Implied Event Risk from Option Prices*

Eben Lazarus† Richard Stanton‡ Johan Walden§

May 18, 2026

Abstract

We introduce a nonparametric method for recovering risk distributions of anticipated events of publicly traded firms, such as earnings announcements, securities litigation events, and the resolution of takeover bids and proxy fights. We first apply our method to study specific corporate events, namely Elon Musk's takeover bid of Twitter and the AT&T–Time Warner antitrust case. In both cases our recovered risk distributions are informative about the market beliefs of the outcome of the event. We then study a large sample of earnings announcements. Empirically, the recovered earnings announcement risk distribution is a strong predictor of realized announcement-day return moments, outperforming conventional option-implied measures that conflate announcement risk and other risks. Our method provides new insight on the effect of anticipated event risk on asset prices, how to measure such risk, and — more specifically — a fruitful way of using option prices to extract market beliefs about event risk.

*We thank Jonathan Napoles for excellent research assistance, and participants at CFE-CMStatisticcs 2025 for valuable comments and suggestions.

†Haas School of Business, U.C. Berkeley. Email: lazarus@berkeley.edu.

‡Haas School of Business, U.C. Berkeley. Email: richard.stanton@berkeley.edu.

§Haas School of Business, U.C. Berkeley. Email: walden@haas.berkeley.edu.

1 Introduction

A large number of extensions and variations of the original Black-Scholes option pricing model (Black and Scholes, 1973) have been introduced to explain deviations between that model's predictions and actual derivatives prices, e.g., manifested as implied volatility smiles and smirks across option strike prices.[1]

One well-known reason for such deviations is event risk (see, for example, Aït-Sahalia, 2004), which, in contrast to the continuous price innovations in the Black-Scholes model, leads to discontinuous jumps in asset prices at the point in time when the event occur. In some cases, the arrival time of such event risk is unknown. The event could, e.g., represent an unanticipated announcement of a merger with another company or an unanticipated lawsuit filed against the company. In many situations, the timing of the event is anticipated but the effect it will have on an asset's price is not. An earnings announcements is one such example of with known timing, that will likely affect the assets price. Other examples include the verdict of a lawsuit, which may be known in advance, political uncertainty that is resolved by an election (see Kelly et al., 2016), or an FOMC announcement (see Lucca and Moench, 2015).

An important distinction is whether the event risk is idiosyncratic or systematic. An earnings announcement for a small or medium-sized company would most likely have little or no impact on the market (see Leung and Santoli, 2014), whereas political uncertainty would likely have aggregate impact.

How should asset prices be adjusted to take anticipated event risk into account, especially when the event lies in the near future? Specifically, is it possible to distinguish the effect event has risk on the price, in advance, from the effects of other, non-event risk? A method to infer the market's view about the event risk distribution would obviously be helpful for forecasters, traders, and potentially for policy makers. It would also have practical value for traders, facing the well-known challenges of calibrating implied volatility curves around earnings announcements.

In this paper, we introduce a parsimonious nonparametric framework, to study the information about implied event risk distributions contained in option prices. We denote this the *event risk distribution recovery problem*. Specifically, given option prices at different maturities, our method recovers the event risk distribution, as well as the distribution of non-event risk.

Our approach can be applied to arbitrary discrete events that occur at a known point in time. In general, risk neutral event distributions are recovered. When event risk is idiosyncratic, unless there are strong frictions, risk neutral and physical probabilities coincide and the interpretation of the results becomes particularly simple. We therefore initially focus on events for which it is

[1] Early extensions include (see Rubinstein, 1994; Dupire, 1994; Derman and Kani, 1994), constant elasticity of variance models (CEV) (Cox (1975)), stochastic volatility models (Heston, 1993), models that incorporate isolated jumps (see Merton, 1976; Kou, 2002; Carr and Wu, 2003), the variance gamma model Madan et al. (1998), and more generally time-changed Lévy processes, see Carr et al. (2002).

reasonable to assume that event risk is idiosyncratic.

We study several examples of anticipated event risk. We first analyze two illustrative examples that are well-suited for our approach, namely Elon Musk's takeover bid of Twitter, and the AT&T–Time Warner antitrust case. We then focus on event risk associated with earnings announcements. An advantage with studying such events is data availability. Beyond recovering the announcement distribution, our decomposition also yields the normal-period distribution: the risk-neutral density for the week following the announcement. We validate this recovery empirically by comparing the recovered normal-period density to the ex-post risk-neutral distribution observed from options traded after the announcement, and show that it outperforms naive alternatives that do not separate announcement from normal-period risk.

Our paper relates to multiple strands of the literature. An extensive strand studies option prices around earnings announcements (see, e.g., Ravagli et al., 2014; Donders et al., 2000; Patell and Wolfson, 1981; Truong et al., 2012; Anagnostopoulou and Tsekrekos, 2015, 2017; Chung and Louis, 2017; Diavatopoulos et al., 2012; Ederington and Lee, 1996; Isakov and Pérignon, 2001; Jiang et al., 2012; Lei et al., 2020; Dubinsky et al., 2019; Alexiou et al., 2025; Liu et al., 2025). Our distinct contribution is our focus on recovering the whole event risk distributions from observed option prices.

A direct application of our method in this setting is to quantify the asymmetric tail risk surrounding earnings announcements and to relate that risk to how option markets price announcement-related skewness and downside exposure. This is closely related to recent evidence that skew premia around earnings announcements are economically meaningful in equity options markets (Neururer and Papadakis, 2026). More broadly, firm-level differences in announcement tail risk may also be related to mechanisms emphasized in the literatures on disagreement, short-sale constraints, and overvaluation (Boehme et al., 2006; Diether et al., 2002).

A second application is to connect the implied event distribution to evidence from option markets that higher risk-neutral moments measured before earnings announcements contain information about subsequent returns. In particular, Diavatopoulos et al. (2012) show that changes in option-implied skewness and kurtosis prior to earnings announcements predict subsequent stock and option returns, while Xing et al. (2010) show more generally that the shape of the individual-option volatility smirk predicts future equity returns and is linked to subsequent earnings shocks.

Another related strand studies the general recovery of objective probability distributions and pricing kernels from observed state prices (see Ross, 2015; Carr and Yu, 2012; Backwell, 2015; Jensen et al., 2019; Qin and Linetsky, 2016; Borovička et al., 2016; Walden, 2017; Schneider and Trojani, 2019; von Sydow and Walden, 2020). Our method has similarities with this strand, but our environment is somewhat different. Especially, we do not rely on time-homogeneous state transitions, since a key characteristic of the announcement is that it is different from normal innovations.

Finally, our paper is related to the literature on jumps in asset prices (see Aït-Sahalia, 2004; Busch et al., 2011; Lee, 2012; Seo and Wachter, 2019; Duong and Swanson, 2015; Jiang and Tian, 2010; Pan, 2002; Eraker et al., 2003; Eraker, 2004; Subramanian, 2004). Our approach differs from this literature by focusing on anticipated events, and by using option prices to draw inferences about the distribution functions of specific future jumps, under weak parametric assumptions.

The rest of the paper is structured as follows: In section 2, we introduce the model. In section 3, we study several event risk environments—specifically some firm-specific events, as well as a larger set of earnings announcements—all of which have in common that the timing of the event is ex ante known but its impact on asset prices is not. We document that the recovered event risk distribution is informative about actual event risk in all of these environments, and that the recovered normal-period distribution predicts the ex-post risk-neutral distribution observed after the announcement. Section 4 concludes. Proofs and additional examples are provided in the appendix.

2 Model

Time is discrete, except for one non-integer time $u > 0$ at which an announcement occurs that changes market beliefs about the asset's payoff. Thus

$$t \in \{0, 1, 2, \ldots, t^*, u, t^* + 1, t^* + 2, \ldots\}, \qquad t^* < u < t^* + 1.$$

The interval $(t^*, t^* + 1)$ is the *announcement period*. All other intervals $(t, t + 1)$, $t \neq t^*$, are *normal periods*. We assume the event time u is known.

The state at date t is

$$s_t = (a_t, m_t) \in \mathbb{Z} \times \mathbb{Z},$$

where a_t summarizes firm-specific, payoff-relevant information for the asset and m_t summarizes aggregate information relevant for the pricing kernel.

2.1 Preferences and pricing

A representative agent has expected utility preferences

$$EU_t = E_t \left[\sum_{v=t}^{T} \rho^{v-t} U(\tilde{C}_v) \right],$$

where $\rho \in (0, 1)$ is the subjective discount factor, U is increasing and strictly concave, and aggregate consumption is

$$\tilde{C}_t = C_t(m_t).$$

The announcement time u is a pure information arrival and is therefore not included in the summation.

Consider a contingent claim with terminal payoff

$$\tilde{F}_T = F(a_T),$$

where the payoff depends only on the payoff-relevant state at time T. Standard equilibrium arguments (see Lucas, 1978) imply that its time-t price is

$$P_t = E_t \left[\rho^{T-t} \frac{U'(\tilde{C}_T)}{U'(\tilde{C}_t)} \tilde{F}_T \right]. \tag{1}$$

Although the terminal payoff depends only on a_T, prices before T may also depend on m_t through the pricing kernel. An equivalent risk-neutral representation is

$$P_t = E_t^{\mathbb{Q}} \left[\left(\prod_{r=t}^{T-1} R_{r,r+1}^{-1} \right) \tilde{F}_T \right], \tag{2}$$

where $R_{r,r+1}$ is the one-period gross risk-free rate and $\mathbb{Q}$ is the risk-neutral measure. We assume the risk-free rate is non-stochastic, so

$$P_t = \left(\prod_{r=t}^{T-1} R_{r,r+1}^{-1} \right) E_t^{\mathbb{Q}} \left[\tilde{F}_T \right]. \tag{3}$$

2.2 Normal periods versus announcement periods

The model has two types of state transitions.

Normal periods. During a normal period $(t, t+1)$, the state evolves as

$$(a, m) \mapsto (a + k, m + k), \qquad |k| \leq K^N,$$

where $K^N > 0$ is a strictly positive constant. Thus payoff-relevant and pricing-kernel-relevant information move together in normal periods, and moreover may not move more than k^N steps in either direction. The risk-neutral probabilities for these moves are

$$\mathbb{Q}(s_{t+1} = (a + k, m + k) \mid s_t = (a, m)) \equiv q_k.$$

We assume these transition probabilities are time-invariant and state-independent,[2] so the risk-neutral state transition behaves like a time-invariant Markov process.

Announcement period. During the announcement period (t^*, t^*+1), the state changes in two steps. First, at the announcement time u, the state evolves as

$$(a, m) \mapsto (a + k, m), \qquad |k| \leq K^p,$$

with

$$\mathbb{Q}(s_u = (a + k, m) \mid s_{t^*} = (a, m)) = \mathbb{P}(s_u = (a + k, m) \mid s_{t^*} = (a, m)) \equiv p_k,$$

where $\sum_k p_k = 1$ and $K^p \geq 0$. The case $K^p = 0$ corresponds to a no-news announcement, so that $p_0 = 1$.

The announcement moves only the payoff-relevant state. It does not affect the pricing-kernel state m, so physical and risk-neutral announcement probabilities coincide. Economically, the event is firm-specific and has negligible implications for aggregate consumption.

Second, at (or right before) $t^* + 1$, the economy again evolves according to the normal-period transition probabilities q_k. Hence the distribution $\{q_k\}$ captures ordinary background news, while $\{p_k\}$ captures the discrete announcement risk that we seek to recover.

As discussed in the appendix, both normal and announcement period events can be modeled by Laurent operators $\mathcal{N} : \mathbb{R}^{\mathbb{Z}} \to \mathbb{R}^{\mathbb{Z}}$ and $\mathcal{A} : \mathbb{R}^{\mathbb{Z}} \to \mathbb{R}^{\mathbb{Z}}$, defined by

$$(\mathcal{N}x)_n = \sum_k x_k q_{n-k}, \qquad (\mathcal{A}x)_n = \sum_k x_k p_{n-k}.$$

The operator $\mathcal{N}$ represents a normal transition and $\mathcal{A}$ an announcement transition. These are both convolution operators, which implies that they commute, i.e., the order in which they are applied does not matter. The sequence of events and corresponding Laurent operators are shown in Figure 1.

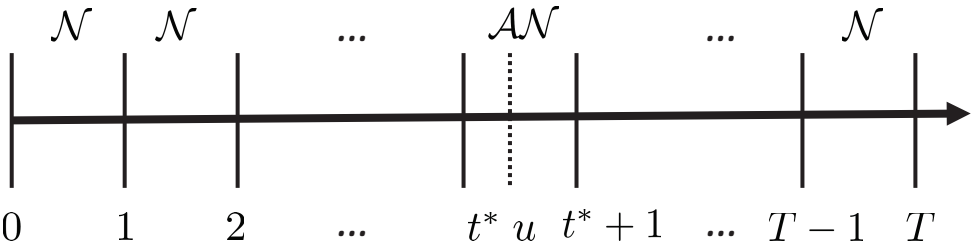

Figure 1: *Sequence of information events and corresponding Laurent operators.*

[2]This may seem like a restrictive specification, but our applications will be over fairly short time periods, say $t \leq 4$, where t is measured in hours, days, weeks, or months, in which case such a time invariance assumption may be reasonable.

2.3 Identification intuition

A single cross section of state prices or option prices is generally not enough to identify the announcement distribution. The reason is that the risk-neutral distribution at any given maturity combines two sources of uncertainty: ordinary background risk, summarized by q_k, and announcement risk, summarized by p_k (whenever that maturity spans the event date).

Our identification strategy uses prices at multiple maturities. Shorter-dated claims contain less accumulated background risk than longer-dated claims. Comparing state-price distributions across maturities allows the common normal evolution to be differenced out, thereby isolating the event distribution.

The formal recovery result is stated below. Its proof uses the Laurent-operator representation of normal transitions and announcement transitions. Since that notation is quite technical, we defer it to Appendix A.

2.4 Recovering the announcement distribution from state prices

Consider the initial state $s_0 = (0, 0)$, and let δ_k^t denote the Arrow–Debreu security that pays one dollar in payoff-relevant state k at time t, for $t = 1, \ldots, T$. Let its date-0 price be $[\Psi^t]_k$, so that $\Psi^t \in \mathbb{R}^{\mathbb{Z}}$. Define the associated risk-neutral probabilities by

$$q^t = \left(\prod_{r=0}^{t-1} R_{r,r+1}^{-1} \right) \Psi^t.$$

Thus q^t is the date-0 risk-neutral distribution over payoff states at horizon t.

If the announcement occurs after date t, then q^t reflects only normal transitions. If the announcement occurs before date t, then q^t reflects both normal transitions and the announcement shock. This is exactly what creates the identification problem—and, with multiple maturities, also what makes identification possible.

The following result provides recovery formulas for the announcement distribution.

Theorem 1. *Assume $T = rx$, where $r, x \in \mathbb{N}$. From the observed risk-neutral probabilities q^r, q^{T-x}, and q^T, define the Laurent polynomials*

$$X(z) = \sum_k q_k^r z^k, \qquad V(z) = \sum_k q_k^{T-x} z^k, \qquad W(z) = \sum_k q_k^T z^k.$$

(i) If $u > r$, then the announcement distribution can be recovered as

$$p_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-ik\omega} \frac{W\left(e^{i\omega}\right)}{X\left(e^{i\omega}\right)^x} \, d\omega. \tag{4}$$

(ii) If $u < T-1$ and $r = T$, then the announcement distribution can be recovered as

$$p_k = \frac{1}{2\pi}\int_{-\pi}^{\pi} e^{-ik\omega}\frac{V\left(e^{i\omega}\right)^T}{W\left(e^{i\omega}\right)^{T-1}}\,d\omega. \tag{5}$$

The theorem shows that the event distribution can be identified from state-price information at multiple maturities.

For implementation, the integrals in (4) and (5) can be evaluated efficiently using the fast Fourier transform (FFT) and its inverse (IFFT). These transforms for risk neutral probabilities, q, with $q_k = 0$, when $|k| \geq K$, are defined by

$$\hat{q}_\omega = \mathcal{F}[q](\omega) = \sum_{k=-K}^{K} q_k e^{-2\pi ik\omega/(2K+1)}, \qquad \omega = -K, -K+1, \ldots, K. \tag{6}$$

for the FFT, and

$$q_k = \mathcal{F}^{-1}[\hat{q}](k) = \frac{1}{2K+1}\sum_{\omega=-K}^{K} \hat{q}_\omega e^{2\pi ik\omega/(2K+1)}, \qquad k = -K, -K+1, \ldots, K. \tag{7}$$

for the IFFT. Henceforth we use the notation $q \in \ell_K^0$, if $q_k = 0$ when $|k| \geq K$. Further details are given in Appendix B.

2.5 Recovering the announcement distribution directly from option prices

So far we have described recovery from risk-neutral probabilities. In practice, however, one typically observes option prices rather than state prices. One can first infer risk-neutral probabilities from option prices, following Breeden and Litzenberger (1978), and then apply Theorem 1. It is preferable, however, to work directly with option prices, because repeated differencing tends to amplify noise.

We therefore derive a direct FFT-based recovery method from option prices. We focus on the case $T = 2$ and $0 < u < 1$, which is the case used in the empirical section. Then one- and two-period option prices are sufficient to identify the event distribution.

Consider time-t risk-neutral probabilities $q^t \in \ell_K^0$, and an underlying asset with prices

$$S_k = \begin{cases} k+K, & k \geq -K, \\ 0, & k < -K. \end{cases} \tag{8}$$

A put option with maturity t and strike S_r, $r \geq -K$, has price

$$P_r = \begin{cases} 0, & r = -K, \\ \sum_{k=-K}^{r-1}(r-k)q_k, & r = -K+1, -K+2, \ldots, K, \\ P_K + (r-K), & r > K. \end{cases} \tag{9}$$

Hence

$$q_k^t = \begin{cases} P_{k+1}, & k = -K, \\ P_{k-1} - 2P_k + P_{k+1}, & k = -K+1, -K+2, \ldots, \end{cases}$$

that is, risk-neutral probabilities are obtained from put prices by second differencing.[3]

We also define the $(2K+1)$-periodic second-difference operator $\widehat{D_+D_-}$, which coincides with the usual second-difference operator in the interior and is completed periodically at the boundaries:

$$\begin{aligned} \left(\widehat{D_+D_-}P\right)_{-K} &= P_K - 2P_{-K} + P_{-K+1}, \\ \left(\widehat{D_+D_-}P\right)_{K} &= P_{K-1} - 2P_K + P_{-K}. \end{aligned}$$

The following theorem shows how to recover the event distribution directly from put prices.

Theorem 2. *Consider the model with $T = 2$ and $0 < u < 1$, with put prices $P^1, P^2 \in \ell_K^0$. Define adjusted prices $\dot{P}^1, \dot{P}^2 \in \ell_K^0$ by*

$$\begin{aligned} \dot{P}_k^1 &= P_k^1 - \frac{k+K}{2K+1}\left(P_K^1 - \frac{1}{2}(2K+2)\right) - \frac{1}{2}\frac{(k+K+1)(k+K+2)}{2K+1}, \\ \dot{P}_k^2 &= P_k^2 - \frac{k+K}{2K+1}\left(P_K^2 - \frac{1}{2}(2K+2)\right) - \frac{1}{2}\frac{(k+K+1)(k+K+2)}{2K+1}, \end{aligned}$$

for $k = -K, -K+1, \ldots, K$. Define

$$\hat{P}^i(\omega) = \mathcal{F}\left[\dot{P}^i\right](\omega), \qquad i = 1, 2, \tag{10}$$

and

$$F \stackrel{\text{def}}{=} \mathcal{F}^{-1}\left[\frac{\hat{P}^1(\omega)^2}{\hat{P}^2(\omega)}\right]. \tag{11}$$

Then the event distribution $p \in \ell_K^0$ satisfies

$$p_k = \begin{cases} \left(\widehat{D_+D_-}F\right)_k + \frac{1}{2K+1}, & |k| \leq K, \\ 0, & |k| > K. \end{cases} \tag{12}$$

This direct option-price method has several advantages over first recovering state prices and

[3] So P contains the same information as the risk neutral probabilities, q^t, in line with the observation in Breeden and Litzenberger (1978).

then applying Theorem 1. Most importantly, it uses second differencing only once, in the final step, which reduces noise amplification. In addition, because asynchronous or noisy option prices can generate negative implied state prices, it is useful in practice to impose convexity constraints on P^1, P^2, and F. We return to this issue in the empirical section. Finally, the method uses only FFT operations and is therefore computationally very fast.

2.6 Examples

We first consider a market with $K = 30$, and $T = 4$. The risk neutral probability distribution during a normal period, q, has a double exponential distribution, as shown in the left panel of Figure 2. The announcement risk distribution is shown in the right panel, and is the mixture of two exponential distributions; it is bimodal. It could, for example, reflect the anticipated outcome of a lawsuit, where the company may be held liable and forced to pay a large fine. In this case, the company's future cash flows will decrease a lot. On the other hand, if the company is cleared of liability, cash flows will be unaffected. The company's current valuation incorporates the risk of a potential negative outcome, so if it is cleared of liability, its stock price will increase.

The announcement occurs at $u = 2.5$. It follows that both approach (i) and (ii) in Theorem 1 are feasible in this case: (i) with $r = 2$, $x = 2$, and (ii) with $r = 4$, $x = 1$, can be used to recover the announcement risk distribution. Specifically, method (i) uses q^2 and q^4 to recover p, whereas method (ii) uses q^3 and q^4 to recover p.

Figure 3 shows the risk neutral probabilities associated with contingent claims with payoffs between 1 and 4, q^1, q^2, q^3, and q^4. The leftmost panel, q^1, is the same as that in Figure 2, since $u > 1$, and therefore $q^1 = q$. The effect of the announcement risk is quite pronounced in the third panel, for q^3, the first probabilities that include the announcement at $u = 2.5$, is q^3. The smoothing effect of the normal event risk over three periods is sufficient for the risk neutral probabilities to no longer be bimodal, but the effect of the bimodal announcement operator $\mathcal{A}$ still has a major effect. At $T = 4$, further smoothing has occurred, and the effect of of the announcement operator is less pronounced.

The recovered event risk distribution is shown in the top panel of Figure 4, using method (i) in Theorem 1. These are the red dots, which look indistinguishable from the actual announcement distribution. In the middle and lower panel, the differences between recovered and true announcement probabilities are shown for the two methods, method (i) in the middle panel and method (ii) in the lower panel. The point-wise errors are of the order 10^{-11}, so the method works very well in this case.

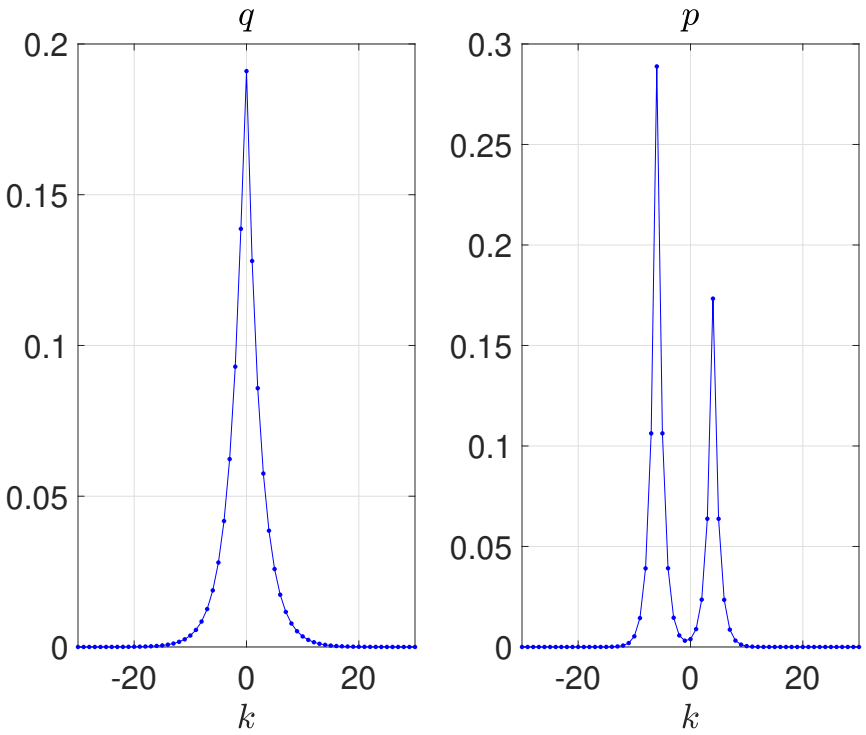

Figure 2: *Left: Risk neutral probabilities during normal period, q. Right: Announcement probabilities. Specification:* $p_k = C_1 \left(e^{-|k+6|} + 0.6e^{-|k-4|}\right)$, $q_k = C_2 e^{-0.4|k+0.1|}$.

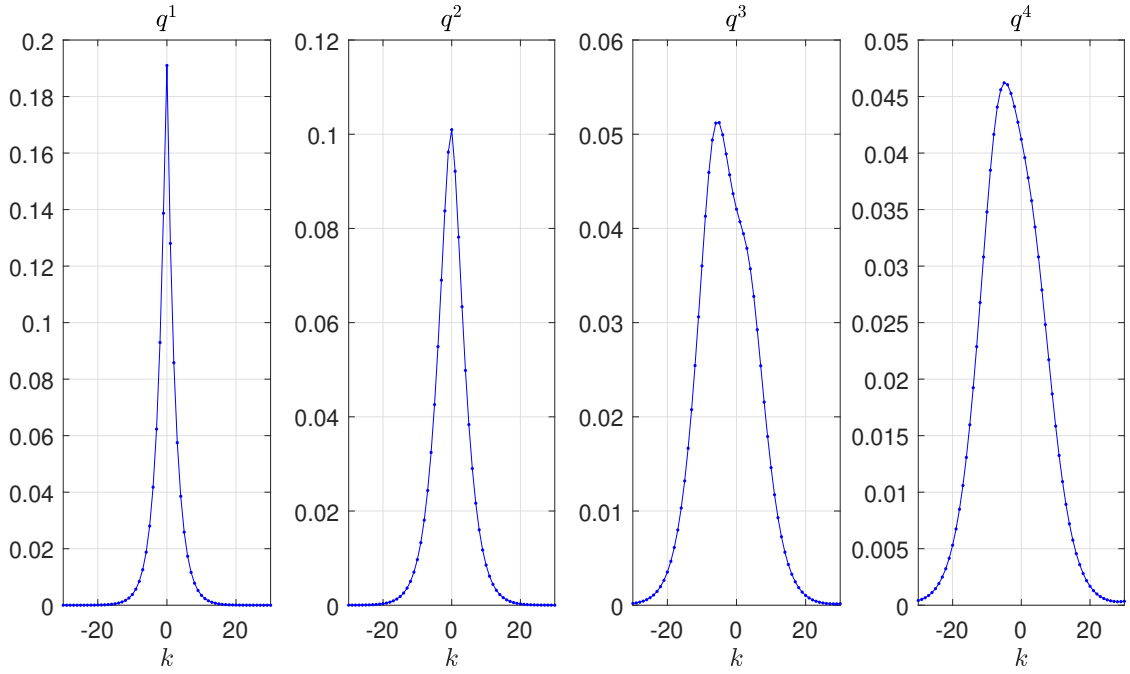

Figure 3: *Risk neutral probabilities,* q^1, q^2, q^3, q^4.

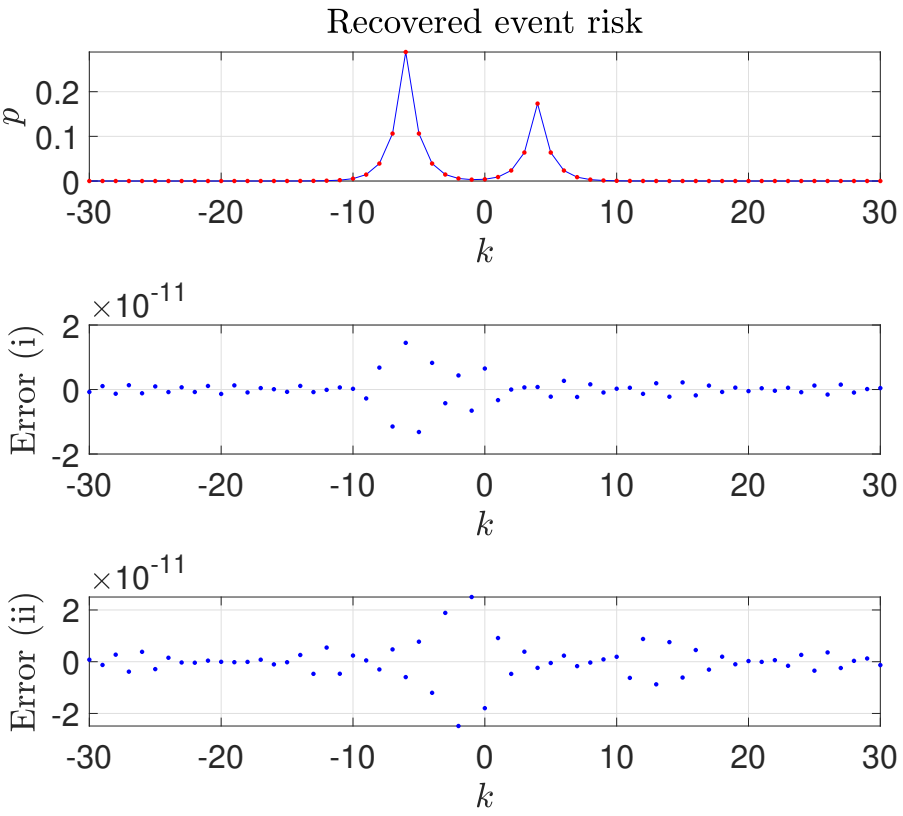

Figure 4: *Top panel: Recovered event risk distribution (top panel). Middle and lower panel and errors of two methods of Theorem 1.*

Binomial model with event risk

A second example is given by the binomial tree approximation of the Black Scholes model with constant coefficients, with $K = 1$, $q_1 = \frac{R-d}{u-d}$ and the risk neutral probabilities during a normal period are $q_{-1} = \frac{u-R}{u-d}$, and where u and $d = 1/u$ are the sizes of up and down moves respectively in the binomial tree model. In this case, the state represents the value of a stock, $S_t = u^k S_0$, in payoff-relevant state k at time t.

To this we add announcement risk. We study an example with $T = 28$ time periods, each representing a day, and where the announcement occurs before $t = 14$. Thus, method (ii) is applicable, with $r = 2$ and $x = 14$. We choose $u = 1/d = e^{0.02}$, corresponding to a 2% continuously compounded volatility per day, i.e., 38.2% per year. The risk-free rate is normalized to $R = 1$.

Figure 5 shows four different examples, each with a different type of event risk. In row 1, the topmost row, there is no announcement, corresponding to $a = \delta_0$. The two leftmost panels show the risk neutral probabilities for 14 and 28 days into the future. These are basically binomially distributed, except for that only probabilities for an even number of moves are non-zero, due to the recombining nature of the binomial tree method. The next panel shows the 28-day implied (daily) volatilities associated with these risk neutral probabilities, and different strike prices, K. We see that they are close to 0.02, the difference stemming from the approximation the binomial tree model provides of the Black-Scholes continuous time model.

Finally, the rightmost panel shows the actual and recovered event risk, from the risk neutral probabilities, q^{14} and q^{28} using method (ii) in Theorem 1. We see that the method basically perfectly identifies the event risk. Note that the method is basically nonparametric, in the sense that it does not make any assumptions about the form of $\mathcal{N}$ and $\mathcal{A}$, when recovering the announcement risk, a. Specifically, the fact that the normal-period dynamics are generated by binomial-tree dynamics is not used anywhere in the recovery method.

In row 2, a bell-shaped announcement risk distribution is used. This increases the spreads of the risk-neutral distributions (two leftmost panels), and increases the level of the implied volatility curve (third panel from the left) compared with the benchmark case. Qualitatively these panels look similar as if the per-period volatility were higher than 0.02.[4] Again, the event risk distribution is very well recovered.

In row 3, a bimodal announcement risk distribution is used. Specifically, the stock price either moves up by a quite specific amount ($k \mapsto k + 4$ or $k \mapsto k + 5$, with equal probability) represented by the sharp spike on the right, or moves down more diffusively by 1 to 11 steps, represented by the left hump. The effect on the implied volatility curve is to make it more "hump"-shaped on the

[4] A couple of technical differences also arise. First, the risk-neutral distributions are now smooth, as seen in row 2. They are no longer zero for odd-numbered moves, since the announcement may shift the state both an even number and an odd number of steps. Second, the implied volatility curve is meaningfully defined over a wider range of strike prices, K, since the announcement may shift the state, and thereby the stock price, further than when there are only binomial tree moves.

left, because of the more narrowely defined downside compared with the wider bell-shaped risk in row 2, as well as flatter in the right part, given the more limited potential upside move compared with that figure. As in the previous example, the recovered event risk distribution is very close to the actual risk distribution.

Finally, in row 4, a two-point announcement risk distribution is used. The corresponding implied volatility curve is now even further affected, having a "hat"-like shape. As in the previous cases, the recovered event risk distribution coincides closely with the actual one, as shown in the rightmost panel.

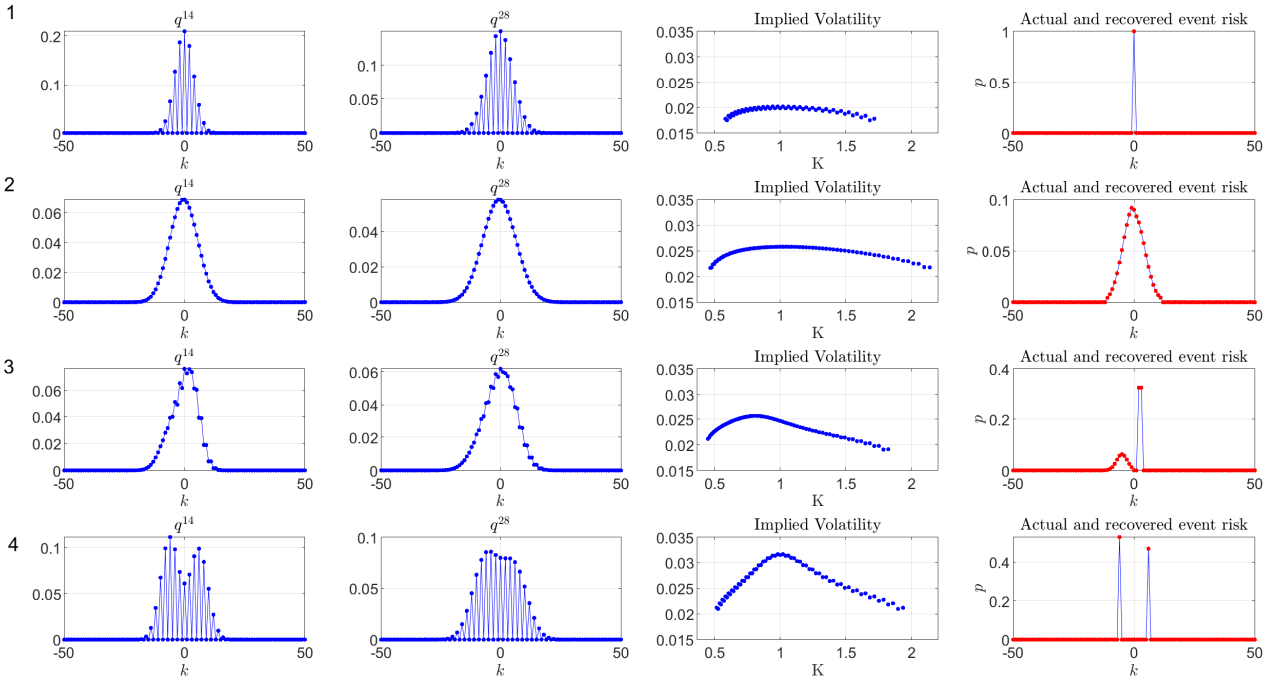

Figure 5: *The figure shows risk neutral probabilities, implied volatilities and announcement risks in Black Scholes model with event risk. Row 1 panels have no event risk; Row 2 panels have bell-shaped event risk; Row 3 panels have point distribution plus bell shape event risk; Row 4 has two-point distributed event risk. Column 1 shows risk neutral probabilities for* $t = 14$ *; Column 2 shows risk neutral probabilities for* $t = 28$*; Column 3 shows associated implied volatilities with call option prices; Column 4 shows actual (blue solid line) and recovered (red dotted line) event risk.*

In the appendix, we provide a third example, corresponding to a binomial tree version of the constant elasticity of variance (CEV) model, with additional event risk added. It is shown that this model is also efficiently handled by our recovery approach.

3 Empirical results

The empirical analysis asks whether option prices before a firm-level announcement contain useful information about the distribution of announcement-day stock returns. We first provide two empirical examples of high-profile firm-related announcements that illustrate the method. We then describe our larger-scale analysis for firm-level earnings announcements, and we study how the recovered announcement distributions help predict realized returns. In both cases, the timing of the events are known ex ante, but the events' impacts on share prices are not.

3.1 Two illustrative examples

To build intuition, we start with two firm-level examples to illustrate what the recovered announcement distributions look like in practice, before turning to the broader panel evidence.

A binary takeover announcement: Twitter and Elon Musk

Our first example considers the takeover of Twitter by Elon Musk in 2022. Musk agreed in April 2022 to acquire Twitter at $54.20 per share, but he attempted to walk away from the agreement in May of that year after a share-price decline. Twitter responded by suing to enforce the merger. In early October, Musk signaled his willingness to complete the merger at the original price. The Delaware judge overseeing the case paused the lawsuit to provide time for the two parties to close the merger. The judge set an explicit deadline of October 28 for Musk to close the deal at the agreed-upon price, after which the case would go to trial.[5] This decision created a clear, publicly known date by which the key announcement about the deal — whether the acquisition would close on the agreed terms — would be resolved. As of October 21, a week prior to the Friday, October 28 deadline, the deal had not closed. The week from October 21 to October 28 therefore provides an ideal setting for our approach: investors faced a single, well-defined announcement window, with little ambiguity about the timing of the news.

Figure 6 shows the recovered announcement distribution for Twitter based on options observed on Friday, October 21 and expiring on the following two Fridays (October 28 and November 4). We describe the recovery procedure in detail in the following subsection. The distribution has a pronounced mode very close to $54.20, the price specified in the purchase agreement and eventually agreed to. The distribution also indicates meaningful risk of the deal collapsing, leading to large negative returns from the October 21 spot price of $49.9 per share. The Twitter episode thus provides a clean illustration of how our method recovers a discrete, announcement-specific risk-neutral distribution that concentrates mass on economically meaningful states.

[5] See "Twitter lawsuit halted so Elon Musk can close deal by Oct. 28," Reuters, October 6, 2022.

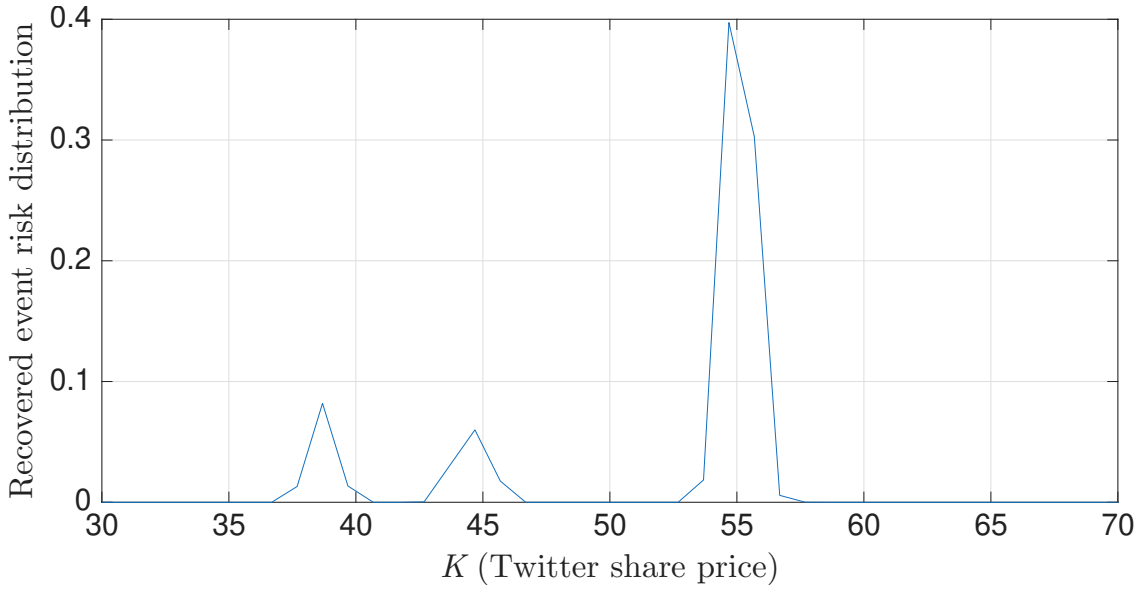

Figure 6: *Recovered event risk distribution for Twitter on October 21, 2022.*

An antitrust court decision: AT&T and Time Warner

Our second example examines the proposed merger between AT&T and Time Warner. AT&T agreed in 2016 to acquire Time Warner for roughly $80 billion, but the U.S. Department of Justice filed suit to block the deal on antitrust grounds. The case went to trial in a federal district court. The judge in the case stated at the end of April 2018 that he planned to issue a decision on June 12.[6] The underlying merger agreement was set to expire on June 21, so the court's guidance effectively created a narrow window in which the key announcement — whether the deal would be allowed to proceed — had to be made.

Figure 7 plots the recovered announcement distribution for Time Warner using options observed on Friday, June 8, expiring on the following two Fridays (June 15 and June 22). As in the Twitter example, this is the last weekly pre-decision trading day before the expected ruling. The distribution again places substantial mass on a price region consistent with the deal being approved (at $107 per share). It also assigns meaningful probability to negative return outcomes (with prices below the June 8 spot price of $95.34 per share) that would prevail if the acquisition were blocked or delayed. This pattern is consistent with investors viewing approval as the most likely outcome (as was eventually realized), but not a foregone conclusion. This episode thus gives a second illustration of how our method recovers a discrete, event-specifc distribution in a setting where the timing of a legal decision is known ex ante.

[6]See "AT&T-Time Warner Trial: Judge Says He'll Have Decision in June," Variety, April 30, 2018.

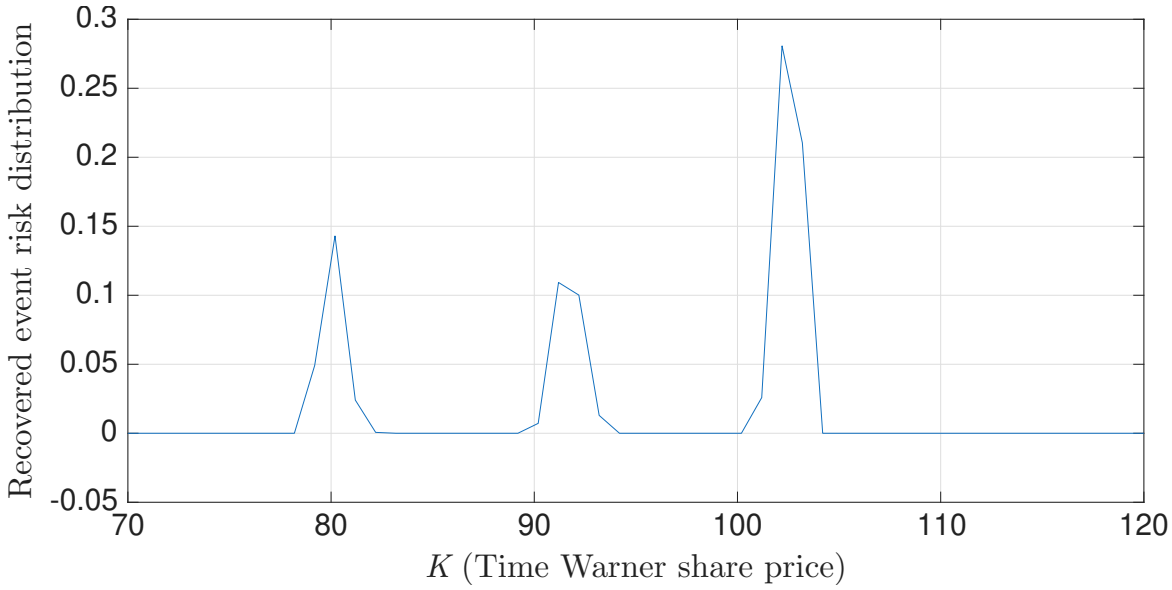

Figure 7: *Recovered event risk distribution for Time Warner on June 8, 2018.*

3.2 Earnings announcement data

To broaden the analysis, we now study quarterly earnings announcements for large U.S. firms in the S&P 500. We test whether our recovered ex ante announcement distribution helps predicts the distribution of returns on the day of the earnings announcement. Our sample covers earnings announcements over the 2010s.[7] We obtain quarterly earnings announcement dates from Compustat and match these to daily announcement-day stock returns from CRSP. Option quotes come from OptionMetrics Ivy DB US.

Our analysis sets a period to be one week, to avoid confounding from other discrete risks that might arise with longer horizons. Weekly options expire on Fridays at market close. For each earnings announcement, we therefore identify three Fridays. The first is the Friday before the announcement, which we treat as trading date $t = 0$. The second ($t = 1$) is the first Friday after the announcement,[8] and the third ($t = T = 2$) is the following Friday. On the pre-announcement trading day, we collect all options that expire on either of the two post-announcement Fridays. This gives us a pre-announcement option panel with two nearby maturities, for which the announcement occurs between the trading date and $t = 1$.

We apply a set of simple filters to the option data: we require a strictly positive bid price, a strictly positive Black-Scholes implied volatility, and strictly positive open interest or trading volume. We work with mid price quotes. For each expiry we require at least eight distinct strikes

[7] Firm-level weekly options were not generally traded in 2010–2011. Our analysis considers a weekly horizon, so our sample effectively starts in 2012 and runs through the end of 2019. Some of our analysis filters by firm market capitalization, also obtained from CRSP.

[8] A small subset of announcements happen on Friday. In all cases, these occur before the market open, so for these we use the announcement date as the first expiration date.

with valid quotes. Events that do not meet these criteria are excluded. Applying these filters yields nearly 4,300 earnings events with usable pre-announcement options over our sample period. This is a subset of the universe of nearly 20,000 S&P 500 earnings announcements over the 2010s, as many firms do not have sufficient (if any) option quotes at the one- or two-week horizon. This particularly affects small stocks and early-sample observations.

3.3 From option prices to announcement distributions

The recovery results in Section 2, in particular equations (10)–(12), show how to recover the announcement distribution from option prices. To implement this empirically, we first face the issue that observed single-stock option prices have meaningful measurement noise. Because the recovery algorithm ultimately takes second differences of put prices, small pricing errors can generate large oscillations in the implied distribution. We therefore smooth the raw option quotes and impose the monotonicity and convexity restrictions that hold for no-arbitrage put-price schedules before applying Theorem 2.

Our smoothing procedure operates as follows. First, for each maturity, we start from the Black-Scholes implied volatilities for puts and calls reported by OptionMetrics. As out-of-the-money options are typically more liquid than in-the-money options, we use puts at strike prices K below the current underlying price S_0, and calls at strikes $K > S_0$. To avoid a discrete jump in the implied volatility and price schedule at $K = S_0$, we blend the put and call observations in a symmetric band $[K_\ell, K_h]$ around S_0. In our implementation we set $K_\ell = 0.9S_0$ and $K_h = 1.1S_0$, and within this range, we consider all available strikes for which both put and call implied volatilities are available.

Formally, denote the put and call implied volatilities as $\sigma^P(K)$ and $\sigma^C(K)$, and denote our final blended implied volatility as $\sigma^{\text{blend}}(K)$. For strikes $K < K_\ell$, we take $\sigma^{\text{blend}}(K) = \sigma^P(K)$ whenever put implied volatilities are available.[9] For strikes $K > K_h$, we take $\sigma^{\text{blend}}(K) = \sigma^C(K)$ whenever call implied volatilities are available. On strikes in the band $[K_\ell, K_h]$ where both puts and calls have valid implied volatilities, we combine the two according to

$$\sigma^{\text{blend}}(K) = w(K)\,\sigma^P(K) + \big(1 - w(K)\big)\,\sigma^C(K), \tag{13}$$

where $w(K) = 1$ for $K \le K_\ell$, $w(K) = 0$ for $K \ge K_h$, and $w(K) = \frac{K_h - K}{K_h - K_\ell}$ for $K_\ell < K < K_h$. Puts thus dominate in the left tail and calls dominate in the right tail, with a smooth linear transition near at-the-money strikes where both sides trade actively. Using $\sigma^{\text{blend}}(K)$, the spot price S_0, the time to maturity, and the risk-free yield curve, we then evaluate the Black-Scholes put-price formula to obtain the no-arbitrage price $P(K) = P_K$ at each strike. The collection $\{K_j, P(K_j)\}_{j=1}^J$ serves as the input to the next step.

[9] In implementing Theorem 2, we consider strike prices with moneyness ranging from −40% to +40%. When this range is wider than the range of strikes with observable quotes, we impute implied volatility from the lowest available put strike for the left tail and from the highest available call strike for the right tail.

Second, to ensure positive risk neutral probabilities, we convexify the put-price schedule with respect to strike price. Fix an expiry and write $\{K_j, P_j\}_{j=1}^{J}$ for the strikes and observed (blended) put prices after filtering. We fit a least-squares quadratic spline approximation $\tilde{P}(K)$ to the observed put prices,

$$\min_{\tilde{P}} \sum_{j=1}^{J} \left(\tilde{P}(K_j) - P_j \right)^2, \tag{14}$$

subject to the convexity restriction

$$\tilde{P}''(K) \geq 0 \quad \text{for all } K. \tag{15}$$

In addition, we impose that the risk-neutral probabilities implied by the double difference of $\tilde{P}$ sum to one by enforcing that the terminal value of the cumulative distribution function associated with $-D_+D_-\tilde{P}$ equals one on a fine strike grid, in line with the construction in equations (10)–(12). This constrained spline approximation problem is a quadratic program in the spline coefficients. This projection replaces the noisy quotes by the closest smooth, convex curve with respect to strike price, such that the associated risk-neutral distribution is positive and integrates to one. It eliminates no-arbitrage violations and reduces the microstructure noise that would otherwise be magnified when taking second differences of prices in the last step of the procedure. We apply this convexification separately to each maturity.

Finally, we use the convexified put prices for the two post-announcement expiration dates to implement the algorithm in Theorem 2. Let $\tilde{P}^1$ and $\tilde{P}^2$ denote the fitted put-price functions for the shorter and longer maturity, respectively. We evaluate these functions on an evenly spaced grid of strikes to construct the discrete vectors P^1 and P^2 in (10)–(11), which we then use to obtain the announcement distribution $p = (p_k)_k$ through the mapping in (12). In the final step we again apply our convexified-spline approximation to the put-price schedule associated with p, which ensures that the implied state prices remain nonnegative and that small oscillations in the recovered announcement distribution are eliminated. Our final announcement distribution is over a range of strikes with moneyness between -40% and +40%, spaced \$1 apart from each other.

For each event, we then compute the variance, skewness, kurtosis, and crash probability implied by the ex ante recovered announcement distribution. The crash probability is defined as the probability that the announcement-day return will fall below the 5^{th} percentile of all earnings-announcement realized returns; in our sample, this corresponds to an announcement-day return of -6.0% (non-annualized).

3.4 Summary statistics and earnings announcement return predictability

Table 1 reports summary statistics for the announcement-day equity returns and for features of the recovered announcement distributions. Our baseline sample consists of 4,288 earnings announce-

	Mean	Median	Std. dev.	Obs.
Realized announcement return (%)	0.12	0.27	4.24	4,288
Ex ante announcement variance	0.0033	0.0021	0.0035	4,288
Implied crash probability (%)	8.34	4.89	11.07	4,288

Table 1: *Summary statistics for earnings announcements.*

Notes: Returns are daily earnings-announcement-day equity returns (non-annualized). Announcement variance is the ex-ante variance of the recovered announcement distribution, measured in squared-return units. Implied crash probability is the model-implied probability that the announcement return falls below −6% (the empirical 5th percentile).

ments. Realized announcement returns tend to be volatile, with a daily standard deviation of above 4% for non-annualized returns. This corresponds to realized variance of 0.002. In comparison, the ex ante variance implied by the recovered announcement distribution is slightly higher, at 0.004.[10] The model-implied probability of a crash below −6% (the empirical 5th percentile for announcement returns) is 9.1%, slightly elevated relative to the true crash frequency of 5%.

We next relate the ex ante objects to realized announcement outcomes. We start by regressing squared announcement-day returns on the recovered announcement distribution variance. We benchmark the predictive power of our recovered measure by comparing it to a prediction based on the one-week implied variance for that stock, constructed identically to the one-week squared VIX but at the single-stock level. We refer to this implied variance as VX^2, following the CBOE's convention for referring to single-stock volatility indices.[11] All regressions cluster standard errors by ticker and calendar quarter. Results are shown in Table 2.

Column (1) of Table 2 shows that the recovered announcement-distribution variance is a strong and significant predictor of squared returns. Each unit increase in ex ante variance corresponds to a 0.34-unit increase in squared announcement-day returns on average. This estimate likely reflects some degree of attenuation bias resulting from noise in the option prices and resulting announcement distribution, but overall the results indicate strong predictability. This predictability is comparable to the predictability in squared returns obtained from using the squared VX at the one-week horizon, as in column (2). In a joint regression (column (3)), both significantly help predict realized squared returns.

We also consider a subsample of larger firms (top 50% by market equity within our baseline sample), which are likely to have more liquid options. Results for this subsample are shown in columns (4)–(6) of Table 2. Our recovered ex ante variance continues to strongly predict squared

[10] In the analysis below, we also use higher moments of the announcement distribution. On average, these recovered distributions feature modest negative skewness and positive excess kurtosis, as in the data for realized returns.

[11] We construct VX as of the same trading date at which we calculate the announcement distribution, using the observed option prices expiring the Friday following the announcement. Our VX calculation follows the CBOE's VIX methodology, described at `https://www.cboe.com/tradable-products/vix/faqs/`. For an example of a single-stock VX index, see `https://www.cboe.com/us/indices/dashboard/vxapl/`.

	(1)	(2)	(3)
Announcement variance	0.4624***		0.3845***
	(0.0932)		(0.1375)
VX^2		0.3192***	0.0826
		(0.0651)	(0.0946)
Obs.	4288	4288	4288
R^2	0.065	0.047	0.066

Table 2: *Regressions for squared realized announcement-day returns.*

Notes: Dependent variable is the announcement-day centered squared return $(r-\mu)^2$, where μ is the mean of the ex ante announcement distribution. "Announcement variance" is the ex ante announcement distribution variance. "VX^2" is the squared pre-announcement one-week single-stock volatility index, constructed using the VIX formula. The sample uses all earnings announcements with available options. Standard errors are two-way clustered by ticker and calendar quarter. Statistical significance at the 10%, 5%, and 1% levels denoted by *, **, and ***, respectively.

returns, and it now significantly outperforms the squared VX. In the joint regression in column (6), it fully drives out the effect of the VX.

Table 3 considers predictions for other aspects of the realized return distribution. Column (1) regresses the cubed announcement-day return on the ex ante skewness implied by the recovered announcement distribution; column (2) regresses the fourth moment of the realized return on the ex ante kurtosis; and column (3) regresses a crash indicator (equal to 1 for a realized return below the 5^{th} percentile of the return distribution, -6%) on the recovered ex ante crash probability. The announcement-implied skewness has little predictive power for the third moment, but announcement-implied kurtosis loads positively on the realized fourth moment, indicating that fatter ex-ante tails are associated with fatter realized tails. The ex-ante crash probability is also positively related to realized crashes, showing that our recovered distribution helps predict crash risk.

Taken together, the results provide evidence that our methodology allows one to measure useful ex ante information regarding announcement-specific risks. It also removes the normal-period confounding that affects measures like implied volatility or the VX, as indicated by its strong performance as compared to the VX shown in Table 2.

3.5 Predicting future implied volatility

In addition to predicting event risk distributions, $p \in \ell^0$, our method can also be used to predict post-event normal period risk distributions, $q \in \ell^0$. This is useful, e.g., if one wants to understand what implied volatility will look like after an announcement event.

We compare our estimate of q with alternative specifications, which do not incorporate the fact that the announcement only affects dynamics during the announcement period. Consider a

	(1) Third moment	(2) Fourth moment	(3) Crash indicator
Corresponding ex ante recovered moment	-0.0514	0.1925**	0.2015***
	(0.1064)	(0.0855)	(0.0476)
Obs.	4288	4288	4288
R^2	0.000	0.016	0.009

Table 3: *Regressions for higher moments and crash indicators.*

Notes: Dependent variables are the centered third and fourth moments of the announcement-day return and an indicator for an announcement-day crash (return below -6%, the 5^{th} percentile of the realized announcement return distribution). Each column regresses the corresponding realized outcome on the matching announcement-implied moment: skewness, kurtosis, and the ex-ante crash probability. We use all events with available options. Standard errors are two-way clustered by ticker and calendar quarter. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

situation with $T = 2$ and $u = 0.5$, so that the announcement occurs in the first period. It follows that the operator in the first period is $\mathcal{AN}$, and the operator in the second period is $\mathcal{N}$.

In the absence of an announcement, the second-period operator would be the same as the first-period operator (both would be $\mathcal{N}$), and it could therefore be inferred from the first period operator, or equivalently from the risk-neutral probabilities q^1. The operator could also be inferred from the two-period risk-neutral probabilities, q^2, by finding the operator $\mathcal{N}$ that satisfies leads to q^2 when applied twice, since no announcement event would occur in the first period. Both these estimates are incorrect in the presence of event risk.

In Figure 8, we provide a specific example. The upper panels show the normal (left) and announcement (right) event risk distributions. The lower panels show the estimated normal period risk distributions from q^1 (left) and q^2 (right). Both estimates significantly misrepresent the normal period risk distributions, q. Specifically, both estimates to some extent include the announcement risk distribution, p, in the estimate. In contrast, our method almost perfectly recovers p in this example (the mean-squared error between estimated and actual q is 8×10^{-15}, and is solely due to double precision arithmetics errors in this theoretical example).

Empirical validation of the recovered normal-period distribution

We now test this empirically. For each earnings announcement in our sample, we compare the recovered q to the ex-post risk-neutral distribution q^{post}, obtained from options traded on the first post-announcement Friday for the second Friday — that is, the option-implied density for a week that does not span any announcement. As benchmarks, we use the two naive alternatives described above: the one-week density q^1 (which conflates announcement and normal-period risk) and the

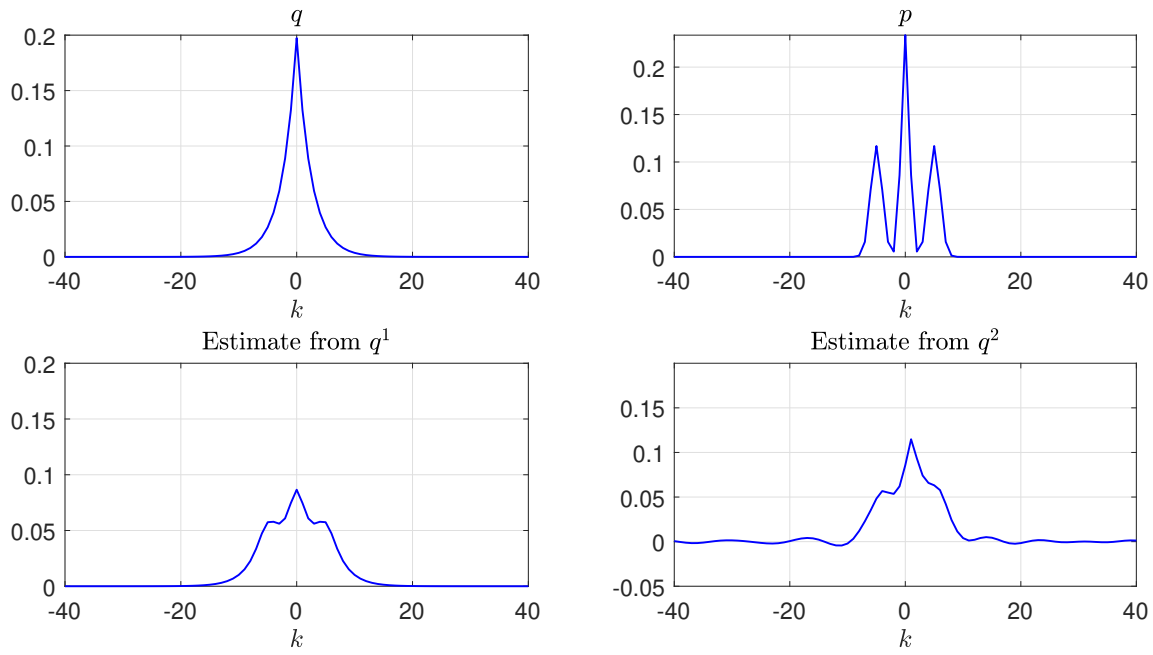

Figure 8: *Upper left panel: Normal period risk distribution; Upper right panel: Event risk distribution; Lower left panel: Estimated normal period risk distribution from q^1; Lower right panel: Estimated normal period risk distribution from q^2.*

convolution square root $\sqrt{q^2}$ (which assumes identical weeks).[12]

One practical issue is that the Laurent recovery involves pointwise division in the Fourier domain, which amplifies noise at frequencies where the denominator is near zero. To address this, we apply Tikhonov regularization — formally identical to ridge regression — which damps high-frequency noise while preserving the low-frequency structure of the recovered density.[13] Without regularization, many Fourier coefficients of $\hat{q}^1(\omega)$ are near zero, producing recovered densities with substantial negative mass that is then distorted by the convexification step.

Table 4 reports results for $N = 2{,}960$ events.[14] The key comparison is the variance regression:

[12] All distributions — forecasts and ex-post truth — are centered to zero mean on the comparison grid before scoring. This removes small location differences arising from the use of different forward prices for the pre-announcement and post-announcement densities, and ensures that the comparison focuses on distributional shape and spread rather than location. We evaluate forecasts on a common moneyness grid $m \in [-0.20, 0.20]$ with spacing $\Delta m = 0.005$ (81 points).

[13] Specifically, we replace the pointwise ratio $\hat{q}^2(\omega)/\hat{q}^1(\omega)$ with the regularized ratio $\hat{q}^2(\omega)\,\overline{\hat{q}^1(\omega)}\,/\,(|\hat{q}^1(\omega)|^2 + \gamma)$, where $\gamma = \gamma_0 \cdot \overline{|\hat{q}^1(\omega)|^2}$ and γ_0 is a scalar penalty parameter. The analogy to ridge regression is exact: γ_0 plays the role of the penalty on regression coefficients, and $|\hat{q}^1(\omega)|^2 + \gamma$ replaces the near-singular "moment matrix" $|\hat{q}^1(\omega)|^2$. We select γ_0 via cross-validation on the first two years of the sample (2013–2014) and evaluate out of sample on the remaining years (2015–2018). The cross-validated optimum is $\gamma_0 \approx 0.5$–1.0, depending on the scoring criterion. We also considered Lavrentiev regularization, which adds a real constant c directly to the denominator — replacing $\hat{q}^1(\omega)$ with $\hat{q}^1(\omega) + c$ — rather than regularizing via the squared modulus. This approach performed consistently worse, likely because adding a real constant to complex-valued Fourier coefficients shifts them rather than uniformly damping their influence.

[14] This sample is smaller than the full earnings announcement sample because the validation requires sufficient post-announcement option data. Events without adequate post-announcement option data are excluded.

Method	Quad. Score	KS	Var. MSE ($\times 10^6$)	β	R^2
q (Tikhonov, $\gamma_0 = 0.5$)	-0.000271	0.139	3.35	0.537	0.191
q (Tikhonov, $\gamma_0 = 1.0$)	-0.000222	0.143	4.16	0.483	0.224
q (unregularized)	-0.000776	0.221	3.92	0.547	0.112
One-week density (incl. annc.)	-0.000176	0.156	8.48	0.250	0.086
$\sqrt{q^2}$ (i.i.d. assumption)	-0.000277	0.170	9.66	0.235	0.042

Table 4: *Normal-period distribution: comparison with ex-post risk-neutral distribution.*

Notes: $N = 3{,}735$ scoring events. All distributions are centered to zero mean before scoring. The regularization parameter γ_0 is selected via time-split cross-validation (training on 2013–2014, testing on 2015–2018). Quadratic score: higher is better. KS: Kolmogorov–Smirnov statistic (lower is better). Variance regression: $\sigma^2_{\text{post}} = \alpha + \beta\hat{\sigma}^2 + \varepsilon$, where σ^2_{post} is the variance of the ex-post risk-neutral distribution and $\hat{\sigma}^2$ is the variance of the forecast density.

$\sigma^2_{\text{post}} = \alpha + \beta\hat{\sigma}^2 + \varepsilon$, where σ^2_{post} is the variance of q^{post} and $\hat{\sigma}^2$ is the variance of the forecast density. A perfect forecast implies $\beta = 1$ and $R^2 = 1$. The regularized Laurent recovery explains substantially more cross-sectional variation in post-announcement variance than either benchmark: $R^2 = 0.21$–0.24, compared to 0.09 for the one-week density and 0.05 for the convolution square root. The regression slope is also closer to unity ($\beta = 0.53$–0.59 versus 0.25–0.26), indicating a better-calibrated forecast of event-specific variance. The Kolmogorov–Smirnov statistic tells a similar story: the regularized recovery achieves a mean KS of 0.139, compared to 0.152 for q^1 and 0.169 for $\sqrt{q^2}$. The unregularized Laurent inversion performs worst on both the quadratic score and KS statistic, confirming that regularization is essential for the deconvolution to yield useful distributional forecasts.

These results are consistent with the interpretation that the Laurent decomposition successfully extracts normal-period risk-neutral information that is distinct from the raw option-implied densities contaminated by announcement risk.

The same Tikhonov regularization can also be applied to the recovered announcement distribution p, where a qualitatively similar noise-amplification issue can arise from near-zero Fourier coefficients.[15]

4 Conclusions

We have introduced a nonparametric method for recovering event risk distributions from option prices at two maturities. In simulations, the method recovers the announcement distribution with essentially zero errors. Empirically, the recovered announcement distribution is a strong predictor of realized announcement-day return moments, outperforming conventional option-implied measures that conflate announcement and normal-period risk. The decomposition also yields a normal-period

[15] Applying Tikhonov regularization to the announcement distribution raises the R^2 for variance prediction modestly, and raises the R^2 for kurtosis more substantially to 0.152. The improvement is driven by the same mechanism.

risk-neutral distribution that substantially outperforms naive benchmarks — including the one-week option-implied density and the convolution square root of the two-week density — at predicting the post-announcement risk-neutral distribution. Tikhonov regularization further improves recovery quality for both the normal-period and announcement distributions.

Beyond earnings announcements, our method applies naturally to any scheduled event with bracketing option expirations. One such type of events that has received substantial attention in recent years are FOMC announcements, and the effect these have on the market. FOMC events are hardly idiosyncratic, so an additional challenge here is that our method recovers the risk-neutral announcement distribution, whereas one is typically interested in the objective announcement distribution. Our method may still be informative, especially with respect to higher moments of these distributions.

Another future application area is to prediction markets, where rapidly expanding companies like Polymarket, Kalshi, and DraftKings, now structure their products as binary options that can be traded in a market. It is an open question how informationally efficient these new markets are — for example, concerns over insider trading have been voiced. Our method could be used to compare implied event distributions from market data for event that are known to affect the market, e.g., elections and the Brexit vote. An interesting application of our method would be to compare the prediction from market data with those in prediction market, to study whether one is superior to the other.

Appendices

A Operator representation

This appendix introduces the operator notation used in the proofs and in the recovery formulas.

A.1 Sequence spaces

Let

$$\mathbb{R}^{\mathbb{Z}} = \{x : \mathbb{Z} \to \mathbb{R}\}.$$

We use the Banach spaces

$$\ell^p = \left\{ x \in \mathbb{R}^{\mathbb{Z}} : \sum_{k\in\mathbb{Z}} |x_k|^p < \infty \right\}, \qquad p \geq 1,$$
$$\ell^\infty = \left\{ x \in \mathbb{R}^{\mathbb{Z}} : \sup_{k\in\mathbb{Z}} |x_k| < \infty \right\},$$

with norms

$$\|x\|_p = \left(\sum_k |x_k|^p \right)^{1/p}, \qquad \|x\|_\infty = \sup_k |x_k|.$$

We also define ℓ^0 as the subspace of ℓ^1 with finite support, and ℓ^0_K as the subspace whose elements satisfy $x_k = 0$ for $|k| > K$.

For $1 \leq p < q$, we have $\ell^p \subset \ell^q$, and therefore

$$\ell^0_0 \subset \ell^0_1 \subset \ell^0_2 \subset \cdots \subset \ell^0 \subset \ell^p \subset \ell^q \subset \ell^\infty \subset \mathbb{R}^{\mathbb{Z}}.$$

Specific elements are $\delta_k \in \ell^0_k$, with $(\delta_k)_k = 1$, and $(\delta_k)_{k'} = 0$, $k \neq k'$. We identify the transition probabilities p and q with elements in ℓ^0_K, $p = (\ldots, p_K, p_{K-1}, \ldots, p_{-K}, \ldots) \in \ell^0_K$, $q = (\ldots, q_K, q_{K-1}, \ldots, q_{-K}, \ldots) \in \ell^0_K$. The special case when there is no announcement risk arises when $p = \delta_0$. The standard binomial tree model is a special case of our model, in which no announcements occur, $K = 1$, and only the normal-period risk neutral probabilities q_{-1} and $q_1 = 1 - q_{-1}$ are nonzero.

A.2 Laurent operators

Define the Laurent operators $\mathcal{N}, \mathcal{A} : \mathbb{R}^{\mathbb{Z}} \to \mathbb{R}^{\mathbb{Z}}$ by

$$(\mathcal{N}x)_n = \sum_k x_k q_{n-k}, \qquad (\mathcal{A}x)_n = \sum_k x_k p_{n-k}.$$

The operator $\mathcal{N}$ represents a normal transition and $\mathcal{A}$ an announcement transition. Since both are convolution operators, they commute:

$$\mathcal{N}\mathcal{A} = \mathcal{A}\mathcal{N}.$$

Their adjoints are

$$(\mathcal{N}^* x)_n = \sum_k x_k q_{k-n}, \qquad (\mathcal{A}^* x)_n = \sum_k x_k p_{k-n},$$

and these also commute:

$$\mathcal{N}^* \mathcal{A}^* = \mathcal{A}^* \mathcal{N}^*.$$

Technically, the adjointness relation is defined on the subspace ℓ^2, with inner product

$$\langle v, w \rangle = \sum_k v_k w_k.$$

Since both physical and risk-neutral probabilities sum to one, the constant vector

$$\mathbf{1} = (\ldots, 1, 1, 1, \ldots) \in \ell^\infty$$

is an eigenfunction with eigenvalue one for both operators:

$$\mathbf{1} = \mathcal{A}\mathbf{1}, \qquad \mathbf{1} = \mathcal{N}\mathbf{1}. \tag{16}$$

A.3 Pricing recursion

Let $f \in \mathbb{R}^{\mathbb{Z}}$ be the payoff vector with components $f_k = F(k)$. Define

$$v_t = \frac{1}{R^{T-t}} \mathcal{A}^{\iota(t)} \mathcal{N}^{T-t} f, \qquad \iota(t) = \begin{cases} 1, & t < u, \\ 0, & t > u. \end{cases}$$

Then the pricing relation can be written as

$$P_t = [v_t]_{a_t}.$$

It follows that

$$v_t = \frac{1}{R_{t,t+1}} \mathcal{N} v_{t+1}, \qquad t \neq t^*, \qquad \text{and} \qquad v_{t^*} = \frac{1}{R_{t^*,t^*+1}} \mathcal{A}\mathcal{N} v_{t^*+1}. \tag{17}$$

The sequence of information events and corresponding operators is shown in Figure 1.

A.4 A stationary-price special case

In the special case in which the payoff-relevant state itself indexes the time-invariant price of an asset, so that

$$P_t = G(a_t), \qquad t = 0, 1, \ldots, T,$$

for some $G \in \mathbb{R}^{\mathbb{Z}}$, we have $v_t = v_{t+1} \equiv v$. If in addition the risk-free rate is constant, $R_{t,t+1} \equiv R$, then (17) implies

$$v = \frac{1}{R} \mathcal{N} v, \qquad v = \mathcal{A} v. \tag{18}$$

Defining the return vector $\mu = v_{t+1}/v_t$, this becomes

$$\mathbf{1} = \frac{1}{R}\mathcal{N}\mu, \qquad \mathbf{1} = \mathcal{A}\mu. \tag{19}$$

This is the case for the underlying asset in the classical binomial tree.

A.5 Date-0 state-price distributions

Let Ψ^t denote the date-0 Arrow–Debreu price vector at horizon t, and let

$$q^t = \left(\prod_{r=0}^{t-1} R_{r,r+1}^{-1}\right)\Psi^t$$

be the corresponding date-0 risk-neutral distribution. Then

$$q^t = (\mathcal{A}^*)^{1-\iota(t)}(\mathcal{N}^*)^t\delta_0, \qquad \iota(t) = \begin{cases} 1, & t < u, \\ 0, & t > u. \end{cases} \tag{20}$$

If $u > 1$, then $q_k^1 = q_k$ for all k. Also, if $q \in \ell^0_{K^N}$ and $p \in \ell^0_{K^P}$, then $q^t \in \ell^0_{tK^N+K^P}$.

A.6 Proof intuition

The identification problem is analogous to recovering a scalar a from observations of an^T: this is impossible from one maturity alone, but feasible from multiple maturities. Specifically, the scalar version of the challenge is the following: consider the challenge of recovering the real positive number $a \neq 0$ from observations of an^T, where $T \neq 0$ is a whole number and $n \neq 0$. This is impossible.

However, if one also observes an^t, where $t \neq T$ is also a whole number, then it is easy to recover a, e.g., by using the formula $a^{T-t} = \frac{(an^t)^T}{(an^T)^t}$, and then calculating $a = (a^{T-t})^{1/(T-t)}$. Note that when a may also be negative, the formula only provides a unique solution if $T-t$ is odd, since there will be two possible solutions when $T-t$ is even, $\pm a$. More generally, if a belongs to the complex field, $a \in \mathbb{C}$, uniqueness is only obtained when $T-t = 1$.

Theorem 1 exploits exactly this logic.

B FFT implementation

For $f \in \ell^0_K$, define the FFT by

$$\hat{f}_\omega = \mathcal{F}[f](\omega) = \sum_{k=-K}^{K} f_k e^{-2\pi i k\omega/(2K+1)}, \qquad \omega = -K, -K+1, \ldots, K. \tag{21}$$

The inverse FFT is

$$f_k = \mathcal{F}^{-1}[\hat{f}](k) = \frac{1}{2K+1} \sum_{\omega=-K}^{K} \hat{f}_\omega e^{2\pi i k\omega/(2K+1)}, \qquad k = -K, -K+1, \ldots, K. \tag{22}$$

Since

$$\mathcal{F}[f](\omega) = \mathcal{L}[f]\Big(e^{-2\pi i\omega/(2K+1)}\Big),$$

Theorem 1 can be implemented directly using FFT methods.

C Proofs

Proof of Theorem 1: The results follow immediately from the properties of the Z-transform on $p, q \in \ell^0$. Specifically, the Z-transform associates the Laurent polynomial $A(z) = \sum_k p_k z^k$ with p, and the Laurent polynomial $Q(z) = \sum_k q_k z^k$ with q. Moreover, $U(z) = A(z)^{\iota(r)} Q(z)^r$, $V(z) = A(z)^{\iota(T-x)} Q(z)^{T-x}$, and $W(z) = A(z)Q(z)^T$. So, when $u > r$,

$$\frac{W(z)}{U(z)^x} = \frac{A(z)Q(z)^T}{(Q(z)^r)^x} = A(z),$$

and when $u < T-1$,

$$\frac{V(z)^T}{W(z)^{T-1}} = \frac{(A(z)Q(z)^{T-1})^T}{(A(z)Q(z)^T)^{T-1}} = \frac{A(z)^T Q(z)^{T(T-1)}}{A(z)^{T-1} Q(z)^{(T-1)T}} = A(z).$$

Finally, by the Cauchy residue theorem from complex analysis, it follows that $p_k = \frac{1}{2\pi}\int_{-\pi}^{\pi} e^{-ik\omega} A(z) d\omega$, completing the proof. ∎

Proof of Theorem 2: First, note that since $P^2 \in \ell^0_K$, and $q^2 \in \ell^0_{2K^N+K^p}$, it follows that $2K^N + K^p \leq K$. This in turn implies that $q_{-K} = 0$, and $p_{-K} = p_{-K+1} = 0$.

Now, the translation properties of the FFT implies that $\mathcal{F}\left[\widehat{D_+D_-}f\right] = 2(\cos(\omega) - 1)\hat{f}(\omega)$, $\omega \neq 0$, for any $f \in \ell^0_K$.

It is straightforward to verify that $\dot{P}^i$ is constructed so that $\left(\widehat{D_+D_-}\dot{P}^i\right)_k = q^i_k - \frac{1}{2K+1}$ when $k = -K, -K+1, \ldots, K$. Indeed, in the interior, for $|k| < K$, the $\widehat{D_+D_-}$-operator is identical to the D_+D_--operator, and the property follows immediately from the facts that $(D_+D_-P^i)(k) = q^i_k$ in the interior, that D_+D_- maps linear functions $[ak+b]_k \in \mathbb{R}^{\mathbb{Z}}$ to 0 and the quadratic function $[k^2]_k \in \mathbb{R}^{\mathbb{Z}}$ to 2, just like the second-order differential operator satisfies $\frac{d^2}{dx^2}(ax+b) = 0$ and $\frac{d^2}{dx^2}x^2 = 2$ in the continuous setting.

At $k=-K$, we have

$$\begin{aligned}\left(\widehat{D_+D_-}\dot{P}^i\right)_{-K} &= \dot{P}^i_K - 2\dot{P}^i_{-K} + \dot{P}^i_{-K+1}\\ &= \frac{1}{2K+1}P^i_K + \frac{K(2K+2)}{2K+1} - \frac{1}{2}(2K+2)\\ &\quad - 2\left(P^i_{-K} - \frac{1}{2}\times\frac{2}{2K+1}\right)\\ &\quad + P^i_{-K+1} - \frac{1}{2K+1}\left(P^1_K - \frac{1}{2}(2K+2)\right) - \frac{1}{2}\times\frac{2\times 3}{2K+1}\\ &= q^i_{-K} - K - 1\\ &\quad + \frac{1}{2K+1}\left(K(2K+2)+2+K-3\right)\\ &= q^i_{-K} - \frac{1}{2K+1}.\end{aligned}$$

A similar argument implies that

$$\left(\widehat{D_+D_-}\dot{P}^i\right)_K = q^i_K - \frac{1}{2K+1}.$$

It therefore follows that

$$\mathcal{F}(p)(\omega) = \begin{cases} 2(\cos(\omega)-1)\times\frac{\hat{P}^1(\omega)^2}{\hat{P}^2(\omega)}, & \omega\neq 0,\\ 1 & \omega = 0.\end{cases}$$

Equivalently, in the physical domain

$$p_k = \begin{cases} \left(\widehat{D_+D_-}F\right)_k + \frac{1}{2K+1}, & |k|\le K,\\ 0, & |k|>K.\end{cases}$$

We are done ∎

D Additional examples

CEV model with event risk

We also study a model that mimics constant elasticity of variance (CEV) dynamics of the underlying asset in normal periods. At time T, the asset price satisfies $v_T(k) = v_T(k-1)\left(1+bv_T(k-1)^{\alpha-1}\right)$, $0<\alpha\le 1$. Risk-neutral transition probabilities are given by $q_{-1}=q_0=q_1=\frac{1}{3}$. Prices at earlier points in time are derived from (17), and as before we normalize the risk-free rate to $R=1$. Normal-period transition are given by a trinomial tree, with one-period volatility $\hat{b}=b\sqrt{\frac{2}{3}}$, and the asset's return dynamics are then a discrete approximation of the continuous time CEV diffusion process

$$\frac{dv_t}{v_t} = \hat{b}v_t^{\alpha}d\omega_t,$$

where ω_t is a standardized Brownian motion. This is similar to how the standard binomial tree model provides a discrete approximation of the continuous time constant-coefficient Black-Scholes model.

We let each period represents half a day, set $T = 56$, corresponding to 28 days, and assume that the risk-neutral probabilities q^{28} and q^{56} are available. As before, the announcement is assumed to occur before $t = 28$, so method (ii) in Theorem 1 is applicable. We set the volatility coefficient to $\hat{b} = 0.015$, corresponding to an instantaneous annualized volatility of $\sigma = \hat{b} \times \sqrt{2 \times 365} \approx 0.4$ at the asset price $v = 1$. The elasticity parameter is $\alpha = 0.4$. The announcement occurs before $t = 28$, and we apply method (ii) in Theorem 1.

Figure 9 show the state prices q^{28} and q^{56}. The effect of the announcement is quite pronounced in q^{28}, but much less so in q^{56}. In Figure 10, several volatilities are compared. The black line corresponds to the instantaneous volatility of the CEV model, which is decreasing and convex in the current spot price, and has the value $\hat{b} = 0.015$ when the spot price is 1. The red line corresponds to the implied volatility curve of the associated with the CEV model, as a function of the option strike price, when there is no announcement risk. This curve is also convex and decreasing, but less so than the spot-price curve, because the averaging effect, since the implied volatility is a function of the expected average volatility between 0 and T, rather than on the instantaneous volatility. The blue dotted line represents the implied volatilities associated with option prices in the CEV model with announcement risk. We see that it deviates significantly from the red line, due to the announcement risk: It (i) lies significantly above the red line, (ii) is partly increasing and partly decreasing, and (iii) has a "hat"-shape, with both concave and convex regions.

Finally, in Figure 11, we show actual and recovered event risk. The two curves are, again, indistinguishable. We stress, again, that the only input to the method are the state prices, q^{28}, and q^{56}. Although the CEV model and the event risk specifications are used to generate these state prices, the recovery method itself does not "know" this. It simply applies Theorem 1 to whatever state prices are given.

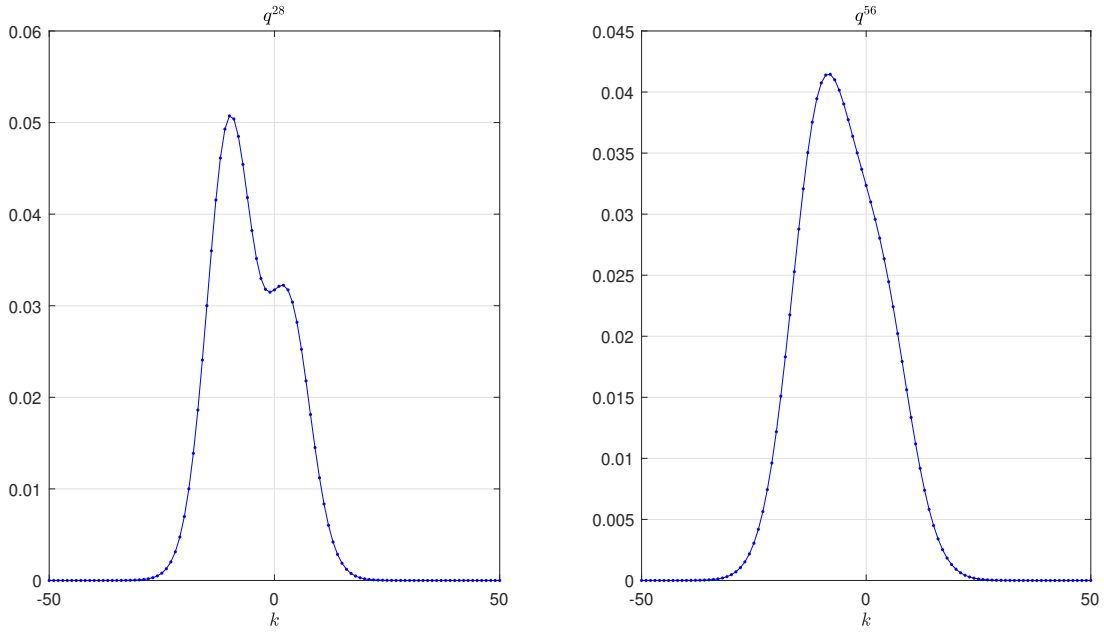

Figure 9: *Left Panel: 28 period (half days) state prices. Right Panel: 56 period (half days) state prices.*

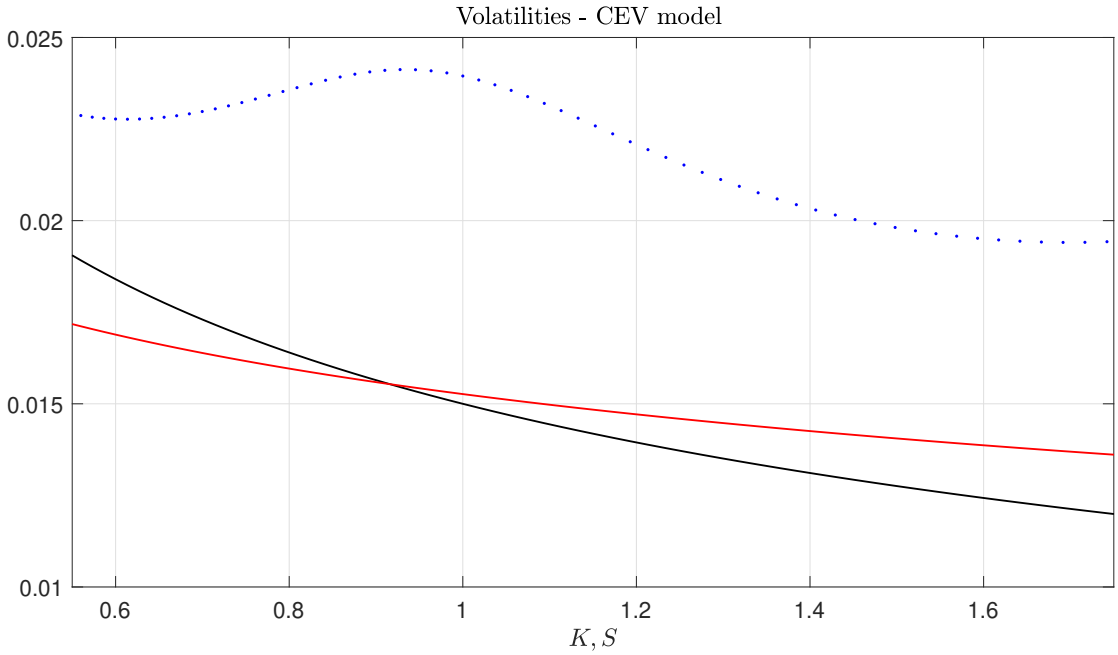

Figure 10: *Instantaneous volatility of CEV model (solid black line), implied volatilities of prices generated by CEV model (solid red line), and implied volatilities of model that includes both CEV risk and announcement risk (dotted blue line).*

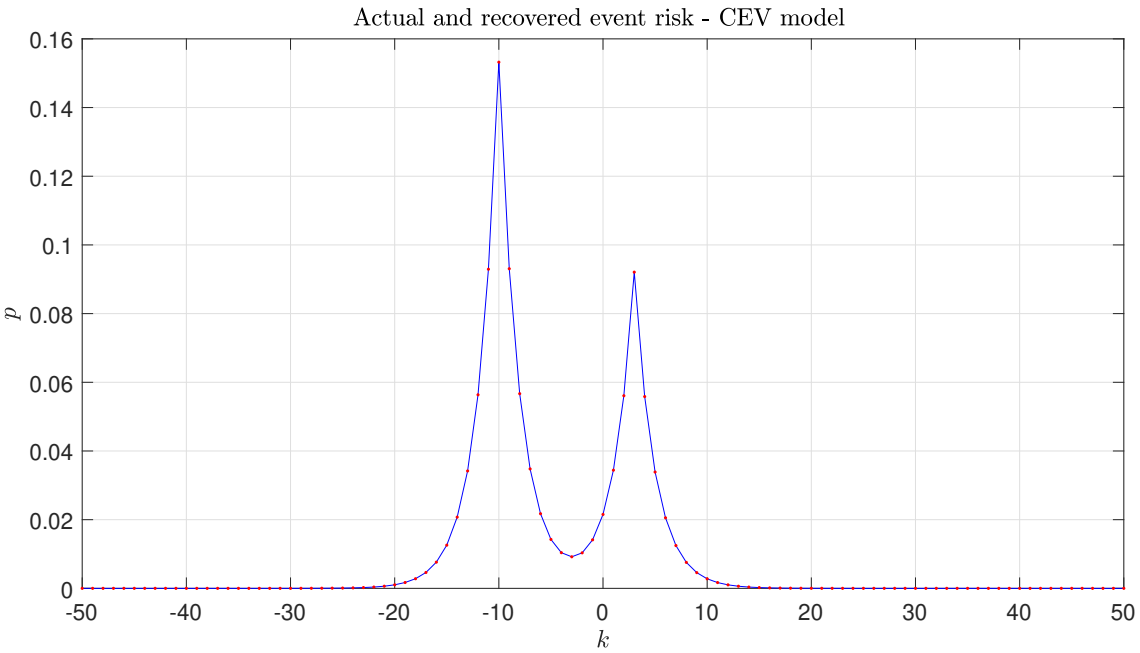

Figure 11: *Recovered event risk distribution, using method (ii) of Theorem 1.*

References

Aït-Sahalia, Yacine, 2004, Disentangling diffusion from jumps, *Journal of Financial Economics* 74, 487–528.

Alexiou, Lykourgos, Amit Goyal, Alexandros Kostakis, and Leonidas Rompolis, 2025, Pricing event risk: Evidence from concave implied volatility curves, *Review of Finance* 29, 963–1007.

Anagnostopoulou, Seraina C., and Andrianos E. Tsekrekos, 2015, Accounting quality, information risk and implied volatility around earnings announcements, *Journal of International Financial Markets, Institutions and Money* 34, 188–207.

Anagnostopoulou, Seraina C., and Andrianos E. Tsekrekos, 2017, Accounting quality, information risk and the term structure of implied volatility around earnings announcements, *Research in International Business and Finance* 41, 445–460.

Backwell, Alex, 2015, State prices and implementation of the Recovery Theorem, *Journal of Risk and Financial Management* 8, 2–16.

Black, Fischer, and Myron S. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–659.

Boehme, Rodney D., Bartley R. Danielsen, and Sorin M. Sorescu, 2006, Short-sale constraints, differences of opinion, and overvaluation, *Journal of Financial and Quantitative Analysis* 41, 455–487.

Borovička, Jaroslav, Lars Peter Hansen, and José A. Scheinkman, 2016, Misspecified recovery, *The Journal of Finance* 71, 2493–2544.

Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621–651.

Busch, Thomas, Bent Jesper Christensen, and Morten Ørregaard Nielsen, 2011, The role of implied volatility in forecasting future realized volatility and jumps in foreign exchange, stock, and bond markets, *Journal of Econometrics* 160, 48–57.

Carr, Peter, Hélyette Geman, Dilip P. Madan, and Marc Yor, 2002, Stochastic volatility for Lévy processes, *Mathematical Finance* 13, 345–382.

Carr, Peter, and Liuren Wu, 2003, What type of process underlies options? A simple robust test, *Journal of Finance* 58, 2581–2610.

Carr, Peter, and Jiming Yu, 2012, Risk, return, and Ross recovery, *The Journal of Derivatives* 20, 38–59.

Chung, Sung Gon, and Henock Louis, 2017, Earnings announcements and option returns, *Journal of Empirical Finance* 40, 220–235.

Cox, John C., 1975, Notes on option pricing I: Constant elasticity of variance diffusion, Working paper, Stanford University.

Derman, Emanuel, and Iraj Kani, 1994, Riding on a smile, *Risk* 7, 32–39.

Diavatopoulos, Dean, James S. Doran, Andy Fodor, and David R. Peterson, 2012, The information content of implied skewness and kurtosis changes prior to earnings announcements for stock and option returns, *Journal of Banking and Finance* 36, 786–802.

Diether, Karl B., Christopher J. Malloy, and Anna Scherbina, 2002, Differences of opinion and the cross-section of stock returns, *Journal of Finance* 57, 2113–2141.

Donders, Monique W. M., Roy Kouwenberg, and Ton C. F. Vorst, 2000, Options and earnings announcements: An empirical study of volatility, trading volume, open interest and liquidity, *European Financial Management* 6, 149–171.

Dubinsky, Andrew, Michael Johannes, Andreas Kaeck, and Norman J Seeger, 2019, Option pricing of earnings announcement risks, *The Review of Financial Studies* 32, 646–687.

Duong, Diep, and Norman R. Swanson, 2015, Empirical evidence on the importance of aggregation, asymmetry, and jumps for volatility prediction, *Journal of Econometrics* 187, 606–621.

Dupire, Bruno, 1994, Pricing with a smile, *Risk* 7, 18–20.

Ederington, Louis H., and Jae Ha Lee, 1996, The creation and resolution of market uncertainty: The impact of information releases on implied volatility, *Journal of Financial and Quantitative Analysis* 31, 513–539.

Eraker, Bjørn, 2004, Do stock prices and volatility jump? Reconciling evidence from spot and option prices, *Journal of Finance* 59, 1367–1403.

Eraker, Bjørn, Michael Johannes, and Nicholas Polson, 2003, The impact of jumps in volatility and returns, *Journal of Finance* 58, 1269–1300.

Heston, Steve, 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327–343.

Isakov, Dušan, and Christophe Pérignon, 2001, Evolution of market uncertainty around earnings announcements, *Journal of Banking and Finance* 25, 1769–1788.

Jensen, Christian Skov, David Lando, and Lasse Heje Pedersen, 2019, Generalized recovery, *Journal of Financial Economics* 133, 154–174.

Jiang, George J., Eirini Konstantinidi, and George Skiadopoulos, 2012, Volatility spillovers and the effect of news announcements, *Journal of Banking and Finance* 36, 2260–2273.

Jiang, George J., and Yisong S. Tian, 2010, Misreaction or misspecification? A re-examination of volatility anomalies, *Journal of Banking and Finance* 34, 2358–2369.

Kelly, Bryan, Ľuboš Pástor, and Pietro Veronesi, 2016, The price of political uncertainty: Theory and evidence from the option market, *Journal of Finance* 71, 2417–2480.

Kou, S. G., 2002, A jump-diffusion model for option pricing, *Management Science* 48, 1086–1101.

Lee, Suzanne S., 2012, Jumps and information flow in financial markets, *Review of Financial Studies* 25, 439–479.

Lei, Qin, Xuewu Wesley Wang, and Zhipeng Yan, 2020, Volatility spread and stock market response to earnings announcements, *Journal of Banking & Finance* 119, 105126.

Leung, Tim, and Marco Santoli, 2014, Accounting for earnings announcements in the pricing of equity options, *Journal of Financial Engineering* 1, 1–46.

Liu, Hong, Yingdong Mao, Xiaoxiao Tang, and Guofu Zhou, 2025, Ex-ante risk premia on earnings announcements: Evidence from the options market, Working paper, Washington University in St. Louis.

Lucas, Robert, 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.

Lucca, David O., and Emanuel Moench, 2015, The pre-FOMC announcement drift, *The Journal of Finance* 70, 329–371.

Madan, Dilip B., Peter P. Carr, and Eric C. Chang, 1998, The variance gamma process and option pricing, *European Finance Review* 2, 79–105.

Merton, Robert C., 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.

Neururer, Thaddeus, and George Papadakis, 2026, Skew premiums around earnings announcements, *Financial Review* (forthcoming).

Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.

Patell, James M., and Mark A. Wolfson, 1981, The ex ante and ex post price effects of quarterly earnings announcements reflected in option and stock prices, *Journal of Accounting Research* 19, 434–458.

Qin, Likuan, and Vadim Linetsky, 2016, Positive eigenfunctions of Markovian pricing operators: Hansen-Scheinkman factorization, Ross recovery, and long-term pricing, *Operations Research* 64, 99–117.

Ravagli, Lorenzo, Sandrine Ungari, and Julien Turc, 2014, Isolating a risk premium on the volatility of volatility: A general framework for short-volga strategies, Technical report, Société Générale Cross Asset Research.

Ross, Steve, 2015, The recovery theorem, *Journal of Finance* 70, 615–648.

Rubinstein, Mark, 1994, Implied binomial trees, *Journal of Finance* 69, 771–818.

Schneider, Paul, and Fabio Trojani, 2019, (Almost) model-free recovery, *The Journal of Finance* 74, 323–370.

Seo, Sang Byung, and Jessica A. Wachter, 2019, Option prices in a model with stochastic disaster risk, *Management Science* 65, 3449–3469.

Subramanian, Ajay, 2004, Option pricing on stocks in mergers and acquisitions, *Journal of Finance* 59, 795–829.

Truong, Cameron, Charles Corrado, and Yangyang Chen, 2012, The options market response to accounting earnings announcements, *Journal of International Financial Markets, Institutions and Money* 22, 423–450.

von Sydow, Lina, and Johan Walden, 2020, Numerical Ross recovery for diffusion processes using a PDE approach, *Applied Mathematical Finance* 27, 46–66.

Walden, Johan, 2017, Recovery with unbounded diffusion processes, *Review of Finance* 21, 1403–1444.

Xing, Yuhang, Xiaoyan Zhang, and Rui Zhao, 2010, What does the individual option volatility smirk tell us about future equity returns?, *Journal of Financial and Quantitative Analysis* 45, 641–662.