holds
P4	Residuals largest in top VIX and top spread terciles	VIX 3.48/3.29/3.23; spread 2.60/3.39/3.91 (Kou-SEJ-Q RMSE by tercile)	fails
P5	Implied announcement variance beats historical	DJ 2.97 vs 8.08; Kou 3.00 vs 8.18	holds
P6	DE shape helps mainly in the wings; $< 5\%$ at the money	gain $1 - \text{RMSE}$ ratio, wing -5.6% , ATM -11.2% (Kou-SEJ-Q vs DJ-implied; negative: Kou-SEJ-Q worse)	fails
P7	BS-HV under-prices; DJ under-prices wings relative to ATM	BS-HV bias $-7.4/-7.1/-9.2$; DJ wing-ATM -0.5 (vol points)	holds
P8	Implied announcement sd exceeds realized; more so in high VIX; P forecasts moves better	ratio 1.01 [0.88, 1.16]; VIX low/high 0.77/1.29; MAE P 4.3 vs Q 4.3	fails




Figure 19: Mean signed implied-volatility error (model minus market) on the out-of-sample expiration of the recent sample by standardized moneyness, averaged within events and then across events, with 90% intervals.

Table 55: Pre-event RMSE in implied-volatility points by firm (out-of-sample expiration, event average; Events: the firm’s events in the common sample of Table 14), and the median straddle-implied and absolute realized moves over all the firm’s events of Table 52 (%).

Firm	Events	BS-HV	DJ-implied	Kou-SEJ-Q	Implied move	Realized
AAPL	6	12.6	3.1	3.1	5.6	2.3
ABNB	6	14.5	2.7	2.9	9.7	2.6
ADBE	6	12.1	2.8	2.6	10.3	3.7
AMD	6	15.2	1.8	1.8	11.8	7.1
AMZN	6	13.1	1.9	1.8	8.7	7.2
AVGO	6	13.0	2.1	2.1	9.5	7.6
BA	5	9.2	5.0	4.9	6.7	4.5
COIN	1	11.5	5.1	5.0	12.9	13.9
CRM	6	8.5	2.0	2.1	9.5	3.4
CRWD	5	14.6	3.2	3.2	10.3	4.5
DIS	5	13.3	3.5	3.2	7.8	7.2
GOOGL	5	11.4	1.9	1.8	7.3	2.2
INTC	5	8.3	3.5	3.4	13.4	11.2
LULU	6	17.6	6.3	7.4	12.1	12.8
META	5	9.8	1.4	1.1	8.6	10.3
MRNA	3	11.3	10.6	10.6	14.5	4.0
MSFT	5	10.7	2.0	1.9	6.3	5.0
MU	6	8.9	2.6	2.5	12.5	3.9
NFLX	5	14.9	2.0	1.8	8.6	7.8
NOW	5	12.1	2.1	2.0	10.4	4.6
NVDA	6	8.0	1.9	2.3	8.4	2.8
ORCL	6	10.5	3.2	2.6	13.2	11.9
PLTR	4	13.9	1.8	2.1	12.8	8.5
PYPL	4	12.4	3.9	3.8	9.5	8.0
SMCI	6	15.7	3.3	4.4	17.4	12.6
TSLA	5	4.7	1.7	1.7	8.6	3.3
TTD	1	26.6	5.1	4.1	17.5	6.6
UBER	6	10.1	3.4	3.3	9.0	5.2

Table 56: Post-freeze analysis of Table 14 (not pre-registered): the third expiration by the length of the chain. Three expirations: 82 events (1,871 quotes), in which the longest quoted expiration is the third, so BS-IVlong uses the ATM implied volatility of the evaluation expiration and coincides with Mkt-ATM; four expirations: 59 events (1,313 quotes), in which BS-IVlong uses a longer expiration. RMSE in implied-volatility points (event average) and mispricing captured relative to BS-IVlong with 90% event bootstrap intervals (1,000 draws, the bootstrap of Table 53).

Model	RMSE (event avg.)		MC relative to BS-IVlong	
	Three exp.	Four exp.	Three exp.	Four exp.
BS-HV	13.10	10.30	–	–
BS-IVlong	2.39	3.43	–	–
Mkt-ATM	2.39	2.41	–	–
DJ-hist	9.25	6.45	–5.37 [–6.67, –4.32]	–2.16 [–2.82, –1.42]
DJ-implied	3.17	2.69	–0.24 [–0.39, –0.13]	0.44 [0.24, 0.58]
Kou-SEJ-P	9.31	6.60	–5.22 [–6.54, –4.13]	–2.14 [–2.78, –1.42]
Kou-SEJ-Q	3.24	2.67	–0.23 [–0.38, –0.12]	0.43 [0.21, 0.58]
Kou-SEJ-Q-Am	2.96	2.66	–0.16 [–0.29, –0.05]	0.43 [0.21, 0.57]