

What Were the Odds? Estimating the Market's Probability of Uncertain Events*

Ashley Langer[†] and Derek Lemoine[‡]

December 16, 2020

An event study generates only a lower bound on the full effect of an event unless researchers know the probability that investors assigned to the event before it occurred. We develop two model-free methods for recovering the market's priced-in probability of events. These methods require running event studies in financial options to complement the standard event study in stock prices. Validating both approaches, we estimate that the 2016 U.S. election outcome had a 12% chance of occurring. This probability is consistent with contemporary polling, bookmaker, and prediction market estimates. Demonstrating the usefulness of our approaches, we show that many OPEC meetings' outcomes were well-anticipated. OPEC retained substantial influence on world oil prices even as the U.S. increased oil production.

JEL: C58, G13, L71, Q43

Keywords: event study, anticipation, policy uncertainty, options, derivatives, oil, cartel, OPEC

*We are grateful to David Brown, David Card, Gautam Gowrisankaran, Hidehiko Ichimura, Diego Kaenzig, David Schreindorfer, and Tiemen Woutersen for feedback, to Justin Wolfers for conference discussion, to the University of Arizona's Eller Center for Management Innovations in Healthcare and Institute of the Environment for financial support, and to Wendan Zhang and Frances Slater for research assistance. We also thank seminar participants at Arizona State University, ASSA 2020, HEC Montreal, TREE, University of Arizona, University of California Berkeley, and Yale. We gratefully acknowledge an allocation of computer time from the University of Arizona's High Performance Computing (HPC) facility.

[†]University of Arizona and NBER; alanger@email.arizona.edu

[‡]University of Arizona and NBER; dlemoine@email.arizona.edu

1 Introduction

Event studies are standard techniques for estimating financial markets' response to news releases. However, it is widely recognized that event studies capture the full market valuation of news only when that news is a complete surprise.¹ Such cases are rare: few things that happen were previously judged to have zero probability. When stock prices partially anticipate events, event studies measure the effect of becoming sure about the news rather than of learning news that is completely new. The event study methodology then provides only a lower bound on the consequences of news—and that bound can be arbitrarily loose.

In this paper we develop two new model-free techniques for estimating the market probability of realized events from widely traded financial options. We validate these techniques on the 2016 U.S. election and then apply them to a series of announcements by the Organization of the Petroleum Exporting Countries (OPEC) oil market cartel. Our new methods enable researchers to recover the full effect of events, which is important for applications ranging from minimum wage studies (e.g., Card and Krueger, 1995; Bell and Machin, 2017) to taxation (e.g., Auerbach and Hassett, 2007) to health policy (e.g., Al-Ississ and Miller, 2013) to environmental policy (e.g., Bushnell et al., 2013). Our new methods are also important because the probabilities of the events can themselves be of direct interest. For instance, we can estimate firms' differential exposure to uncertainty, which could be critical to understanding the role of uncertainty in the economy.

The intuition underlying our first approach is straightforward. Imagine that a financial option has value only if a given event occurs: there is some chance of exercising the option conditional on the event occurring but very little chance of exercising it otherwise. In an oil market application, this might be an option that gives the holder the right to purchase oil at a high “strike” price after an OPEC decision on whether to cut production. If OPEC cuts production, then it is likely that the market price of oil will end up above the strike price, making the option valuable. But if OPEC does not cut production, then the market price of oil is highly unlikely to be greater than the strike price, making the option nearly worthless. On the day before the event, the value of this option is approximately the probability of the event occurring times the value of the option if the event occurs. On the

¹Card and Krueger (1995)[314] explain that “one difficulty in interpreting” their finding that minimum wages have only small effects on stock prices is “the fact that investors might have anticipated the news before it was released”. MacKinlay (1997, 37) laments that while event studies are in principle a promising tool for recovering “the wealth effects of regulatory changes for affected entities”, their usefulness has been limited by the fact that “regulatory changes are often debated in the political arena over time”, with their effects incorporated into stock prices only gradually.

day after the event occurs, the value of this option is simply the value of the option given that the event has occurred. If nothing else changed over the event window, the ratio of the option's price before the event to its price after the event is the priced-in probability of the event occurring. By running an event study in option prices, researchers can estimate what the change in an option's price would have been if nothing but the event had occurred. This approach requires that some liquidly traded options would have had little value if the event had turned out differently but retained value after the realized event. In practice, this approach requires an event to be sufficiently important for the price of an option's underlying asset and requires that sufficiently deep out-of-the-money options (i.e., options with extreme strike prices) are liquidly traded.²

Our second method of estimating an event's probability relies on different identifying assumptions and does not rely on deep out-of-the-money options. If the market knows that an event will occur on a given date and if a stock's volatility in the days and weeks immediately after the event does not depend directly on the event's outcome, then the event-induced change in the stock's expected near-term variance can be used to recover the probability of the event. For instance, traders typically know the day that OPEC will meet and might expect that the supply and demand shocks that will induce price volatility in the few weeks immediately after the meeting are independent of the meeting's outcome. Under minimal assumptions, the price of a portfolio of options reveals the market's expectation of the variance in a stock's price (Martin, 2017). We recover the market's probability of events by estimating the event-induced change in the price of this portfolio. Intuitively, differencing the pre- and post-event expected variances eliminates the post-event variance but retains the variance induced by uncertainty about the event's realization. This latter variance is merely a transformation of the probability we seek.

We validate our methods by estimating the probability of the Republican sweep of the 2016 U.S. election. Prediction markets, bookmakers, and polling-driven models imply probabilities ranging from 0.07 to 0.26. We recover a probability of 0.12. Our two methods generate nearly identical results, despite relying on different identifying assumptions. And the empirical results support the theoretical predictions: we demonstrate that the out-of-the-money option approach requires using only extreme strikes and restricting attention to firms with especially large event-day stock price

²In general, options with extreme strike prices would not be worthless if the event fails to occur because there may still be some chance of reaching an extreme stock price. We show that the estimated probability is then an upper bound on the market's priced-in probability of the event occurring. We describe theoretically motivated restrictions designed to generate a tight bound and find that this bound does appear to be tight in our application to the 2016 U.S. election.

movements, and we demonstrate that the variance swap approach requires removing firms that were not exposed to the election outcome. Recent work has been interested in the implications of event studies of this election (e.g., Mukanjari and Sterner, 2018; Wagner et al., 2018a,b; Ramelli et al., 2019). Our results imply that the full effect of the election is 15% larger than implied by standard event study estimates.

We show that estimating these probabilities can be critical for understanding economic forces. In particular, we estimate probabilities for 13 OPEC meetings from 2011 to 2016. Conventional event studies show that OPEC meetings largely failed to move oil prices between 2012 and 2016, which coincided with a period of sustained oil production quotas by OPEC member countries. However, previous work has speculated that similar null results may merely reflect that anticipation of OPEC announcements masks the true influence of OPEC (e.g., Draper, 1984; Deaves and Krinsky, 1992; Wirl and Kujundzic, 2004; Spencer and Bredin, 2019). Indeed, we find that most meetings did not produce news that was fairly surprising. The few events that produced surprising news did have large effects on oil markets. In particular, markets anticipated OPEC's failure to reach agreement on new production quotas in December 2015 and June 2016. Standard event studies suggest that these meetings did not strongly affect oil markets, but we find that these were in fact two of the more important meetings after adjusting for market anticipation.

The usefulness of priced-in event probabilities

Recovering event probabilities is widely acknowledged to be important for measuring the full magnitude of events and for understanding whether events are unimportant or just anticipated (e.g., Card and Krueger, 1995; Karpoff et al., 1996; Hallock, 1998). Measuring the full magnitude of events is critical for testing theoretical predictions (e.g., Hennessy and Strebulaev, 2020), conducting cost-benefit analysis (e.g., Meng, 2017), and estimating damages in court settings (e.g., Cornell and Morgan, 1990).

Many other lines of economic inquiry would also benefit from credibly estimated probabilities. First, substantial recent research has attempted to measure the level and impact of policy uncertainty (e.g., Baker et al., 2014; Dorsey, 2019; Bianconi et al., 2020). Our methods recover new measures of this uncertainty. Further, our methods can identify how exposure to uncertainty varies over firms, as when firms are exposed to different aspects of an election, court ruling, or legislation. Recent work has used linguistic analysis to connect firm-level uncertainty to decisions on hiring, investment, and lobbying decisions (Hassan et al., 2019; Handley and Li, 2020). Our methods are a valuable complement to this textual work both in relying on preferences revealed through transactions and in measuring exposure to uncertainty

about more narrowly defined events.

Second, some research seeks to aggregate or compare different events over time in a way that is similar to our OPEC analysis (e.g., Bronars and Deere, 1990; Farber and Hallock, 2009; Kogan et al., 2017). Without knowing the underlying probability of each event, it is impossible to distinguish between events that have different effects (whether over time or across firms) and events that are better-anticipated under some conditions than others. Researchers have long noted that cross-sectional analyses could be severely biased—even to the point of estimating the wrong sign—when the market’s ability to forecast events varies with firm characteristics (e.g., Lanen and Thompson, 1988; MacKinlay, 1997; Bhagat and Romano, 2002). Our methods enable future studies to control for the priced-in probability of the event.

Third, researchers have used close elections as randomized experiments (e.g., Lee, 2008), but Caughey and Sekhon (2011) show that the outcomes of close elections may not be random. Our methods allow researchers to identify the elections that market participants viewed as effectively random,³ allowing researchers to use only those elections that were truly unexpected as randomized experiments.

Finally, researchers are interested in the risk premia placed on different states of the world, determined by variation in the stochastic discount factor. Our methods allow for new means of identifying how the stochastic discount factor varies with event outcomes. We recover risk-neutral probabilities, which reweight “objective” or “physical” probabilities by marginal utility and are the probabilities needed to correct event study estimates.⁴ If researchers use other approaches to pin down physical probabilities, then these events’ risk-neutral probabilities inform us about investors’ expectations of consumption in the realized state versus in other possible states.⁵

Previous approaches to correcting event studies for partial anticipation

Researchers have long recognized the need to measure or control for the likelihood of an event (e.g., Ball, 1972; Binder, 1985; Malatesta and Thompson, 1985; MacKinlay,

³This identification requires that the elections be too “small” to bear much of a risk premium (so that the estimated risk-neutral probabilities roughly correspond to physical probabilities) but be important enough to affect some firms’ stock prices.

⁴Risk-neutral probabilities differ from objective probabilities when there is a risk premium. The existence or magnitude of risk premia is immaterial to the present paper.

⁵For instance, if some elections are decided by a coin flip (e.g. Virginia’s 94th District in 2017, which maintained the Republican majority in the House of Delegates) or are shown to be effectively random through the types of balance tests described by Caughey and Sekhon (2011), then the risk-neutral probability of these elections tells us which candidate or party was anticipated to have more favorable consequences for aggregate consumption.

1997; Lamdin, 2001). Previous literature has adopted several approaches. The simplest solution is to select events that the researcher judges to be relatively surprising. But this assumption can be difficult to test and limits the applicability of the event study approach.⁶ A related stream of research attempts to reduce the effects of partial anticipation by extending the event window to include earlier time periods in the hope that the extended event window will capture any information leakage (e.g., Jayachandran, 2006; Auerbach and Hassett, 2007; Linn, 2010; Lee and Mas, 2012; Al-Ississ and Miller, 2013). In some cases, the event window becomes years-long. However, extending the event window decreases the plausibility of the event study methodology's identifying assumption that no information is revealed in the event window other than the event itself and also reduces the power to detect true effects (Brown and Warner, 1985; Kothari and Warner, 1997). For these and other reasons, many recommend keeping the event window as short as possible (e.g., Bhagat and Romano, 2002; Kothari and Warner, 2007).

Instead of trying to minimize the market probability of an event, two strands of research seek to recover that probability directly. First, many researchers use prediction market contracts to measure event probabilities (e.g., Roberts, 1990; Heron, 2000; Hughes, 2006; Knight, 2006; Snowberg et al., 2007; Lange and Linn, 2008; Wolfers and Zitzewitz, 2009; Imai and Shelton, 2011; Snowberg et al., 2011; Lemoine, 2017; Meng, 2017). Prediction markets can be a valuable source of information when the proper contracts exist, but this method faces significant hurdles in many applications. In particular, prediction market contracts are not available for many events of interest or may be written on only part of an event (e.g., the outcome of a presidential election but not of each down-ballot race). Further, prediction market contracts can be quite thinly traded, rapidly fluctuating prices can be hard to map to end-of-day stock market data, and prediction market participants might not hold the same beliefs as financial market participants. Prediction markets are powerful sources of information, but it is important to have other means of estimating event probabilities.

We show that bets made in ordinary options markets can encode the same information as specialized prediction markets. Options markets are attractive because they can be quite thickly traded, are closely linked to stock markets of interest in event studies, and may be more informationally efficient than stock markets (Black,

⁶Dube et al. (2011) find that even top-secret coup authorizations leak to the markets. Auerbach and Hassett (2007) combine estimates from several legislative events in the hope that the cumulative effect of incrementally increasing the probability of successful legislation approximates the effect of truly surprising legislation. Of course, some researchers do claim that their events were complete surprises (e.g., Bell and Machin, 2017). Our methods allow them to test such claims.

1975; Easley et al., 1998; Pan and Poteshman, 2006; Johnson and So, 2012). A few papers in the finance literature have previously inferred the priced-in probability of events from the prices of financial options (Gemmill, 1992; Barraclough et al., 2013; Borochin and Golec, 2016; Carvalho and Guimaraes, 2018). These papers assume that options are priced according to specific parametric models and search for the event probability that reconciles observed option prices and theoretical option prices. However, as recognized by these authors, the distributional assumptions about stock prices that underpin these parametric models are commonly violated in practice.⁷ For instance, discrepancies between actual option prices and theoretical option prices generate well-known “anomalies” such as implied volatility smirks and smiles. When a theoretical model does not correctly predict option prices, including a probabilistic event adds at least one additional parameter that can improve the fit to observed option prices even if there were in fact no chance of an event.

Our new methods retain the advantages of using options markets but do not impose parametric assumptions.⁸ They are therefore not vulnerable to conventional option pricing “anomalies”. Our methods for estimating the event probability require only the absence of arbitrage and either (i) that some out-of-the-money options have nontrivial value following the realized event outcome but not following other possible outcomes or (ii) that expected volatility in the days after an event does not depend on the event’s realization. The latter is among the assumptions imposed in Gemmill (1992) and Carvalho and Guimaraes (2018), so our second method can be viewed as generalizing prior methods to eliminate parametric assumptions.⁹

⁷In fact, each of these papers uses a different option pricing model. Their various assumptions conflict with each other.

⁸Some other finance literature is more loosely related to the present paper. Gürkaynak et al. (2020) show how heteroskedastic responses to events identify non-headline news. Kelly et al. (2016) use financial options to estimate how much news is likely to be released by upcoming events. Whereas they seek the spread of possible outcomes, we seek the probability of the realized outcome, and whereas they need the date at which news will be released to be known well in advance (they study national elections and global summits), we propose a method that allows the date to be unknown in advance. Acharya (1993) studies the probability of endogenous events such as corporate announcements. We focus on policy events that affect a cross-section of firms and are not endogenous to any one firm. van Tassel (2016) and Fisman and Zitzewitz (2019) use stock and option prices, respectively, to recover investors’ post-event beliefs. We recover investors’ pre-event beliefs. Finally, Martin (2017), among others, shows how to use variation in option prices across strikes to recover the probability of sufficiently large losses. We show how to use variation over time to recover the probability of realized events.

⁹In concurrent work, Grinblatt and Wan (2020) argue that one can in principle back out risk-neutral probabilities from the prices of options traded before an event. Their approach requires the full state space to be specified. In contrast, we use time series variation in option prices to recover the risk-neutral probability of the realized event without needing to explicitly specify the possible

Outline

The next section describes the setting and defines the bias present in standard event studies. Section 3 derives the two new approaches to recovering an event’s probability from options data. Section 4 explains how we take these theoretical approaches to the data, and Section 5 recovers probabilities for the 2016 U.S. election and for OPEC meetings. Section 6 concludes. The appendix contains supporting empirical analysis, theoretical extensions, and formal derivations.

2 Setting

We begin by developing a generic reduced-form representation of stock prices around an event. Consider a case with only a single firm. Its stock price reflects information about an upcoming event and all information about the state of the world beyond the event (such as earnings forecasts and interest rates). Formally, the price of the firm’s stock is $S(\omega_t, e)$, where $\omega_t \in \mathcal{R}^N$ captures all payoff-relevant non-event news and $e \in \{L, H\}$ captures all event-related news.¹⁰ We write S_t for the observed stock price, S_t^L for $S(\omega_t, L)$, and S_t^H for $S(\omega_t, H)$.

We study the beliefs of a representative market investor. Investors know ω_t at all times t , but they do not know whether $e = L$ or $e = H$ until time τ .¹¹ At all times after τ , observed stock prices S_t are either S_t^H or S_t^L . Let the time t representative agent assign risk-neutral probability p_t^H to $e = H$ and risk-neutral probability $p_t^L = 1 - p_t^H$ to state $e = L$.¹² Assume that event uncertainty affects payoffs only through the realized event e (i.e., p_t^H is independent of ω_t), as when changes in event uncertainty do not lead firms to take actions that directly change their exposure to other shocks.

states or the form of the event’s effect.

¹⁰We write the discrete component as binary, but this choice is not restrictive. If, for instance, outcome H is realized and extreme, then we can aggregate all of the other possible outcomes into a single indicator L and treat $S(\omega_t, L)$ as an expected value over these outcomes. For the derivation when the event is not extreme, see Appendix C.2

¹¹Until Section 3.2, we do not specify whether investors know the value of τ in advance.

¹²Absence of arbitrage ensures the existence of a risk-neutral measure, and the risk-neutral measure is unique if markets are complete. The risk-neutral measure can be interpreted as embedding risk into the probability weights by adjusting “physical” probabilities for the representative agent’s risk aversion. For more on risk-neutral pricing, see standard asset pricing texts such as Björk (2004) or Cochrane (2005).

2.1 The Bias in the Standard Event Study Methodology

Before we discuss how to recover event probabilities, we now review how partially anticipated events challenge the standard event study methodology. Without loss of generality, assume that event H occurs. An event study aims to recover the full effect of the event, $S_{\tau-1}^H - S_{\tau-1}^L$, where $S_{\tau-1}^H$ and $S_{\tau-1}^L$ indicate what the time $\tau - 1$ stock price would have been if the upcoming event's outcome were known to be H or L , respectively. In order to estimate this difference, researchers regress stock returns ($S_{\tau}/S_{\tau-1}$) on controls and an indicator for the first day (τ) on which the event outcome is known. The identifying assumption is that the controls account for anything other than the event that occurs between time $\tau - 1$ and time τ (i.e., for the difference between ω_{τ} and $\omega_{\tau-1}$). If there are multiple firms, then the identifying assumption becomes that, once the researcher includes controls, the expected value across firms of all residual non-event stock price news between $\tau - 1$ and τ is zero.¹³

If the identifying assumption is satisfied, then the standard approach recovers $S_{\tau-1}^H - S_{\tau-1}$, the ceteris paribus change in the stock price at the time of the event. In the absence of arbitrage, the observed time $\tau - 1$ stock price must be the probability-weighted average of the stock price if event L occurs and the stock price if event H occurs:

$$S_{\tau-1} = p_{\tau-1}^L S_{\tau-1}^L + p_{\tau-1}^H S_{\tau-1}^H. \quad (1)$$

Rearranging and adding $S_{\tau-1}^H$ to both sides yields:

$$S_{\tau-1}^H - S_{\tau-1} = (1 - p_{\tau-1}^H) [S_{\tau-1}^H - S_{\tau-1}^L]. \quad (2)$$

As is well known (e.g., Snowberg et al., 2011), the estimated event effect ($S_{\tau-1}^H - S_{\tau-1}$) is less than the full event effect ($S_{\tau-1}^H - S_{\tau-1}^L$) because some of the event outcome is already “priced in” before the event occurs. Researchers recover the true value of the event from the change in stock prices only when the event outcome is completely unexpected at $\tau - 1$ (i.e., only as $p_{\tau-1}^H \rightarrow 0$). For this reason, researchers have sought events that are surprises. Unfortunately, an event study measures only an arbitrarily small fraction of the true event effect as the probability of the realized event outcome becomes large (i.e., as $p_{\tau-1}^H \rightarrow 1$). Event studies provide only a lower bound—and potentially only a very loose lower bound—for the implications of an event in the absence of information about $p_{\tau-1}^H$.

¹³See Campbell et al. (1997, Chapter 4), MacKinlay (1997), and Kothari and Warner (2007), among others, for reviews of event study methods.

3 Two Model-Free Approaches to Recovering the Event Probability from Options Data

We propose two new approaches to recovering $p_{\tau-1}^H$ from time series variation in options prices. The first approach identifies this probability from changes in tail probabilities, and the second approach identifies the probability from changes in a stock's expected variance. We formally derive each approach in this section before detailing how we empirically implement each approach in Section 4.

3.1 Using the Change in Tail Probabilities

A call (put) option on a stock confers the right—but not the obligation—to buy (sell) the stock at a defined “strike” price K on a defined expiration date T .¹⁴ A call option's value derives from the chance that the underlying stock's price will be higher than the strike price at the expiration date, in which case the option holder can buy the stock at the strike price, sell it at the market price, and keep the difference. If, on the other hand, the stock price at the expiration date is less than the strike price, the option holder should allow the option to expire unexecuted. Thus, the date x value of the call option, $C_{x,T}(S_x, K)$, is the expected difference between the stock price and the strike conditional on that difference being positive:

$$C_{x,T}(S_x, K) = \frac{1}{R_{x,T}} \int_K^\infty (S_T - K) [[1 - p_x^H]f_x(S_T|S_x^L, L) + p_x^H f_x(S_T|S_x^H, H)] dS_T, \quad (3)$$

where we explicitly condition the risk-neutral distribution $f_x(\cdot)$ of S_T on the time x stock price S_x and event outcome e and where $R_{x,T} \geq 1$ is the gross risk-free rate from time x to T . A put option is priced analogously, recognizing that its value derives from the chance that the underlying stock's price will be less than the strike price at the expiration date.

Now consider the effect of an event that occurred at time τ . From the previous expression, the time $\tau - 1$ price of a call option is

$$C_{\tau-1,T}(S_{\tau-1}, K) = p_{\tau-1}^H C_{\tau-1,T}^H(S_{\tau-1}^H, K) + (1 - p_{\tau-1}^H) C_{\tau-1,T}^L(S_{\tau-1}^L, K), \quad (4)$$

¹⁴We model options as “European”, even though most traded options are “American” options that allow the holder to exercise the option before T . This distinction is unlikely to be quantitatively important. Appendix C.1 extends the analysis to American options, showing that the results converge to the case of a European option as the time to maturity shrinks. In the empirical application, we focus on options with the shortest time to maturity, drop firms with high dividend yields that might make early exercise attractive, and, for other reasons given below, emphasize call options. These restrictions combine to limit the early exercise premium and thus to limit any difference between the prices of American and European options.

where $C_{\tau-1,T}^e(S_{\tau-1}^e, K)$ is the value of the option at time $\tau - 1$ if it was already known that the event's outcome would be e . Assume, without loss of generality, that outcome H was realized at time τ . We can then estimate $C_{\tau-1,T}^H(S_{\tau-1}^H, K)$ under standard event study assumptions. Dividing both sides of equation (4) by $C_{\tau-1,T}^H(S_{\tau-1}^H, K)$, we have:

$$\frac{C_{\tau-1,T}(S_{\tau-1}, K)}{C_{\tau-1,T}^H(S_{\tau-1}^H, K)} = p_{\tau-1}^H + (1 - p_{\tau-1}^H) \frac{C_{\tau-1,T}^L(S_{\tau-1}^L, K)}{C_{\tau-1,T}^H(S_{\tau-1}^H, K)} \triangleq \bar{p}. \quad (5)$$

The ratio of the observed and counterfactual pre-event option prices is equal to the risk-neutral probability of the event outcome plus a “bias” term that depends on $C_{\tau-1,T}^L(S_{\tau-1}^L, K)$, the unobserved value of the option if the counterfactual event outcome had occurred. We label the ratio \bar{p} . Standard arbitrage bounds (e.g., Cochrane, 2005) require $C_{\tau-1,T}^L(S_{\tau-1}^L, K) \geq 0$ and $C_{\tau-1,T}^H(S_{\tau-1}^H, K) \geq 0$. Therefore \bar{p} is at least as large as the true event probability $p_{\tau-1}^H$ and approaches this probability as the value of the option under the counterfactual event outcome approaches 0.

The bias term must be small if \bar{p} is to be useful for recovering $p_{\tau-1}^H$. Assume that the event increases stock prices in sufficiently extreme states of the world: there exists some \bar{S} such that $S(\omega_t, H) > \bar{S}$ implies $S(\omega_t, H) > S(\omega_t, L)$.¹⁵ Given that infinite prices are not possible, the conditional distributions $f_x(S_T|S_x^L, L)$ and $f_x(S_T|S_x^H, H)$ converge to zero as S_T grows. $C_{\tau-1,T}^L(S_{\tau-1}^L, K)$, the value of the call option conditional on outcome L , must converge to zero in the strike price K more quickly than does $C_{\tau-1,T}^H(S_{\tau-1}^H, K)$. Equation (5) then shows that \bar{p} converges to $p_{\tau-1}^H$ for some K that is sufficiently large but not so large as to drive $C_{\tau-1,T}^H(S_{\tau-1}^H, K)$ to zero. Finally, conditional on a given event probability and a given strike, the difference $C_{\tau-1,T}^H(S_{\tau-1}^H, K) - C_{\tau-1,T}^L(S_{\tau-1}^L, K)$ is likely to be larger when the event has a larger effect on the underlying stock price. This means that the bias in equation (5) is likely to be small for a broader range of strikes when the event increases the stock price by a large amount.

Figure 1 depicts the intuition. Imagine that event H leads to high stock prices and event L leads to low stock prices. Conditional on the information available just before the event occurs (at time $\tau - 1$), the risk-neutral distribution of prices for the stock at the expiration date T is the dashed line $f_{\tau-1}(S_T|S_{\tau-1})$, which is a mixture of the distribution conditional on event H occurring (the long-dashed

¹⁵This assumption does not require that the event can have only two possible outcomes. Instead, this assumption requires that the realized outcome is extreme: if we define L to indicate a set of outcomes $\{L_1, \dots, L_N\}$, then this assumption requires that $S(\omega_t, H) > \bar{S}$ implies $S(\omega_t, H) > \sum_{i=1}^N [p_{\tau-1}^{L_i}/(1 - p_{\tau-1}^H)]S(\omega_t, L_i)$. Appendix C.2 relaxes the assumption that the realized outcome is extreme.




Figure 1: Illustration of how changes in option prices identify the probability $p_{\tau-1}^H$ when event H is realized at time τ .

line $f_{\tau-1}(S_T | S_{\tau-1}^H, H)$ and the distribution conditional on the counterfactual event occurring (the solid line $f_{\tau-1}(S_T | S_{\tau-1}^L, L)$). The value of a call option with strike K is the discounted expected value of the difference between the time T stock price and K , conditional on the stock price being in the area labeled A. Once the event occurs, the density of stock prices at the expiration date becomes the dotted line $f_{\tau-1}(S_T | S_{\tau-1}^H, H)$ and the option's value is calculated as the integral over the larger area $A + B$. If, as depicted, the distribution $f_{\tau-1}(S_T | S_{\tau-1}^L, L)$ contains very little mass above K , the jump in the option price between date $\tau - 1$ and τ identifies the extent to which the distribution of stock prices conditional on the realized outcome ($f_{\tau-1}(S_T | S_{\tau-1}^H, H)$) was downweighted by the probability $p_{\tau-1}^H$ of that outcome.

In Figure 1, the bias of this approach is very close to zero for the indicated K . For smaller K , $f_{\tau-1}(S_T | S_{\tau-1}^L, L)$ may have nontrivial mass in region A. In this case, the change in the price of the option reflects both the resolution of uncertainty surrounding the event and the loss of this unobserved probability mass. The possibility of unobserved probability mass drives the bias term in equation (5).

Equation (5) further shows that the bound \bar{p} will be especially tight when $p_{\tau-1}^H$ is large. This is precisely the case in which standard event studies suffer arbitrarily large biases (see equation (2)) and therefore is the case in which a tight bound is most needed. The bias vanishes as $p_{\tau-1}^H$ grows both because $1 - p_{\tau-1}^H$ shrinks and because $S_{\tau-1}^L$ shrinks (for given observables $S_{\tau-1}$ and $S_{\tau-1}^H$). In Figure 1, large

$p_{\tau-1}^H$ corresponds to a case in which the distribution conditional on L receives little weight. The long-dashed distribution then has little effect on the price of the option just before the event occurs.

Put options can also recover the event probability, which is useful when an event reduces the price of the underlying asset. Trivially adapting equations (3) through (5), \bar{p} is again the ratio of the option price the day before the event to the option price the day before the event if the event outcome were already known. The bias is small if $P_{\tau-1,T}^L(S_{\tau-1}^L, K)$, the value of the put option under the counterfactual outcome, is small. However, now this condition is more likely to hold when an event reduces the stock price.¹⁶

Put options are therefore useful when an event reduces the price of an option's underlying asset, while call options are useful when an event increases the price of an option's underlying asset. However, much work has conjectured that out-of-the-money put options carry a premium because they offer protection against low-probability crashes.¹⁷ The possibility of such disasters is plausibly independent of the event outcome. In that case, deep out-of-the-money put options may retain much of their value even if event outcome L occurs. The bias in a version of equation (5) using puts therefore does not converge to zero as K becomes small.¹⁸ Further, if the value of the deepest out-of-the-money put options is primarily driven by disaster risk that is independent of the event outcome, then the bias from estimating $p_{\tau-1}^H$ may actually increase as the strike price falls. It is no longer clear which strikes should provide the tightest bound.

Estimating \bar{p} may therefore be most informative when the realized event increases

¹⁶Again applying standard arbitrage bounds, we have $\bar{p} \geq p_{\tau-1}^H$. If event H decreases stock prices in sufficiently extreme states of the world (i.e., if there exists \underline{S} such that $S(\omega_t, L) < \underline{S}$ implies $S(\omega_t, H) < S(\omega_t, L)$) and if, in addition, sufficiently low prices would not be possible following event L (i.e., if $f_{\tau-1}(S_T | S_{\tau-1}^L, L) \rightarrow 0$ as S_T approaches some small positive value), then the arbitrage bound $P_{\tau-1,T}^L(S_{\tau-1}^L, K) \geq 0$ holds exactly as K becomes small. In that case, the bias in the version of equation (5) with put options converges to zero for some sufficiently small K . The ratio of the observed and counterfactual put option prices at time $\tau - 1$ then accurately recovers $p_{\tau-1}^H$.

¹⁷Since the 1987 stock market crash, out-of-the-money put options on the S&P index have carried a premium (identified via the implied volatility “smirk”) reflecting an implied risk-neutral distribution that heavily weights the possibility of a crash (e.g., Rubinstein, 1994; Jackwerth and Rubinstein, 1996; Bates, 2000). Kelly et al. (2016) find that the crash or tail-risk premium can become especially large around political events, such as the elections we consider in our applications below. Others have explored whether the possibility of rare disasters can explain the equity premium puzzle (e.g., Rietz, 1988; Barro, 2006; Barro and Ursúa, 2012).

¹⁸With respect to footnote 16, it may in fact not be true that $f_{\tau-1}(S_T | S_{\tau-1}^L, L) \rightarrow 0$ as S_T approaches some small positive value.

a firm's stock price. In that case, researchers can estimate the event's risk-neutral probability from call options, for which the bias is more likely to decrease monotonically in observed strike prices and plausibly vanishes for sufficiently large strike prices.¹⁹

3.2 Using the Change in Expected Variance

The previous method of estimating the priced-in event probability $p_{\tau-1}^H$ relied on changes in the tail of the distribution of S_T , as reflected in option prices. We now derive a second method of estimating the priced-in probability, using changes in the expected variance of stock prices. This method identifies the event probability under a different set of circumstances, now requiring advance knowledge of the date that the event will happen and requiring that the expected variance of the stock price in the days immediately after an event be independent of the event's realization.

Assuming that market participants know that the event will occur at time τ , the time $\tau - 1$ variance of the time τ stock price includes both the variance coming from the resolution of the event and the variance coming from the evolution of ω_τ . Conditioning on ω_τ , the variance generated by the event itself is:

$$Var_{\tau-1}[S_\tau|\omega_\tau] = p_{\tau-1}^H(S_{\tau-1}^H - S_{\tau-1})^2 + (1 - p_{\tau-1}^H)(S_{\tau-1}^L - S_{\tau-1})^2.$$

Using equation (1), the variance generated by the event itself can be written as:

$$Var_{\tau-1}[S_\tau|\omega_\tau] = \frac{p_{\tau-1}^H}{1 - p_{\tau-1}^H}(S_{\tau-1}^H - S_{\tau-1})^2, \quad (6)$$

This conditional variance is a function only of the event-induced jump in stock prices and the market's probability of the event as of $\tau - 1$. This suggests that we can recover $p_{\tau-1}^H$ if we can estimate the variance generated by the event.

Now consider the unconditional variance of stock prices from time $\tau - 1$ to time τ . Temporarily fixing the one-day gross risk-free rate to $R_{\tau-1,\tau} = 1$, Appendix D.1 shows that

$$Var_{\tau-1}[S_\tau] = Var_{\tau-1}[S_\tau|\omega_\tau] + p_{\tau-1}^H Var_{\tau-1}[S_\tau^H] + (1 - p_{\tau-1}^H) Var_{\tau-1}[S_\tau^L]. \quad (7)$$

¹⁹Of course, the deeper out-of-the-money call options' value could be driven by unusually favorable circumstances that make the event irrelevant, but it is more likely that an event will move stock prices in both slightly and strongly favorable scenarios than that an event will move stock prices in both slightly and severely unfavorable scenarios.

The first term on the right-hand side is the conditional variance from equation (6), generated by resolution of the event uncertainty. The second and third terms capture the expectation of the variance induced by the resolution of ω_t .

Subtracting $Var_{\tau-1}[S_\tau^H]$ from each side of equation (7), substituting from equation (6), and rearranging, we find:

$$\frac{p_{\tau-1}^H}{1 - p_{\tau-1}^H} = \frac{Var_{\tau-1}[S_\tau] - Var_{\tau-1}[S_\tau^H]}{(S_{\tau-1}^H - S_{\tau-1})^2} + (1 - p_{\tau-1}^H) \frac{\overbrace{(Var_{\tau-1}[S_\tau^H] - Var_{\tau-1}[S_\tau^L])}^{\Delta Var}}{(S_{\tau-1}^H - S_{\tau-1})^2}. \quad (8)$$

This formulation relates the event probability on the left-hand side to a function of the event's effect on stock prices and its effect on the variance of the stock price, both of which are in the first term on the right-hand side and can be estimated. The second term on the right-hand side is proportional to the difference in the variance of ω_t under realized and counterfactual events, which we label ΔVar . If the event outcome does not affect the variance of the stock price, then ΔVar is equal to zero and this second term vanishes. In that case, the change in variance at the event and the change in the stock price at the event together identify the priced-in event probability.

If we had a portfolio of options that could replicate $Var_{\tau-1}[S_\tau]$, then we could potentially construct a similar replicating portfolio for $Var_{\tau-1}[S_\tau^H]$ by applying event study methods and then use equation (8) to recover an estimate of $p_{\tau-1}^H$. Martin (2017) provides a critical result. He constructs the replicating portfolio for a related object, a “simple variance swap”.²⁰ The variance strike $V_{\tau-1,T}$ that sets the value of a simple variance swap to zero is:

$$V_{\tau-1,T} = E_{\tau-1} \left[\sum_{j=0}^{T-\tau} \left(\frac{S_{\tau+j} - S_{\tau+j-1}}{\tilde{R}_{\tau-1,\tau+j-1} S_{\tau-1}} \right)^2 \right],$$

where expectations are, as elsewhere, taken under the risk-neutral measure and where $\tilde{R}_{t,y}$ is the net-of-dividend gross rate from time t to y .²¹ Martin (2017) prices the simple variance swap under the assumptions of a constant interest rate, a constant

²⁰The long position in a variance swap pays a fixed amount (the “strike”) at some future time T in exchange for payments linked to the realized variance of a stock's price between times t and T . The time t variance swap rate is the strike that sets the value of the swap to 0 at time T . This strike is equal to the risk-neutral expected variance between times t and T .

²¹Note that $\tilde{R}_{\tau-1,y} S_{\tau-1}$ is the time $\tau - 1$ forward price of S_y .

dividend rate, and small timesteps, without assuming away the possibility of jumps.²² Martin (2017) shows that

$$V_{\tau-1,T} = \frac{2R_{\tau-1,T}}{[\tilde{R}_{\tau-1,T}S_{\tau-1}]^2} \left\{ \int_0^{\tilde{R}_{\tau-1,T}S_{\tau-1}} P_{\tau-1,T}(S_{\tau-1}, K) dK + \int_{\tilde{R}_{\tau-1,T}S_{\tau-1}}^{\infty} C_{\tau-1,T}(S_{\tau-1}, K) dK \right\}. \quad (9)$$

The variance strike $V_{\tau-1,T}^H$ for a conditional simple variance swap can be recovered using equation (9) and option prices that are estimated to hold under the realized outcome H .

The following proposition relates the event probability $p_{\tau-1}^H$ to the values of simple variance swaps $V_{\tau-1,T}$ and $V_{\tau-1,T}^H$:

Proposition 1. *Define*

$$\tilde{V} \triangleq (S_{\tau-1})^2 V_{\tau-1,T} - (S_{\tau-1}^H)^2 V_{\tau-1,T}^H, \quad \tilde{p} \triangleq \frac{\tilde{V}}{\tilde{V} + [2\tilde{R}_{\tau-1,\tau} - 1] [S_{\tau-1}^H - S_{\tau-1}]^2}. \quad (10)$$

Then:

1. $\tilde{p} \rightarrow p_{\tau-1}^H$ as either $[S_{\tau-1}^H]^2 V_{\tau-1,T}^H - [S_{\tau-1}^L]^2 V_{\tau-1,T}^L \rightarrow 0$ or $p_{\tau-1}^H \rightarrow 1$.
2. If $\tilde{V} > 0$, then $p_{\tau-1}^H \geq \tilde{p}$ if and only if $[S_{\tau-1}^L]^2 V_{\tau-1,T}^L \leq [S_{\tau-1}^H]^2 V_{\tau-1,T}^H$.
3. If $\tilde{V} < 0$, then \tilde{p} is an uninformative bound on $p_{\tau-1}^H$.

Proof. See appendix. □

The proposition defines an estimator \tilde{p} of $p_{\tau-1}^H$. The first result establishes that \tilde{p} becomes an arbitrarily good approximation to $p_{\tau-1}^H$ as the difference in the (scaled) conditional simple variance swap rates goes to zero ($[S_{\tau-1}^H]^2 V_{\tau-1,T}^H - [S_{\tau-1}^L]^2 V_{\tau-1,T}^L \rightarrow$

²²The pricing of variance swaps dates back to the early 1990s, but most literature assumes that the underlying stock price cannot jump. See Carr and Lee (2009) for a review. However, we are here interested precisely in the possibility of jumps. Jiang and Tian (2005) and Carr and Wu (2009) describe the approximation error induced by jumps. We follow the approach of Martin (2017), who redefines the variance to be exchanged so that very small stock prices do not cause the payoff to go to infinity. Martin (2017) assumes European options, yet we observe American options in the empirical application. To minimize the importance of this distinction, we will drop firms with high dividend yields, will use options with short maturities, and will clean options of early exercise premia.

0) or as $p_{\tau-1}^H \rightarrow 1$. Even if the difference in the conditional simple variance swap rates is not small, we again have a good approximation in the case of large $p_{\tau-1}^H$, which, from equation (2), is when the full event effect is most sensitive to $p_{\tau-1}^H$. The second result establishes that \tilde{p} is a lower (upper) bound on $p_{\tau-1}^H$ when the scaled conditional simple variance swap rate is smaller (larger) under the counterfactual event outcome L than under the realized event outcome H . The intuition for the result tracks that already given for equation (8), with $[S_{\tau-1}^H]^2 V_{\tau-1,T}^H - [S_{\tau-1}^L]^2 V_{\tau-1,T}^L \rightarrow 0$ serving as the analogue of the difference in the variance resulting from ω_τ under the different potential event outcomes.²³ As $p_{\tau-1}^H \rightarrow 1$, the possibility that $[S_{\tau-1}^L]^2 V_{\tau-1,T}^L$ differs from $[S_{\tau-1}^H]^2 V_{\tau-1,T}^H$ becomes irrelevant.

The sign of \tilde{V} plays a critical role in Proposition 1, where \tilde{V} is a metric readily constructed from observed option prices and from event study estimates. If $\tilde{V} > 0$, then uncertainty (as measured by the scaled conditional variance swap rates) is reduced by learning the event outcome. This is the standard case, which we implicitly assumed in discussing equation (8). In contrast, if $\tilde{V} < 0$, then the variance conditional on event outcome H is greater than the variance of the full compound lottery. From equation (7), this can happen only if expected variance is sufficiently greater under the realized event outcome than under the counterfactual event outcome. In this case, expected variance is far from independent of the event outcome and \tilde{p} provides an uninformative bound on $p_{\tau-1}^H$, being either less than zero or greater than 1.

In sum, the critical condition for \tilde{p} to recover $p_{\tau-1}^H$ is that the expected variance between day τ and the option expiration date T be independent of the event outcome. Previous approaches to recovering event probabilities with parametric option pricing models implicitly imposed an analogous restriction in addition to other distributional assumptions (e.g., Gemmill, 1992; Carvalho and Guimaraes, 2018). This new approach is therefore a direct generalization of prior literature.

3.3 Comparing the Two Estimators

We have developed two new estimators of the risk-neutral probability of an event. Both estimators are model-free, in contrast to the prior literature that recovers event probabilities from option prices by assuming that stock prices evolve according to specific parametric processes. The first estimator (\bar{p}) requires that some options that are valuable when the realized event happens would have been worth very little if other events had happened. The second estimator (\tilde{p}) requires that the expected

²³If there are multiple possible event outcomes, then $[S_{\tau-1}^L]^2 V_{\tau-1,T}^L$ becomes the expected post-event variance conditional on event H not occurring.

variance of the stock price process over a post-event window not be sensitive to the realization of the event.

The strengths of the estimator \bar{p} are that it is straightforward to compute, that it does not require market agents to anticipate that the event was going to occur on a particular date, and that we know which types of options should yield the tightest bound. In contrast, the estimator \tilde{p} requires approximating an integral over option prices, requires market agents to know the event's date at least one day ahead of time, and imposes an identifying assumption that is difficult to test (but may not be especially restrictive when the time to expiration is short). In particular, the integral approximation becomes poorer when liquidly traded options' strike prices become less dense and/or cover a narrower interval. In this case, we may obtain only a noisy estimate of \tilde{p} .

However, the estimator \tilde{p} can perform well in contexts in which the estimator \bar{p} may yield only a loose bound. First, some events will not affect firms with many deep out-of-the-money options that are liquidly traded. Second, some realized events will not be extreme. Appendix C.2 shows that the bound obtained from \bar{p} cannot become arbitrarily tight for such "middle" events. In contrast, \tilde{p} does not depend on the realized event being extreme. Third, some events will only reduce firms' value. In these cases, researchers must use put options to estimate \bar{p} , but, as discussed above, we might expect greater bias when estimating \bar{p} from put options.

In sum, the two estimators are complementary. Since they rely on distinct assumptions, obtaining similar results from both approaches should increase researchers' confidence in having properly recovered an event's probability.

4 Empirical Approach

Both of our approaches to recovering the priced-in probability of an event require estimating what the price of an option would have been if the event's realization had been known a bit earlier. This is the standard event study identification challenge.

We obtain stock prices, quarterly dividends, and earnings dates from Compustat. We obtain equity options data from OptionMetrics, using all firms available in IvyDB US. We calculate an option's price as the average of its closing bid and its closing ask. In our OPEC application, we study options on crude oil futures. S_t then represents the price of the underlying future. We obtain oil futures and options data from Intercontinental Exchange.

For both estimators, we limit the sample to the nearest major expiration date so that the distinction between European-style and American-style options is less

important (see Appendix C.1).²⁴ For a similar reason, we drop firms with a quarterly dividend yield greater than 2% over the estimation window (e.g., Dubinsky et al., 2019). Previous work has shown that options prices respond to earnings announcements (e.g., Patell and Wolfson, 1979, 1981; Dubinsky et al., 2019), so we limit the sample to firms that do not have an earnings announcement in a 3-day window around the event. Finally, we drop firms whose stock price falls below \$5 at any point in either the estimation or event windows (e.g., Dubinsky et al., 2019).

Using these data, we describe our approaches to estimating \bar{p} and \tilde{p} in turn, including additional, theoretically motivated restrictions designed to recover tight bounds on the priced-in event probability.

4.1 Estimating \bar{p} from Out-of-the-Money Options

We first describe how we estimate \bar{p} in the equity option application, where we observe options on many firms. We then describe how we adjust our empirical approach in the oil market application, where we observe a dense set of strikes on a single commodity.

Consider the case with options on many firms exposed to the same event, with firms indexed by i . We estimate \bar{p} from the following regression:²⁵

$$\ln(C_{iK(t-1)}/C_{iKt}) = \alpha_{iK} + \beta Event_t + \theta_{iK} X_{it} + \varepsilon_{iKt}, \quad (11)$$

where we simplify the call option price notation to C_{iKt} . $Event_t$ is a indicator variable for the day the event information is realized. X_{it} is a vector of controls, which includes time to expiration and its square, indicator variables for the days before and after the event, and indicator variables for a three-day window around earnings announcements.²⁶

The prices of observations with greater liquidity and lower transaction costs provide more precise measures of contemporary information. We therefore weight equation (11) by the inverse of the relative bid-ask spread (Madhavan, 2000; Vayanos and

²⁴Equity options in our data overwhelmingly expire on the third Friday of the month. There are some options that expire on other dates within the month, but we focus our analysis on the major expiration dates because the other expiration dates are less liquid. We use the first major expiration date that is at least a week past the end of our estimation window (see Beber and Brandt, 2006; Kelly et al., 2016).

²⁵We use the log of the option price ratio because we find that log-changes in option prices are approximately normally distributed.

²⁶We do not control for the market index. Because we analyze substantial events that may have affected that index, controlling for it could absorb the desired event effect. Controlling for the market index may be appropriate in other applications.

Wang, 2013) averaged over days t and $t - 1$. We assign a weight of zero if either day has a bid of zero. We run the regression on trading days that are within 100 days before the event and all trading days after the event up to 7 days before the option's expiration date (see Beber and Brandt, 2006). Standard errors are clustered by firm and by date. We run an analogous regression when we examine puts.

We estimate \bar{p} by predicting $C_{iK(\tau-1)}^L/C_{iK(\tau-1)}^H$ from $\hat{\beta}$, comparing the option price on the day before the event to what the option price would have been if the event outcome had been known but nothing else had changed. This gives us $\bar{p} = \exp(\hat{\beta})$. We do not let β vary across firms because the event probability should here not vary across firms, although this constraint could be relaxed in other applications.²⁷

The regression in equation (11) will recover an upper bound on the event probability. The bound is tight if $C_{\tau-1,T}^L(S_{\tau-1}^L, K)$ is small, but $C_{\tau-1,T}^L(S_{\tau-1}^L, K)$ is unobservable. We therefore make a few additional restrictions on the set of options we analyze in order to recover a tight bound on the event probability without knowledge of $C_{\tau-1,T}^L(S_{\tau-1}^L, K)$. These restrictions all leverage theoretically motivated insights about how the bias in equation (5) varies with observables.

First, we saw in Section 3.1 that deeper out-of-the-money options will generate tighter bounds than closer-to-the-money options, assuming all are liquid.²⁸ Our preferred specifications therefore limit the sample to the deepest out-of-the-money liquid option for each firm.²⁹

Second, following Section 3.1, we should use put options for firms whose stock price declines with the event, but estimates based on put options will be especially biased if those options account for a disaster risk that is independent of the event. We assess firms' response to the event via firm-by-firm traditional event studies:

$$\ln(S_{it}/S_{i(t-1)}) = \gamma_{i1} + \gamma_{i2}Event_t + \gamma_{i3}X_{it}^S + \varepsilon_{it}, \quad (12)$$

where S_{it} is the closing stock price for firm i on trading date t , X_{it}^S includes the earning announcement and event controls from before, and standard errors are clustered by

²⁷By absence of arbitrage, all firms exposed to the same event must have the same risk-neutral probability of the event. Heterogeneous beliefs about the probability of an event could affect that market probability but should not lead it to differ across firms. Probabilities of composite events (such as complex court rulings) can vary by firm if exposure to various components of the event differs by firm.

²⁸As noted earlier, this effect is especially strong when a realized event increases stock prices because the empirical researcher then analyzes call options, whose value does not include a hedge against disasters.

²⁹We define the set of sufficiently liquid strikes according to the method used in constructing the VIX: we use all strikes with nonzero bids between the forward price and the point at which two strikes in a row have bids of zero. See Cboe (2019).

firm and by date.³⁰ We use our estimate of γ_{i2} to restrict our analysis to firms whose stock prices were positively affected by the event ($\hat{\gamma}_{i2} > 0$) and thus to call options.

Finally, we restrict attention to firms that are strongly affected by the event. A firm that is unaffected by an event does not provide information about $p_{\tau-1}^H$. Moreover, a firm will, all else equal, generate a tighter bound on the event probability if its stock price is especially sensitive to the event: for a given probability of the event, a larger jump in the stock price at the event ($S_{\tau-1}^H - S_{\tau-1}$) implies that the counterfactual value of the stock ($S_{\tau-1}^L$) is low and thus that the counterfactual value of the option ($C_{\tau-1}^L$) is likely to be small.³¹ We use the t-statistic on $\hat{\gamma}_{i2}$ in equation (12) to limit the sample to firms that were especially affected by the event. This limitation should reduce bias but also increases the variance of our estimate by limiting us to fewer option prices. We provide results with increasingly strict t-statistic cutoffs to show how \bar{p} and its standard error vary with the severity of this restriction.

We adapt our approach slightly for the oil market application, where we have only a single underlying asset but substantially greater density of strike prices. We seek to use information at more strikes while still focusing on strikes that are both farther from the money and relatively liquid. We therefore estimate the event's effect on option prices by fitting a spline across strikes in the event-day effect:

$$\ln(C_{K(t-1)}/C_{Kt}) = \beta_0 Event_t + \sum_{j=1}^J \beta_j \min(K - \mu_{j-1}, \mu_j - \mu_{j-1}) Event_t + \varepsilon_{Kt}, \quad (13)$$

where the μ_j are knots that evenly divide the oil option strikes into J evenly sized groups (with $\mu_0 = 0$). We fit splines with $J = 10$ knot points.³² The variables in the vector X_t control for other OPEC meetings that occur in a given meeting's estimation window, for splines in the day before and day after the meeting effect, and for time to expiration and its square. Because we do not observe bid and ask prices, we weight observations by average open interest across days t and $t - 1$. We take the minimum of the estimated spline and its standard error as our estimate of \bar{p} .

4.2 Estimating \tilde{p} from Synthesized Variance Swaps

Now consider estimating \tilde{p} . Begin by considering the probability \tilde{p}_i implied by options on firm i . Our estimate of \tilde{p} (from equation (10)) requires the variance swap

³⁰We use a 200-day estimation window (180 days before and 20 days after the event).

³¹This corresponds to the case in which $f_{\tau-1}(S_T|S_{\tau-1}^L, L)$ in Figure 1 has little mass above K .

³²Appendix B shows that our estimates are not sensitive to the number of knot points.

rate on the day before the event ($V_{i(\tau-1,T)}$) and the counterfactual variance swap rate if the event's outcome were already known ($V_{i(\tau-1,T)}^H$). Equation (9), from Martin (2017), shows that calculating $V_{i(\tau-1,T)}$ requires integrating call and put option prices over the domain of strikes. However, we observe option prices only at discrete strikes. We discretize the integral and calculate the forward price by adapting the methodology used to construct the familiar VIX index (see Cboe, 2019).³³ We use $R_{\tau-1,T}$ from a daily version of the 3-month LIBOR rate. To calculate $V_{i(\tau-1,T)}^H$, we use counterfactual option prices predicted from a regression like (11), modified to allow the event effect (β_{iK}) to vary by firm and strike.³⁴ Similarly, we recover $S_{i(\tau-1)}^H$ from a regression that adds heterogeneity in the event effect by firm and strike to equation (12). The observed forward price and $S_{i(\tau-1)}$ together imply $\tilde{R}_{i(\tau-1,\tau)}$ (see footnote 21).

Following Martin (2017), our approach to discretizing the integral avoids using pricing models at any step. The variance swap pricing literature often more finely discretizes the integral by converting option prices to implied volatilities, interpolating implied volatility on a fine grid between observed strikes, and then converting back to option prices (e.g., Carr and Wu, 2009). This alternate approach has the advantages of stripping early exercise premia from option prices and more finely approximating the integral, but it has the disadvantage of using pricing models to convert to and from implied volatility. Appendix A.2 describes this alternate approach and shows that it yields results similar to our baseline ones.

We estimate \tilde{p} as the average of the \tilde{p}_i . Following Proposition 1, we drop firms for which the central estimate of \tilde{V}_i is negative. We calculate the standard error of \tilde{p} from the generalized method of moments.³⁵ Estimated covariance matrices are

³³For equities, we adapt the VIX methodology (Cboe, 2019) to drop options with a bid of zero on either day $\tau - 1$ or day τ as well as some others likely to be illiquid. We drop any firms that do not have at least three strikes surviving this restriction. Appendix A.2 shows that results are not sensitive to tightening this requirement. For commodities, we drop all options with open interest below 50 and all strikes that are at least as deep out-of-the-money as the first strike at which the price is constant for two strikes in a row.

³⁴As in Section 4.1, we weight by the inverse of the relative bid-ask spread averaged over days $t - 1$ and t , with a weight of zero assigned if either day has a bid of zero. When calculating counterfactual stock prices and option implied volatilities, we use a second-order Taylor expansion that adjusts for the standard error of the estimated event effect.

³⁵We stack the moment conditions underlying all regressions of the forms (11) and (12) into a single system and add the moment conditions $\tilde{p}_i - \tilde{p} = 0$. To apply those last moment conditions, we express each counterfactual option price in a version of equation (9) in terms of the estimated coefficients from the regressions (11) and (12). The system is over-identified when there is more than one firm. In future work, we will estimate the regressions and \tilde{p} simultaneously, including with a two-step estimator, and we will explore running event studies on V directly. When there is

clustered by date.

5 Applications

We apply our new methods to two high-stakes settings. The first setting, the 2016 U.S. election, is one in which we have a rough idea of the probability from prediction markets and polling data. It serves to validate our approach. The second setting is the regular meetings of the Organization of the Petroleum Exporting Countries. We use our methods to understand OPEC's ability to influence world oil prices through the rise of U.S. oil production and show how this understanding changes after controlling for market anticipation of OPEC meeting outcomes.

5.1 The 2016 U.S. Election

On Tuesday November 8, 2016, Donald Trump was elected President of the United States and the Republican party captured both chambers of Congress.³⁶ This extreme outcome was widely surprising. On the morning of November 8, the prediction market PredictIt gave Trump a 22% chance of winning, gave Republicans a 41% chance of controlling the Senate, and gave Republicans a 16% chance of controlling the presidency and both chambers of Congress. The polling-driven New York Times' Upshot forecast gave Trump a 15% chance of winning, whereas the polling-driven FiveThirtyEight forecasts gave him a 28–29% chance of winning.³⁷ Both forecasts gave Republicans around a 50% chance of controlling the Senate. PredictWise, which uses both prediction market data and polling data, gave Trump a 12% chance of winning, gave Republicans a 33% chance of controlling the Senate, and gave Republicans around a 94% chance of controlling the House of Representatives. Using conditional probabilities implied by the PredictIt contracts, the chance of a Republican sweep was 13% according to the Upshot, 26% according to FiveThirtyEight, and 7% ac-

only one firm (or commodity), we stack the regression moment conditions to estimate a covariance matrix for the regression coefficients and apply the delta method to \tilde{p}_i .

³⁶The country learned of this outcome in between markets' close on November 8 and their opening on November 9. This election is particularly attractive for validating our methods because it was an extreme outcome that was largely known by the time markets closed on the day after the election. Other elections can be more ambiguous. For instance, the 2008 election of Barack Obama occurred in tandem with critical Senate elections that were not resolved for some time afterward.

³⁷Whereas prediction markets plausibly give a risk-neutral probability, polling-driven estimates target an objective probability. It is unclear whether the risk-neutral probability of the realized election outcome should be greater or less than the objective probability.

according to PredictWise.³⁸ Finally, the prediction market Betfair gave Trump a 20% chance of winning and the bookmakers Paddy Power and Ladbroke's odds implied that Trump had a 22% and 24% chance of winning, respectively. Using the range of Senate probabilities above, these three sources imply that the chance of a Republican sweep was between 12% and 21%. In sum, the various markets and models imply that the chance of a Republican sweep was somewhere between 7% and 26%.

Table 1 reports our estimates of the probability of the realized 2016 election outcome with our different approaches.³⁹ The top panel reports the \bar{p} obtained from out-of-the-money options, as described in Section 4.1. The first column does not impose any of the theoretically motivated restrictions. As expected, this estimate is subject to severe upward bias: we estimate an unreasonable probability above 1. The second column restricts attention to the deepest out-of-the-money options that are sufficiently liquid. Again, this estimate is biased upwards with an estimated probability above 1. The third column aims to further reduce bias by restricting attention to specifications with call options: it drops firms for which the election lowered stock prices ($\hat{\gamma}_{i2} \leq 0$). The estimated probability falls only slightly, to 0.60.

The remaining columns limit the sample to firms with sufficiently large t-statistics on $\hat{\gamma}_{i2}$. The fourth column drops those firms which were either nearly unaffected by the election (and thus reveal nothing about p_{T-1}^H) or not strongly affected by the election (and thus biased upwards). The estimated probability falls substantially, to 0.39. The remaining columns tighten this restriction further, with the estimated probability appearing to converge around 0.12.⁴⁰ This estimated probability is well within the range of pre-election probabilities surveyed above and is quite close to the PredictIt contract for full Republican control.

One might be concerned that, by attending only to firms with large t-statistics, we select for firms that had *something* unexpected happen to them, whether the

³⁸PredictIt implies that the probability of Trump winning conditional on Republicans winning the Senate was 0.39. Calculations for other sources assume that the ratio of that conditional probability to the unconditional probability of Trump winning was constant across models and markets. We assume that Republicans would not win the Senate without also winning the House of Representatives, which was a common belief at the time.

³⁹Our methods require that some number of liquidly traded firms be sensitive to the election outcome. The 2016 election clearly passes this test. Appendix A shows that most firms' stocks jumped in extremely unusual fashion: 71% (10%) of firms had t-statistics of magnitude greater than 10 (50), which should very rarely happen by chance. The large number of firms affected by the same event is a good setting in which to test our methods. The last row of the top panel of Table 1 reports the combined volume on the election day and the day after the election for the options used in each estimation.

⁴⁰If we allowed firms with $\hat{\gamma}_{i2} < 0$ and imposed a cutoff in the absolute value of the t-statistic, then the estimated probability would be around 2 percentage points larger in each of these columns.

Table 1: Estimated Probability for 2016 Election Outcome

	Standard Restrictions	Extreme Strikes	Positive Events	t-statistic > 25	t-statistic > 50	t-statistic > 75	t-statistic > 100
Out-of-the-Money Options Approach:							
Probability	1.032	1.062	0.602	0.387	0.221	0.132	0.119
Standard Error	(0.003)	(0.005)	(0.007)	(0.006)	(0.007)	(0.006)	(0.024)
# Firms	2,629	2,595	1,835	764	188	43	10
# Options	79,885	5,122	1,835	764	188	43	10
# Option-Days	2,979,681	214,576	77,569	32,057	7,624	1,897	406
Adjusted R^2	0.043	0.050	0.136	0.277	0.340	0.406	0.498
Volume	7,755,853	392,848	5,761,200	1,591,146	475,303	118,293	6,073
Variance Swap Approach:							
Probability	0.615			0.284	0.143	0.118	0.021
Standard Error	(0.026)			(0.022)	(0.011)	(0.017)	(0.035)
# Firms	372			105	37	14	3
# Options	3,973			1,024	356	118	24
# Firm-Days	74,028			20,895	7,363	2,786	597
# Option-Days	152,095			40,365	11,903	3,873	856

The case with standard restrictions weights by the average of the inverse relative bid-ask spread, removes firms with high dividends or low stock prices during the event window, and removes options with bids equal to zero. Extreme strikes only uses one option per firm, either the highest call option or the lowest put option with a positive bid on both election day and the day after that does not have two options with zero bids closer to the money. Positive events restricts attention to only those firms with positive stock price movement after the event. For the Out-of-the-Money Approach, each column shows the estimated probability, \bar{p} , if the stock return event study t-statistic is greater than the cutoff. Standard errors are clustered at both the firm and date level. The Variance Swap Approach is similar, except using \tilde{p} , using the absolute value of the stock return t-statistic, and clustering by date.

event of interest or some other news.⁴¹ Imposing the t-statistic cutoff would then bias our estimated probability down by leading us to estimate the probability of some unusual event that was not our target. This concern is in general legitimate, but its effect is likely minimal in this particular application. First, we examine t-statistics well beyond those that should arise with any frequency in the course of the distribution of daily news that firms experience. The 2016 U.S. election was

⁴¹Appendix A shows that the 2016 election day is a major outlier for these firms' option prices: on typical days, their option prices do not increase in a way that would indicate a probability near 0.12. Instead, as one should expect, option prices appear to be martingales in that the ratio of consecutive days' option prices is typically evenly distributed around 1. On the overwhelming majority of days, this ratio is between 0.75 and 1.25 for nearly all firms.

such an enormous event for such a large number of firms that we can even examine firms with t-statistics greater than 75. Daily news should generate such t-statistics in fewer than 1 in 100,000 firms, whereas Table 1 reports 43 such firms from the 1,835 with positive events. Second, in Appendix A we perform a robustness check in which we select for those firms whose stock price movements were strongly correlated with the PredictIt prediction market contract for a Trump victory in the run-up to the election. We show that restricting attention to those firms does not change our results.

The lower panel of Table 1 reports the \tilde{p} obtained from synthesized variance swaps, as described in Section 4.2. The first column directly imposes all restrictions except for the t-statistic restrictions. We recover a probability of 0.62, consistent with the analogous columns for \bar{p} . The next two columns are empty because they are not relevant to \tilde{p} . The remaining columns of Table 1 demonstrate how tightening the t-statistic cutoff affects the estimated \tilde{p} , where here the cutoff is on the absolute value of the t-statistic rather than its level. As before, imposing some restriction is critical because it eliminates firms that are not affected by the event and thus do not provide an estimate of $p_{\tau-1}^H$: the estimated \tilde{p} falls to 0.28. Tightening the t-statistic cutoff to only 50 recovers a probability consistent with the converged \bar{p} .⁴² Tightening the t-statistic cutoff all the way to 100 does yield a much smaller probability, but at this point there are only three firms left and the standard error grows larger.

The identifying assumption underlying the variance swap approach is that the election outcome did not affect the expected variance over the remaining life of the option contracts. The election of Donald Trump would almost surely violate this assumption if we used long-dated options, but we analyze options that expire only ten days after the election. Trump would not take office for more than two further months, and he would announce his first Cabinet official only on the day these options expired. Table A-3 of Appendix A.2 provides evidence in support of the identifying assumption. It shows how our estimates of \tilde{p} change when we limit our sample to only those firms with realized volatility over the 9 days after the election that is close to the realized volatility over the 9 days before the election.⁴³ We find that results

⁴²The number of firms is generally lower when estimating \tilde{p} because we lose firms with either $\tilde{V} < 0$ or fewer than three usable strikes, but the number of options is nonetheless generally greater when estimating \tilde{p} because this method uses all liquid strikes rather than only a firm's most extreme liquid strike.

⁴³The realized volatility over the 9 days after the election approximates the expected post-election variance conditional on Trump winning while the realized volatility over the 9 days leading up to the election approximates the unconditional expected post-election variance. These two variances are similar if the post-election variance is not sensitive to whether Trump wins or not. Approximating expectations of later volatility by recent realized volatility follows Martin (2017) and Zhou (2018),

are not sensitive to using only those firms with the smallest differences between these approximations of conditional and unconditional expected volatility.

Figure 2 shows how the estimated \bar{p} (solid line) and \tilde{p} (dashed line) converge as the minimum required event study t-statistic increases. The estimated \bar{p} converges for cutoffs above 70, and the estimated \tilde{p} converges for cutoffs above 30. We should not be surprised that \tilde{p} converges faster: the cutoff there serves merely to identify firms exposed to the event, whereas for \bar{p} the cutoff also identifies firms for which observed out-of-the-money call options would have been nearly worthless under alternate election outcomes. The 95% confidence intervals grow in the cutoff. This increasing imprecision reflects the bias-variance tradeoff described above, as higher t-statistic cutoffs simultaneously reduce both the upward bias and the number of observed firms. Most importantly, at a t-statistic cutoff around 80, both estimators have converged to a probability around 0.12 with fairly tight standard errors—despite their different identifying assumptions and despite their reliance on different sets of options per firm. The coherence across the estimates increases our confidence in each of the individual estimates.

In sum, our new, model-free methods of estimating the risk-neutral probability of an event generate estimates that are compatible with each other and that move in expected ways as we apply restrictions meant to reduce their bias. Further, these estimates are broadly consistent with the range of estimates available from prediction markets, bookmakers, and polling-driven models. A number of recent papers have relied on event studies of this election (e.g., Mukanjari and Sterner, 2018; Wagner et al., 2018a,b; Ramelli et al., 2019). While this election was indeed more surprising than most, our results nonetheless suggest that the event study estimates should be inflated by 14% (multiplied by $1/(1 - 0.12)$) to recover the full effect of the election.

5.2 OPEC Meetings

Measuring the extent to which the market anticipated an event is critical for being able to compare the relative magnitude of a series of events. We further demonstrate the power of our new approaches by analyzing a series of events involving Organization of the Petroleum Exporting Countries (OPEC) decisions on oil production quotas.

OPEC is a cartel that aims to achieve higher prices by restricting oil supply.

among others. The calculated pre-event volatility uses a regression of stock prices on the price of the PredictIt contract for Trump’s victory to remove volatility associated with changes in beliefs about the event’s outcome. Finally, note that these approximations are to expectations under the physical measure whereas the assumption is for expectations under the risk-neutral measure.

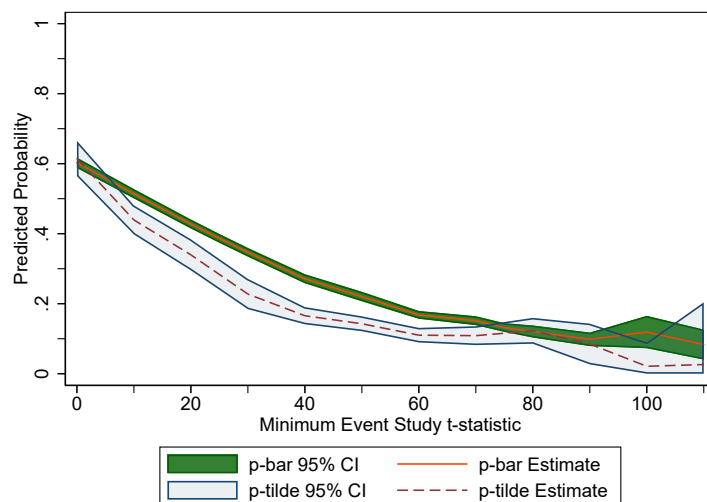


Figure 2: Estimated Probability of the 2016 Election Outcome, By Event Study t-statistic Cutoff

OPEC meets at least twice a year to assess its production quotas, with special meetings as warranted. Its meetings are typically of high interest, as the price of oil has implications throughout the global economy. There have been many event studies of OPEC meetings, but anticipation of meetings has been recognized as a problem for understanding OPEC's effect on markets dating back to the earliest work (Draper, 1984; Deaves and Krinsky, 1992).⁴⁴ Differential anticipation of announcements has been proposed as an explanation for why oil markets appear to react to announcements of higher or unchanged quotas but not to smaller quotas (Hyndman, 2008; Demirer and Kutan, 2010; Loutia et al., 2016). Anticipation could also explain why oil markets' reactions appear smaller than seems reasonable (Wirl and Kujundzic, 2004) and why effects are larger when OPEC maintains the current quota because its members cannot agree as opposed to when it maintains the quota because its members agree to do so (Spencer and Bredin, 2019).

We study the announcements from 13 OPEC meetings in 2011–2016.⁴⁵ Most

⁴⁴Researchers have addressed anticipation by including earlier days in the event window (e.g., Draper, 1984; Loutia et al., 2016; Lauenstein and Simic, 2017) and by using realized movements to classify events as good or bad news relative to expectations (e.g., Deaves and Krinsky, 1992).

⁴⁵Diego Kaenzig kindly provided the times OPEC decisions were formally announced and the times they appeared on Bloomberg news. Horan et al. (2004) show that options' implied volatil-

meetings result in no change in production quotas, and press reports suggest that most of these outcomes are largely expected.⁴⁶ OPEC raised its production target in November 2011 amid global economic troubles. This agreement was seen as largely symbolic, since producers had already been ignoring the old quota. However, some meetings do generate news about quotas. Most observers expected OPEC to cut production in November 2014 (see Plante, 2019) in response to sliding oil prices. Some members indeed proposed such cuts, but, at the behest of Saudi Arabia, OPEC agreed to maintain the quota. By December 2015 and June 2016, disagreement was severe enough to keep OPEC from agreeing on quotas. In September 2016 the members managed to reach a preliminary agreement to cut oil production, which was eventually finalized in a November meeting. Contemporary media reports suggest that these latter two meetings were far from foregone conclusions.

In contrast to our analysis of the 2016 election, we do not attempt to identify which firms are affected by the events. Instead, we use options on the commodity (oil) that is the subject of the meetings. In particular, we study options on front-month Brent oil futures. These futures are often used as the spot price of internationally traded oil.⁴⁷

The left panel of Figure 3 displays the time-series of Brent front-month oil future prices with OPEC meeting dates. The right panel displays the event-study estimates of the abnormal return on oil futures due to each OPEC meeting, with 95% confidence intervals.⁴⁸ These traditional event study estimates suggest that the OPEC meetings

ity anticipated the release of valuable information at the meetings of the Ministerial Monitoring Committee but not at the biannual conference. However, that committee no longer had power over production quotas by the time of our sample (Lin and Tamvakis, 2010), so we follow others in focusing on the biannual conference.

⁴⁶OPEC meetings generate news about continuous measures of quotas and news about non-quota factors. Even though quotas are in principle continuous, they in practice seem to have only a few focal points, often round numbers. Our framework of discrete event outcomes may therefore be an adequate approximation. Regarding non-quota news, our probabilities can be viewed as recovering the joint probability of all news released, much as we recover the joint probability of Congressional and Presidential election outcomes in Section 5.1. Känzig (2019) argues that non-quota news is not very important, and Brunetti et al. (2013) provide evidence in support of this view. See Gürkaynak et al. (2020) for more on identifying non-headline news.

⁴⁷There are two primary oil future prices: Brent and West Texas Intermediate (WTI). Brent futures measure the price of oil to be delivered to the North Sea and are now the primary international benchmark. WTI futures measure the price of a barrel of oil delivered to Cushing, Oklahoma and are the U.S. benchmark. These prices typically track each other but did diverge in the period of interest due to a surplus of non-tradable U.S. oil production. We focus on Brent prices because they are more directly linked to OPEC decisions. We report results with WTI prices in Appendix B.

⁴⁸Event study estimates use a 120-day estimation window, including 90 days before the event and 29 days after. Controls include the future's days to expiration, its square, and an indicator for

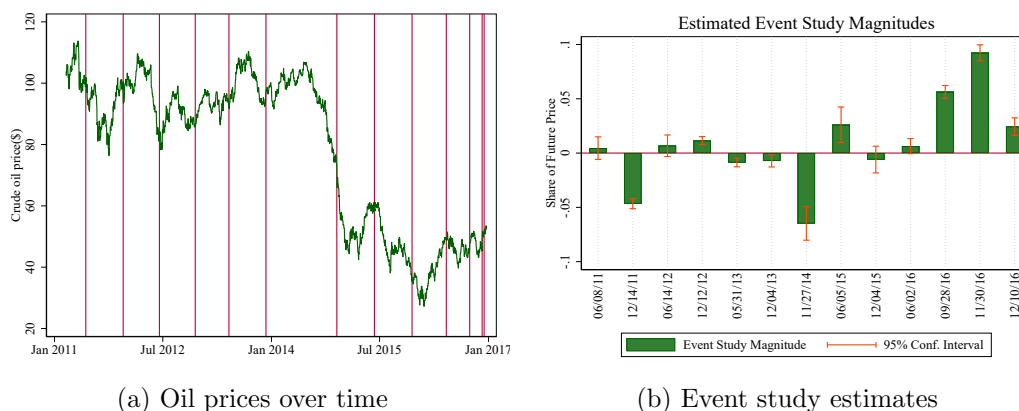


Figure 3: Left: Time series of Brent front-month oil future prices, with vertical lines marking OPEC meetings. Right: Traditionally-estimated event-study change in Brent front-month oil future prices (plus 95% confidence intervals) at the date the outcome of the OPEC meeting is known.

in December 2011, November 2014, and September and November 2016 moved oil future prices substantially. Most other meetings did not have a significant effect on oil prices. Notably, the meetings in December 2015 and June 2016—at which OPEC failed to agree on production targets—generated only minimal changes in oil prices. These small responses could suggest that the recent dramatic surge in U.S. shale oil production had limited OPEC’s influence on oil markets. However, we should hesitate to draw such conclusions unless we also know the extent to which the market anticipated these meeting outcomes.

Figure 4 demonstrates the empirical approach, using the OPEC meeting on September 28, 2016 as an example.⁴⁹ This meeting increased the price of oil, so we study call options. The circles depict the \bar{p} implied by strike-by-strike regressions analogous to equation (13), with sizes proportional to standard errors. The line and shading indicate the estimated spline and its 95% confidence interval. The vertical line indicates the price of the oil future on the day of the event, so out-of-the-money options have strikes to the right of the line. As expected, call options

any other OPEC meeting that falls in the estimation window. Observations are weighted by the future’s average open interest over day t and day $t - 1$.

⁴⁹We show the spline plot for the September 28, 2016 meeting because it has the lowest estimated probability of any of the OPEC meeting in our sample, which makes the variation in the estimated probability across strikes easiest to see. Appendix B contains analogous plots for every OPEC meeting in our sample.




Figure 4: Estimates of \bar{p} using Brent front-month call options for the September 28, 2016 OPEC meeting. Circles represent estimates using only that strike and weighting by the average inverse of the bid-ask ratio. Circle sizes are proportional to the standard error of the estimate. The green line is the spline with 10 knots, and orange shading is the 95% confidence interval. The minimum of the spline is our preferred estimate of the meeting outcome’s probability.

at lower strikes do appear to produce upward-biased estimates of the event probability. On the other hand, deeper out-of-the-money (higher-strike) estimates of the event’s probability exhibit both larger standard errors and more noise across strikes, reflecting a bias-variance trade-off similar to that seen across firms in Figure 2. The spline combines the information across strikes. Its minimum occurs at a probability of 0.29, with a standard error of 0.0129.

Figure B-4 in Appendix B compares the estimated \bar{p} and \tilde{p} . The two approaches yield quite similar estimates despite their distinct identifying assumptions: the unweighted correlation between the two estimates is 0.68, and the correlation is 0.88 when weighting by the inverse of an event’s average standard error across the two approaches. As with the 2016 election, the coherence between the two sets of estimates increases confidence in each. Unfortunately, we recover \tilde{p} only for 6 of the 13 meetings, with the other meetings’ estimated probabilities being uninformative (i.e., \tilde{V} is negative; see Proposition 1). We proceed by using \bar{p} in order to take advantage of the full sample of meetings.

The left panel of Figure 5 depicts the \bar{p} graphically. Green dots indicate how estimated probabilities relate to the effect estimated from a traditional event study (as in Figure 3). Events that we find to have high probability (such as meetings in December 2015 and June 2016) generally have small event study estimates, while

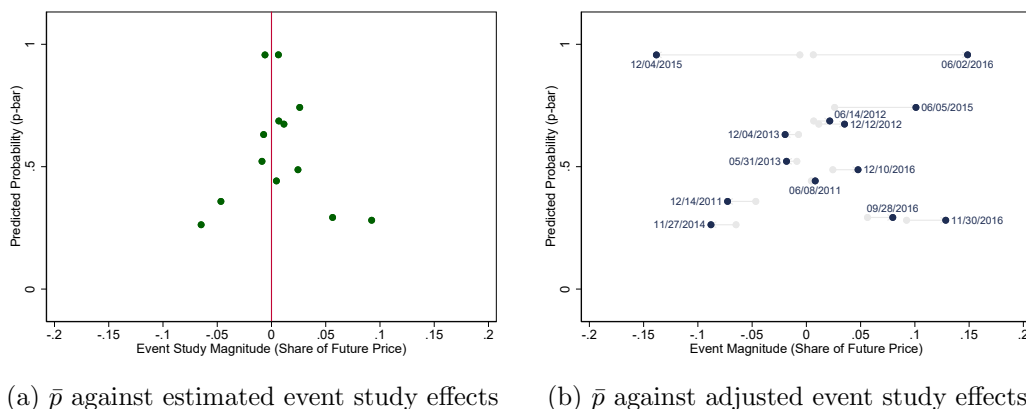


Figure 5: Panel (a): Green circles plot predicted event probabilities \bar{p} against conventional event study estimates. Panel (b): Blue circles plot predicted event probabilities \bar{p} against event study estimates adjusted for the event probabilities (conventional event study estimates are in light gray for comparison).

meetings with large event study estimates (such as November 2014 and November 2016) tend to be more surprising. The blue dots in the right panel of Figure 5 show how adjusting for these estimated probabilities changes our understanding of which meetings were important (conventional event study estimates are in light gray for comparison). After adjusting for the estimated probability of the meeting outcome, there is little correlation between the event-study effect of the meeting on oil future prices and the market's predicted probability of the meeting outcome.

Figure 6 displays the original (orange) and adjusted event study estimates (green), organized according to the timeline of meetings and with 95% confidence intervals. The traditional event study estimates systematically underestimate OPEC's effect on oil markets, sometimes substantially. Meetings in December 2015 and June 2016 may have had the largest impact on oil prices, although these two adjusted estimates are imprecise.⁵⁰ Standard event studies did not detect this impact because the meetings' outcomes were well-anticipated. These meetings are particularly interesting because OPEC failed to agree on new production quotas. That failure was important, even though expected.⁵¹

⁵⁰Confidence intervals become wide when conventionally estimated event effects are small and imprecise because adjusting for the event probability makes larger positive values and larger negative values both more likely.

⁵¹See Spencer and Bredin (2019) for more on the importance of failing to agree on new quotas versus agreeing to maintain quotas.

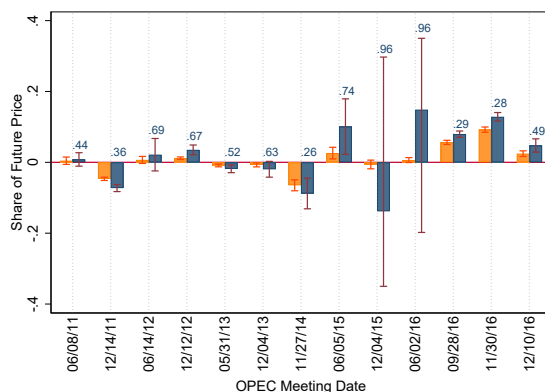


Figure 6: Orange bars plot conventional event study estimates, and blue bars adjust for event probabilities. 95% confidence intervals are truncated at event study magnitudes of -0.35 and 0.35 . \bar{p} is printed above each blue bar.

Overall, these results highlight the importance of estimating event probabilities. We find that OPEC meetings did not strongly affect oil prices in 2012 and 2013 but did substantially impact oil prices from late 2014 through 2016. In standard event study estimates, the market's anticipation of meeting outcomes in 2015 and June 2016 obscures the meetings' true importance for oil prices. As a result, standard event studies detect impacts from OPEC decisions only once OPEC actually begins reaching agreements to cut production. However, our probability-adjusted estimates show that OPEC substantially affected world oil markets even during a period in which it was not changing production and during which U.S. shale oil producers had become a substantial force in world markets.

6 Conclusions

We have demonstrated how to use time series variation in option prices to estimate the priced-in probability of events. In contrast to prior literature, our two new approaches do not impose parametric models on the evolution of stock prices. Both approaches boil down to running event studies in option prices to complement conventional event studies in stock prices. We show that our approaches appear to work in practice. Each approach estimates a probability for President Trump's 2016 election victory that is consistent with the range of probabilities implied by book-makers, prediction markets, and polling-driven models. Importantly, the estimated probabilities are consistent with each other, despite relying on distinct identifying

assumptions. We also estimate probabilities for OPEC meetings that are consistent across the two estimators, that vary in ways consistent with narrative evidence, and that improve our understanding of the importance of OPEC for oil markets.

Our new methods come with two caveats. First, we recover the probability of a realized event, but some event studies seek the probability of a future policy whose odds are merely shifted by the event, as when an election increases the chance of tax reform. Our estimated probabilities are useful in these cases, but they are only part of the adjustment required to recover the full effect of the policy from the event study estimate. Second, we recover the probability of an event as of the day before the event occurred. In some cases, researchers seek a time series of pre-event probabilities in order to analyze the evolution of uncertainty about an event (e.g., Carvalho and Guimaraes, 2018). One could in principle construct such a time series using our methods, but doing so would challenge the identifying assumptions underpinning the event study regressions used to estimate the probabilities.

Future researchers should use our new methods to improve event study estimates for cost-benefit analyses. For instance, many researchers have used event studies to assess the Affordable Care Act and minimum wage laws. Adjusting for estimated event probabilities could substantially revise such assessments. Future researchers should also use our new methods to improve analyses of economic policy uncertainty. Such work often relies on textual evidence and only recently has begun to measure firm-level exposure to uncertainty. Our techniques offer a new revealed preference measure of firm-level uncertainty. This measure could be used to understand and validate existing aggregated measures of policy uncertainty and could provide a cross-sectional dimension when testing for effects of uncertainty on economic activity.

References

- Acharya, Sankarshan (1993) “Value of latent information: Alternative event study methods,” *The Journal of Finance*, Vol. 48, No. 1, pp. 363–385.
- Al-Ississ, Mohamad M and Nolan H Miller (2013) “What does health reform mean for the health care industry? Evidence from the Massachusetts special Senate election,” *American Economic Journal: Economic Policy*, Vol. 5, No. 3, pp. 1–29.
- Auerbach, Alan J. and Kevin A. Hassett (2007) “The 2003 dividend tax cuts and the value of the firm: An event study,” in Alan J. Auerbach, James R. Hines Jr., and Joel Slemrod eds. *Taxing Corporate Income in the 21st Century*, Cambridge: Cambridge University Press, pp. 93–126.

- Baker, Scott R., Nicholas Bloom, Brandice Canes-Wrone, Steven J. Davis, and Jonathan A. Rodden (2014) “Why has U.S. policy uncertainty risen since 1960?” Working Paper 19826, National Bureau of Economic Research.
- Ball, Ray (1972) “Changes in accounting techniques and stock prices,” *Journal of Accounting Research*, Vol. 10, pp. 1–38.
- Barraclough, Kathryn, David T. Robinson, Tom Smith, and Robert E. Whaley (2013) “Using option prices to infer overpayments and synergies in M&A transactions,” *Review of Financial Studies*, Vol. 26, No. 3, pp. 695–722.
- Barro, Robert J. (2006) “Rare disasters and asset markets in the twentieth century,” *The Quarterly Journal of Economics*, Vol. 121, No. 3, pp. 823–866.
- Barro, Robert J. and José F. Ursúa (2012) “Rare macroeconomic disasters,” *Annual Review of Economics*, Vol. 4, No. 1, pp. 83–109.
- Bates, David S. (2000) “Post-’87 crash fears in the S&P 500 futures option market,” *Journal of Econometrics*, Vol. 94, No. 1, pp. 181–238.
- Beber, Alessandro and Michael W. Brandt (2006) “The effect of macroeconomic news on beliefs and preferences: Evidence from the options market,” *Journal of Monetary Economics*, Vol. 53, No. 8, pp. 1997–2039.
- Bell, Brian and Stephen Machin (2017) “Minimum wages and firm value,” *Journal of Labor Economics*, Vol. 36, No. 1, pp. 159–195.
- Bhagat, Sanjai and Roberta Romano (2002) “Event studies and the law: Part II: Empirical studies of corporate law,” *American Law and Economics Review*, Vol. 4, No. 2, pp. 380–423.
- Bianconi, Marcelo, Federico Esposito, and Marco Sammon (2020) “Trade policy uncertainty and stock returns,” working Paper.
- Binder, John J. (1985) “Measuring the effects of regulation with stock price data,” *The RAND Journal of Economics*, Vol. 16, No. 2, pp. 167–183.
- Björk, Tomas (2004) *Arbitrage Theory in Continuous Time*, Oxford; New York: Oxford University Press, 2nd edition.
- Black, Fischer (1975) “Fact and fantasy in the use of options,” *Financial Analysts Journal*, Vol. 31, No. 4, pp. 36–41, Publisher: Routledge eprint: <https://doi.org/10.2469/faj.v31.n4.36>.

- Borochin, Paul and Joseph Golec (2016) “Using options to measure the full value-effect of an event: Application to Obamacare,” *Journal of Financial Economics*, Vol. 120, No. 1, pp. 169–193.
- Bronars, Stephen G. and Donald R. Deere (1990) “Union representation elections and firm profitability,” *Industrial Relations: A Journal of Economy and Society*, Vol. 29, No. 1, pp. 15–37.
- Brown, Stephen J. and Jerold B. Warner (1985) “Using daily stock returns: The case of event studies,” *Journal of Financial Economics*, Vol. 14, No. 1, pp. 3–31.
- Brunetti, Celso, Bahattin Büyükkşahin, Michel A. Robe, and Kirsten R. Sonesson (2013) “OPEC “fair price” pronouncements and the market price of crude oil,” *The Energy Journal*, Vol. 34, No. 4, pp. 79–108.
- Bushnell, James B., Howard Chong, and Erin T. Mansur (2013) “Profiting from regulation: Evidence from the European carbon market,” *American Economic Journal: Economic Policy*, Vol. 5, No. 4, pp. 78–106.
- Campbell, John Y., Andrew W. Lo, and Craig MacKinlay (1997) *The Econometrics of Financial Markets*, Princeton N.J.: Princeton University Press.
- Card, David and Alan B. Krueger (1995) *Myth and Measurement*, Princeton, NJ: Princeton University Press.
- Carr, Peter and Roger Lee (2009) “Volatility derivatives,” *Annual Review of Financial Economics*, Vol. 1, No. 1, pp. 319–339.
- Carr, Peter and Liuren Wu (2009) “Variance risk premiums,” *Review of Financial Studies*, Vol. 22, No. 3, pp. 1311–1341.
- Carvalho, Augusto and Bernardo Guimaraes (2018) “State-controlled companies and political risk: Evidence from the 2014 Brazilian election,” *Journal of Public Economics*, Vol. 159, pp. 66–78.
- Caughey, Devin and Jasjeet S. Sekhon (2011) “Elections and the regression discontinuity design: Lessons from close U.S. house races, 1942–2008,” *Political Analysis*, Vol. 19, No. 4, pp. 385–408.
- Cboe (2019) “VIX White Paper,” Technical report, Cboe Exchange, Inc., <http://www.cboe.com/products/vix-index-volatility/vix-options-and-futures/vix-index/the-vix-index-calculation>.

- Chaudhury, Mohammed M. and Jason Wei (1994) “Upper bounds for American futures options: A note,” *Journal of Futures Markets*, Vol. 14, No. 1, pp. 111–116.
- Cochrane, John H. (2005) *Asset Pricing*, Princeton, N.J: Princeton University Press, revised edition.
- Cornell, Bradford and R. Gregory Morgan (1990) “Using finance theory to measure damages in fraud on the market cases,” *UCLA Law Review*, Vol. 37, pp. 883–924.
- Deaves, Richard and Itzhak Krinsky (1992) “The behavior of oil futures returns around OPEC conferences,” *Journal of Futures Markets*, Vol. 12, No. 5, pp. 563–574.
- Demirer, Rıza and Ali M. Kutan (2010) “The behavior of crude oil spot and futures prices around OPEC and SPR announcements: An event study perspective,” *Energy Economics*, Vol. 32, No. 6, pp. 1467–1476.
- Dorsey, Jackson (2019) “Waiting for the courts: Effects of policy uncertainty on pollution and investment,” *Environmental and Resource Economics*, Vol. 74, No. 4, pp. 1453–1496.
- Draper, Dennis W. (1984) “The behavior of event-related returns on oil futures contracts,” *Journal of Futures Markets*, Vol. 4, No. 2, pp. 125–132.
- Dube, Arindrajit, Ethan Kaplan, and Suresh Naidu (2011) “Coups, corporations, and classified information,” *The Quarterly Journal of Economics*, Vol. 126, No. 3, pp. 1375–1409.
- Dubinsky, Andrew, Michael Johannes, Andreas Kaeck, and Norman J. Seeger (2019) “Option pricing of earnings announcement risks,” *The Review of Financial Studies*, Vol. 32, No. 2, pp. 646–687.
- Easley, David, Maureen O’Hara, and P. S. Srinivas (1998) “Option volume and stock prices: Evidence on where informed traders trade,” *The Journal of Finance*, Vol. 53, No. 2, pp. 431–465, Publisher: [American Finance Association, Wiley].
- Farber, Henry S. and Kevin F. Hallock (2009) “The changing relationship between job loss announcements and stock prices: 1970–1999,” *Labour Economics*, Vol. 16, No. 1, pp. 1–11.
- Fisman, Raymond and Eric Zitzewitz (2019) “An event long-short index: Theory and application,” *American Economic Review: Insights*, Vol. 1, No. 3, pp. 357–372.

- Gemmill, Gordon (1992) “Political risk and market efficiency: Tests based in British stock and options markets in the 1987 election,” *Journal of Banking & Finance*, Vol. 16, No. 1, pp. 211–231.
- Grinblatt, Mark and Kam-Ming Wan (2020) “State pricing, effectively complete markets, and corporate finance,” *Journal of Corporate Finance*, Vol. 60, p. 101542.
- Gürkaynak, Refet S., Burçin Kısacıkoglu, and Jonathan H. Wright (2020) “Missing events in event studies: Identifying the effects of partially-measured news surprises,” *American Economic Review*, Vol. 110, No. 12, pp. 3871–3912.
- Hallock, Kevin F. (1998) “Layoffs, top executive pay, and firm performance,” *The American Economic Review*, Vol. 88, No. 4, pp. 711–723.
- Handley, Kyle and J. Frank Li (2020) “Measuring the effects of firm uncertainty on economic activity: New evidence from one million documents,” Working Paper 27896, National Bureau of Economic Research.
- Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun (2019) “Firm-level political risk: Measurement and effects,” *Quarterly Journal of Economics*, Vol. 134, No. 4, pp. 2135–2202.
- Hennessy, Christopher A. and Ilya A. Strebulaev (2020) “Beyond random assignment: Credible inference of causal effects in dynamic economies,” *Journal of Finance*, Vol. 75, No. 2, pp. 825–866.
- Herron, Michael C. (2000) “Estimating the economic impact of political party competition in the 1992 British election,” *American Journal of Political Science*, Vol. 44, No. 2, pp. 326–337.
- Horan, S.M., J.H. Peterson, and J. Mahar (2004) “Implied volatility of oil futures options surrounding OPEC meetings,” *Energy Journal*, Vol. 25, No. 3, pp. 103–125.
- Hughes, Paul (2006) “Do EPA defendants prefer Republicans? Evidence from the 2000 election,” *Economic Inquiry*, Vol. 44, No. 3, pp. 579–585.
- Hyndman, Kyle (2008) “Disagreement in bargaining: An empirical analysis of OPEC,” *International Journal of Industrial Organization*, Vol. 26, No. 3, pp. 811–828.

- Imai, Masami and Cameron A. Shelton (2011) “Elections and political risk: New evidence from the 2008 Taiwanese Presidential Election,” *Journal of Public Economics*, Vol. 95, No. 7, pp. 837–849.
- Jackwerth, Jens Carsten and Mark Rubinstein (1996) “Recovering probability distributions from option prices,” *The Journal of Finance*, Vol. 51, No. 5, pp. 1611–1631.
- Jayachandran, Seema (2006) “The Jeffords effect,” *The Journal of Law and Economics*, Vol. 49, No. 2, pp. 397–425.
- Jiang, George J. and Yisong S. Tian (2005) “The model-free implied volatility and its information content,” *The Review of Financial Studies*, Vol. 18, No. 4, pp. 1305–1342.
- Johnson, Travis L. and Eric C. So (2012) “The option to stock volume ratio and future returns,” *Journal of Financial Economics*, Vol. 106, No. 2, pp. 262–286.
- Känzig, Diego R. (2019) “The macroeconomic effects of oil supply news: Evidence from OPEC announcements,” working paper.
- Karpoff, Jonathan M., Paul H. Malatesta, and Ralph A. Walkling (1996) “Corporate governance and shareholder initiatives: Empirical evidence,” *Journal of Financial Economics*, Vol. 42, No. 3, pp. 365–395.
- Kelly, Bryan, Lubos Pastor, and Pietro Veronesi (2016) “The price of political uncertainty: Theory and evidence from the option market,” *Journal of Finance*, Vol. 71, No. 5, pp. 2417–2480.
- Knight, Brian (2006) “Are policy platforms capitalized into equity prices? Evidence from the Bush/Gore 2000 Presidential Election,” *Journal of Public Economics*, Vol. 90, No. 4–5, pp. 751–773.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman (2017) “Technological innovation, resource allocation, and growth,” *Quarterly Journal of Economics*, Vol. 132, No. 2, pp. 665–712.
- Kothari, S. P. and Jerold B. Warner (1997) “Measuring long-horizon security price performance,” *Journal of Financial Economics*, Vol. 43, No. 3, pp. 301–339.
- Kothari, S.P. and Jerold B. Warner (2007) “Econometrics of event studies,” in B. Espen Eckbo ed. *Handbook of Empirical Corporate Finance*, San Diego: Elsevier, pp. 3–36, doi:10.1016/B978-0-444-53265-7.50015-9.

- Lamdin, Douglas J. (2001) “Implementing and interpreting event studies of regulatory changes,” *Journal of Economics and Business*, Vol. 53, No. 2–3, pp. 171–183.
- Lanen, William N. and Rex Thompson (1988) “Stock price reactions as surrogates for the net cash flow effects of corporate policy decisions,” *Journal of Accounting and Economics*, Vol. 10, No. 4, pp. 311–334.
- Lange, Ian and Joshua Linn (2008) “Bush v. Gore and the effect of New Source Review on power plant emissions,” *Environmental and Resource Economics*, Vol. 40, No. 4, pp. 571–591.
- Lauenstein, Philipp and André Küster Simic (2017) “Information processing in freight and freight forward markets: an event study on OPEC announcements,” *International Journal of Financial Markets and Derivatives*, Vol. 6, No. 2, p. 149.
- Lee, David S. (2008) “Randomized experiments from non-random selection in U.S. House elections,” *Journal of Econometrics*, Vol. 142, No. 2, pp. 675–697.
- Lee, David S. and Alexandre Mas (2012) “Long-run impacts of unions on firms: New evidence from financial markets, 1961–1999,” *The Quarterly Journal of Economics*, Vol. 127, No. 1, pp. 333–378.
- Lemoine, Derek (2017) “Green expectations: Current effects of anticipated carbon pricing,” *Review of Economics and Statistics*, Vol. 99, No. 3, pp. 499–513.
- Lin, Sharon Xiaowen and Michael Tamvakis (2010) “OPEC announcements and their effects on crude oil prices,” *Energy Policy*, Vol. 38, No. 2, pp. 1010–1016.
- Linn, Joshua (2010) “The effect of cap-and-trade programs on firms’ profits: Evidence from the Nitrogen Oxides Budget Trading Program,” *Journal of Environmental Economics and Management*, Vol. 59, No. 1, pp. 1–14.
- Loutia, Amine, Constantin Mellios, and Kostas Andriosopoulos (2016) “Do OPEC announcements influence oil prices?” *Energy Policy*, Vol. 90, pp. 262–272.
- MacKinlay, A. Craig (1997) “Event studies in economics and finance,” *Journal of Economic Literature*, Vol. 35, No. 1, pp. 13–39.
- Madhavan, Ananth (2000) “Market microstructure: A survey,” *Journal of Financial Markets*, Vol. 3, No. 3, pp. 205–258.

- Malatesta, Paul H. and Rex Thompson (1985) “Partially anticipated events: A model of stock price reactions with an application to corporate acquisitions,” *Journal of Financial Economics*, Vol. 14, No. 2, pp. 237–250.
- Martin, Ian (2017) “What is the expected return on the market?” *The Quarterly Journal of Economics*, Vol. 132, No. 1, pp. 367–433.
- Melick, William R. and Charles P. Thomas (1997) “Recovering an asset’s implied PDF from option prices: An application to crude oil during the Gulf crisis,” *Journal of Financial and Quantitative Analysis*, Vol. 32, No. 1, pp. 91–115.
- Meng, Kyle C. (2017) “Using a free permit rule to forecast the marginal abatement cost of proposed climate policy,” *American Economic Review*, Vol. 107, No. 3, pp. 748–784.
- Mukanjari, Samson and Thomas Sterner (2018) “Do markets trump politics? Evidence from fossil market reactions to the Paris Agreement and the U.S. election,” Working Papers in Economics 728, University of Gothenburg, Department of Economics.
- Pan, Jun and Allen M. Poteshman (2006) “The information in option volume for future stock prices,” *The Review of Financial Studies*, Vol. 19, No. 3, pp. 871–908, Publisher: Oxford Academic.
- Patell, James M. and Mark A. Wolfson (1979) “Anticipated information releases reflected in call option prices,” *Journal of Accounting and Economics*, Vol. 1, No. 2, pp. 117–140.
- (1981) “The ex ante and ex post price effects of quarterly earnings announcements reflected in option and stock prices,” *Journal of Accounting Research*, Vol. 19, No. 2, pp. 434–458.
- Plante, Michael (2019) “OPEC in the news,” *Energy Economics*, Vol. 80, pp. 163–172.
- Ramelli, Stefano, Alexander F. Wagner, Richard J. Zeckhauser, and Alexandre Ziegler (2019) “Stock price rewards to climate saints and sinners: Evidence from the Trump election,” Working Paper 25310, National Bureau of Economic Research.
- Rietz, Thomas A. (1988) “The equity risk premium: A solution,” *Journal of Monetary Economics*, Vol. 22, No. 1, pp. 117–131.

- Roberts, Brian E. (1990) “Political institutions, policy expectations, and the 1980 election: A financial market perspective,” *American Journal of Political Science*, Vol. 34, No. 2, pp. 289–310.
- Rubinstein, Mark (1994) “Implied binomial trees,” *The Journal of Finance*, Vol. 49, No. 3, pp. 771–818.
- Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz (2007) “Partisan impacts on the economy: Evidence from prediction markets and close elections,” *The Quarterly Journal of Economics*, Vol. 122, No. 2, pp. 807–829.
- (2011) “How prediction markets can save event studies,” in L. V. Williams ed. *Prediction Markets: Theory and Applications*, New York: Routledge, pp. 18–34.
- Spencer, Simon and Don Bredin (2019) “Agreement matters: OPEC announcement effects on WTI term structure,” *Energy Economics*, Vol. 80, pp. 589–609.
- van Tassel, Peter (2016) “Merger options and risk arbitrage,” Technical Report 761, Federal Reserve Bank of New York.
- Vayanos, Dimitri and Jiang Wang (2013) “Market liquidity—Theory and empirical evidence,” in George M. Constantinides, Milton Harris, and Rene M. Stulz eds. *Handbook of the Economics of Finance*, Vol. 2: Elsevier, pp. 1289–1361.
- Wagner, Alexander F., Richard J. Zeckhauser, and Alexandre Ziegler (2018a) “Unequal rewards to firms: Stock market responses to the Trump election and the 2017 corporate tax reform,” *AEA Papers and Proceedings*, Vol. 108, pp. 590–596.
- (2018b) “Company stock price reactions to the 2016 election shock: Trump, taxes, and trade,” *Journal of Financial Economics*, Vol. 130, No. 2, pp. 428–451.
- Wirl, Franz and Azra Kujundzic (2004) “The impact of OPEC conference outcomes on world oil prices 1984-2001,” *The Energy Journal*, Vol. 25, No. 1, pp. 45–62.
- Wolfers, Justin and Eric Zitzewitz (2009) “Using markets to inform policy: The case of the Iraq war,” *Economica*, Vol. 76, No. 302, pp. 225–250.
- Zhou, Hao (2018) “Variance risk premia, asset predictability puzzles, and macroeconomic uncertainty,” *Annual Review of Financial Economics*, Vol. 10, No. 1, pp. 481–497.

Appendix

Appendices A and B contain supporting material and supplementary analyses for the 2016 election and OPEC applications, respectively. The final two appendices contain additional theoretical analysis and derivations.

A Supporting Analysis for the 2016 Election

Figure A-1 depicts the distribution of event-day t-statistics (i.e., of t-statistics on $\hat{\gamma}_{i2}$). We see that the election was an unusually large event for an unusual number of firms. Very many firms' t-statistics exceeded even enormous values such as 50, indicating that an event of the election's magnitude is far outside the usual variation. Further the distribution is skewed, containing more large positive than large negative t-statistics. The tendency of stock prices to increase upon learning the election outcome is helpful for our approach to estimating \bar{p} .

A.1 Robustness Tests for \tilde{p}

In order to understand how option prices evolve over time and the extent to which our estimates of the likelihood of the 2016 election outcome using \bar{p} deviate from normal

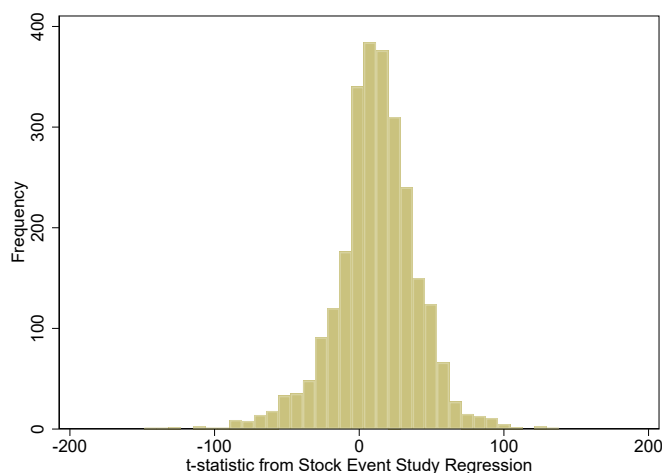


Figure A-1: Distribution of statistical significance of firms' responses to the 2016 U.S. election.

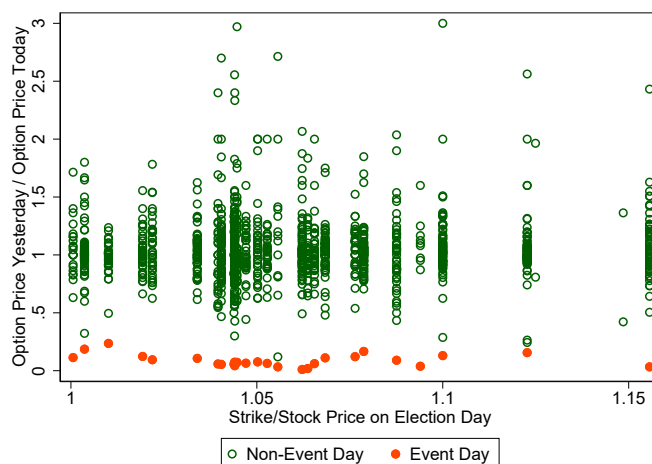


Figure A-2: Distribution of option prices on day $t - 1$ to option prices on day t . For non-event and event days for firms with election stock event-study t-statistics of 75 or higher.

option price changes, Figure A-2 shows how the raw (non-regression-corrected) ratios C_{t-1}/C_t vary on non-event day (green open circles) and event days (orange filled circles) for firms that had an election-day stock price event-study t-statistic above 75. The change in option prices on the day after the election was extreme relative to normal daily variation in option prices for all of these firms.

Table A-1 shows how our results change if we limit our analysis to firms with stock returns that were highly positively correlated with the PredictIt contract that Donald Trump would win the 2016 U.S. election. Our results are not substantially different when we limit our analysis to these firms, and we still see a substantial reduction in bias when we limit our analysis to firms with post-election stock returns that were both positive and highly statistically significant.

A.2 Robustness Tests for \tilde{p}

We now consider robustness of the estimated \tilde{p} to alternate methods and sample selection criteria.

First, consider an alternate approach to discretizing the integrals that lie at the heart of each firm's implied \tilde{p}_i . The approach in the main text remains purely model-free in avoiding any use of pricing models, instead discretizing the integral as is done

Table A-1: Comparison of \bar{p} for Firms with Different for 2016 Presidential Election

	2016 Election Stock Return Event Study			
	t-statistic > 20	t-statistic > 40	t-statistic > 60	t-statistic > 80
PredictIt t-statistic>1				
Probability	0.3941	0.2452	0.1571	0.1260
Standard Error	(0.0090)	(0.0109)	(0.0133)	(0.0198)
# Firms	221	76	20	10
PredictIt t-statistic>1.65				
Probability	0.3820	0.1982	0.1268	0.0914
Standard Error	(0.0183)	(0.0179)	(0.0172)	(0.0255)
# Firms	76	28	13	6
PredictIt t-statistic>1.96				
Probability	0.4415	0.2357	0.1434	–
Standard Error	(0.0230)	(0.0206)	(0.0324)	–
# Firms	51	17	5	0

Note: Estimated \bar{p} for regressions run on different subsets of options. We first ran regressions of stock returns on the value of the PredictIt contract that Trump would win the 2016 Presidential election. In the first panel we limit our analysis to firms with t-statistics on the PredictIt contract in this regression greater than 1 and 2016 election stock return event study t-statistics above the cutoffs in the column headers. The second panel limits our analysis to firms that were more highly correlated with the PredictIt contract, displaying a t-statistic greater than 1.65. The third panel is only for firms with PredictIt t-statistics over 1.96. All regressions are only for options that meet the standard restrictions from Table 1 and are the highest strike, well-traded call option for their firm.

in calculating the VIX. We here present an alternate approach that draws on the model-free variance swap pricing literature (Carr and Wu, 2009):¹

1. We convert observed option prices into implied volatilities. We obtain equity options' implied volatilities directly from OptionMetrics (following Carr and Wu, 2009). This approach strips away early exercise premia, yielding the European-style option prices required by the variance swap formula (9). These premia tend to be small in practice.
2. We calculate the forward price by adapting the methodology used to construct the familiar VIX index.²
3. We drop options that are likely to be illiquid, as in the main text.
4. We construct a fine grid of option prices by linearly interpolating implied volatilities at 10,000 different moneyness levels.³ For moneyness levels below (above) the smallest (largest) observed moneyness, we use the implied volatility from the smallest (largest) observed liquid strike.
5. We convert back to option prices using the Black-Scholes formula for equities (following Carr and Wu, 2009).
6. We calculate the variance swap rate, using $R_{\tau-1,T}$ from a daily version of the 3-month LIBOR rate (as in the main text).

To calculate $V_{i(\tau-1,T)}^H$, we use counterfactual implied volatilities predicted from a regression like (11), modified to allow the event effect (β_{iK}) to vary by firm and

¹This approach has been considered model-free because the critical pricing formula makes no distributional assumptions. It uses traditional pricing models only to transform the data into a form amenable to interpolation on a fine grid and then to transform the interpolated data back to option prices. Importantly, the pricing models' assumptions are not imposed across options. For instance, implied volatility is not restricted to be constant across strikes. If there is no early exercise premium, then all the pricing model does is allow interpolation to be conducted in implied volatility space rather than in price space. The former possess better properties. The main drawback to using this model-based discretization procedure is that calculated early exercise premia could have different types of errors just before and just after the event.

²We linearly interpolate implied volatilities between observed strikes, convert back to option prices, and seek the strike price at which the prices of calls and puts are most similar. See Cboe (2019) for more details.

³Following Carr and Wu (2009), time t moneyness is $\ln(K/F_{it})$, where F_{it} is the forward price for firm i . We bound the domain of strikes by the more extreme of eight standard deviations from the observed price of the underlying and the farthest liquid strike.

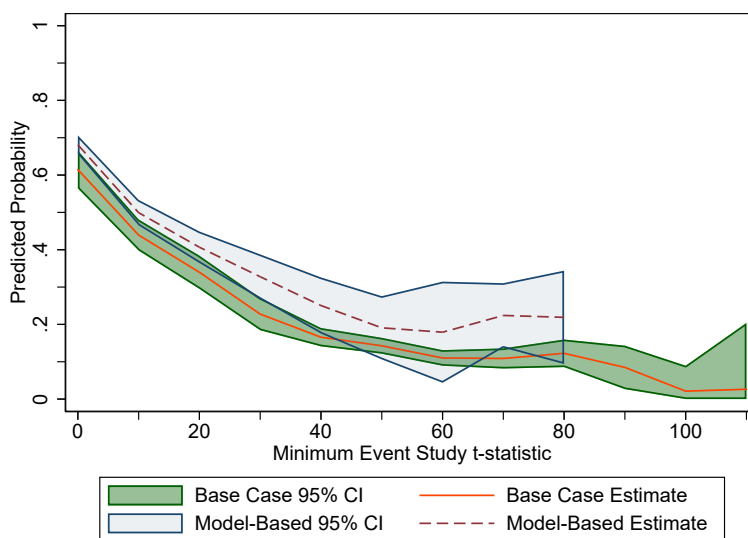


Figure A-3: Sensitivity of \hat{p} to discretization methods for 2016 election

strike and to use implied volatility as the dependent variable.⁴ A regression like (12) yields $S_{i(\tau-1)}^H$.

Figure A-3 compares the two approaches. The central estimate falls more slowly under the model-based discretization procedure and converges to a larger value (around 0.2). However, its confidence interval is also larger, so that the new estimates are not significantly different from the ones presented in the main text. Table A-2 provides more details. We see that the model-based discretization method draws on substantially fewer firms, in large part because of issues with the implied volatility variable. On the whole, given the statistical insignificance of the difference between the converged central estimates, we conclude that changing the discretization procedure does not have first-order consequences.

Second, Table A-3 assesses sensitivity to the minimum number of liquid strikes required for each firm. In this and subsequent tables, the rows vary the requirement on the (absolute value of) the t-statistic on $\hat{\gamma}_{i2}$. The main text required only 3 observed, liquid strikes around the event day (as in the first column here). The table

⁴As before, we weight by the inverse of the relative bid-ask spread averaged over days $t-1$ and t , with a weight of zero assigned if either day has a bid of zero. When calculating counterfactual stock prices and option implied volatilities, we use a second-order Taylor expansion that adjusts for the standard error of the estimated event effect.

Table A-2: Sensitivity of \tilde{p} to Discretization Method for 2016 Presidential Election

	Standard Restrictions	t-statistic > 25	t-statistic > 50	t-statistic > 75	t-statistic > 100
Variance Swap Approach: Base Case					
Probability	0.6148	0.2837	0.1428	0.1183	0.0213
Standard Error	(0.0258)	(0.0220)	(0.0109)	(0.0173)	(0.0347)
# Firms	372	105	37	14	3
# Options	3,973	1,024	356	118	24
# Firm-Days	74,028	20,895	7,363	2,786	597
# Option-Days	152,095	40,365	11,903	3,873	856
Variance Swap Approach: Model-Based Discretization					
Probability	0.6811	0.3712	0.1909	0.1945	0.2773
Standard Error	(0.0128)	(0.0250)	(0.0432)	(0.0468)	(0.0102)
# Firms	360	75	18	5	1
# Options	4,655	978	244	62	14
# Firm-Days	71,640	14,925	3,582	995	199
# Option-Days	158,614	34,210	7,597	1,923	245

shows that the estimated \tilde{p} are largely not sensitive to tightening this requirement. The exception is at the most stringent requirements. In particular, combinations of stringent requirements and high t-statistic cutoffs limit the sample to only 2 firms, and these 2 firms can happen to have very different \tilde{p}_i than the averaged \tilde{p} reported under less stringent requirements. This discrepancy is not troubling because we do not want to base our estimates on cases with so few firms.

Third, Table A-4 provides evidence on the identifying assumption underlying the estimated \tilde{p} . As described in the main text, we can approximate the difference between the conditional and unconditional variance of stock prices by using pre-election and post-election realized stock price volatility. The table restricts attention to the indicated fraction of firms with the smallest (in absolute value) difference between the approximated variances, which theory suggests are the firms for which \tilde{p}_i is a less biased measure of the probability of interest. The first column replicates prior results, keeping all firms regardless of their approximated variance. The second column drops the 25% of firms with the largest difference in realized volatility. Results are largely unchanged, although the smaller sample yields larger standard errors. The third and fourth columns drop the 50% and 75% of firms with the largest difference in volatility, and we again see only small effects, except for at the higher t-statistics where the central estimate is now a bit smaller. While we cannot completely verify the identifying assumption, this test suggests that violations are not driving the

results.

Table A-5 restricts attention only to firms whose stocks were sufficiently correlated with the PredictIt contract for the 2016 election in the period prior to the election, as measured by the absolute value of the t-statistic on a given stock's PredictIt coefficient. If such firms also have large event-day t-statistics, then we can be fairly sure that these were due to the election outcome and not to other events on election day. The first column replicates prior results. The first three rows show that correlation with the PredictIt contract does not completely identify exposed firms, in part because that contract may in turn merely be responding to economic news. There is additional information in stocks' event-day responses. If we require a t-statistic of 1.65 in a regression of stock prices on the PredictIt contract, then we are down to four or fewer firms if we also require the event-day t-statistic to be at least 40. Reassuringly, the estimates that retain a bit more firms (for cutoffs of 30 and 40) are quite close to the estimates that forgo use of PredictIt but impose tight t-statistic requirements. Requiring a PredictIt t-statistic of at least 1.96 leaves more than one firm only at event-day t-statistics of 20 or less, and these estimates do not appear to be useful. The takeaways are that the PredictIt t-statistic does not appear sufficient to identify the firms exposed to the event, it may generate complementary information that softens the event study t-statistic requirement, and it does not generate enough additional power over the event-day t-statistic to be worth the loss of firms from the sample. That tradeoff could work differently for other events in which the event study t-statistics were less extreme, in which case researchers would be less sure that firms with large responses were reflecting the event of interest rather than some other concurrent event.

Table A-3: Sensitivity of estimated \tilde{p} to minimum strike requirement, for 2016 election.

Min Abs Stock t-stat	Minimum Strikes Required				
	3	5	10	15	20
0	0.61 (0.026)	0.59 (0.021)	0.57 (0.035)	0.57 (0.029)	0.56 (0.037)
10	0.44 (0.021)	0.39 (0.02)	0.36 (0.031)	0.38 (0.037)	0.37 (0.046)
20	0.34 (0.023)	0.3 (0.021)	0.26 (0.031)	0.29 (0.037)	0.31 (0.048)
30	0.23 (0.022)	0.21 (0.021)	0.21 (0.035)	0.24 (0.043)	0.26 (0.057)
40	0.17 (0.013)	0.12 (0.015)	0.14 (0.028)	0.17 (0.046)	0.17 (0.051)
50	0.14 (0.011)	0.12 (0.019)	0.14 (0.036)	0.14 (0.048)	0.14 (0.05)
60	0.11 (0.011)	0.098 (0.021)	0.12 (0.04)	0.12 (0.053)	0.12 (0.053)
70	0.11 (0.014)	0.1 (0.029)	0.16 (0.085)	0.18 (0.15)	0.18 (0.15)
80	0.12 (0.019)	0.13 (0.044)	0.16 (0.085)	0.18 (0.15)	0.18 (0.15)

Standard errors in parentheses, clustered by date.

Table A-4: Sensitivity of estimated \tilde{p} to restricting to firms with the least difference between pre- and post-event volatility, for 2016 election.

Min Abs Stock t-stat	Fraction of Firms to Keep			
	1	0.75	0.5	0.25
0	0.61 (0.026)	0.62 (0.029)	0.62 (0.037)	0.63 (0.06)
10	0.44 (0.021)	0.43 (0.028)	0.44 (0.036)	0.44 (0.058)
20	0.34 (0.023)	0.33 (0.029)	0.33 (0.037)	0.34 (0.063)
30	0.23 (0.022)	0.24 (0.031)	0.24 (0.038)	0.29 (0.066)
40	0.17 (0.013)	0.16 (0.013)	0.15 (0.016)	0.18 (0.042)
50	0.14 (0.011)	0.14 (0.015)	0.12 (0.017)	0.17 (0.053)
60	0.11 (0.011)	0.091 (0.013)	0.063 (0.013)	0.075 (0.035)
70	0.11 (0.014)	0.087 (0.018)	0.062 (0.016)	0.089 (0.049)
80	0.12 (0.019)	0.1 (0.027)	0.06 (0.023)	0.11 (0.11)

Standard errors in parentheses, clustered by date.

Table A-5: Sensitivity of estimated \tilde{p} to PredictIt t-statistic requirement, for 2016 election.

Min Abs Stock t-stat	Min Abs PredictIt t-stat			
	0	1.65	1.96	2.6
0	0.61 (0.026)	0.66 (0.046)	0.67 (0.17)	0.64 (0.56)
10	0.44 (0.021)	0.47 (0.075)	0.44 (0.29)	0.4 (1.3)
20	0.34 (0.023)	0.32 (0.092)	0.34 (0.34)	0.4 (1.3)
30	0.23 (0.022)	0.053 (0.032)	0.062 (0.021)	–
40	0.17 (0.013)	0.053 (0.032)	0.062 (0.021)	–
50	0.14 (0.011)	0.056 (0.026)	0.062 (0.021)	–
60	0.11 (0.011)	0.056 (0.026)	0.062 (0.021)	–
70	0.11 (0.014)	0.056 (0.026)	0.062 (0.021)	–
80	0.12 (0.019)	0.053 (0.042)	–	–

Standard errors in parentheses, clustered by date.

B Supporting Analysis for OPEC Meetings

Table B-6 presents our estimated probability for each meeting’s outcome. The estimates align fairly well across the two approaches. However, the variance swap approach can estimate probabilities for only 6 of the 13 meetings, with the other meetings’ estimated probabilities being uninformative (i.e., \tilde{V} is negative; see Proposition 1). Most of the uninformative estimates occur early in our sample, when options on the front-month future do not have as many strikes with sufficient open interest.⁵

Figure B-4 plots the \bar{p} and \tilde{p} against each other, with circle sizes representing the average inverse of the standard error across approaches (so larger circles are more precisely estimated). Estimates tend to lie along the 45°line. Their unweighted correlation is 0.68 and their weighted correlation is 0.88. The coherence between the \bar{p} and \tilde{p} reassures that the \bar{p} are reasonable.

Figures B-5 and B-6 give plots analogous to Figure 4 for every OPEC meeting in our sample. Meetings where oil futures moved negatively with meeting news are plotted with the negative of the strike on the x-axis so that the figures always present deeper-out-of-the-money strikes to the right. Estimated probabilities do generally fall as we move to deeper-out-of-the-money strikes. This reduction in the estimated probability as we move further from the money is the most extreme for events with low predicted probabilities, as we would expect given that the bias in equation (5) is smaller when \bar{p} is larger. Individual strikes can produce noisy estimates, especially for deeper out-of-the-money strikes. Further, the pattern of decreasing \bar{p} can break down at the deepest-out-of-the-money strikes. The spline is important for smoothing out the noise between individual strikes’ estimates.

Figure B-8 plots the \bar{p} estimated from the Brent contract (as in the main text) against the \bar{p} estimated from the Intercontinental Exchange’s WTI contract. Estimates tend to lie along the 45°line, suggesting that the two data sets produce results that are roughly comparable. On average, options on WTI futures predict slightly higher event probabilities than do those on Brent futures, although this is not universally true. We use the Brent contract as our preferred specification because, as discussed in the main text, it is more directly linked to OPEC decisions.

⁵The small number of strikes also inflates standard errors on early years’ estimates in the out-of-the-money approach.

Table B-6: Estimated Probabilities for OPEC Meetings

	Approach	
	Out-of-the Money (\bar{p})	Variance Swap (\tilde{p})
December 10, 2016	0.4876 (0.0566)	0.1329 (0.2550)
November 30, 2016	0.2812 (0.0182)	0.4373 (0.0220)
September 28, 2016	0.2926 (0.0130)	–
June 2, 2016	0.9571 (0.0451)	0.9632 (0.0446)
December 4, 2015	0.9567 (0.0528)	0.9844 (0.0252)
June 5, 2015	0.7422 (0.0593)	0.5868 (0.1585)
November 27, 2014	0.2628 (0.1618)	–
December 4, 2013	0.6313 (0.1577)	–
May 31, 2013	0.5220 (0.0948)	–
December 10, 2012	0.6739 (0.0366)	–
June 14, 2012	0.6868 (0.2445)	–
December 14, 2011	0.3582 (0.0310)	–
June 8, 2011	0.4419 (0.1325)	0.8200 (0.2235)

Predicted OPEC meeting outcome probabilities using Brent front-month oil futures. Standard errors in parentheses. Events with missing \tilde{p} estimates either have negative values of \tilde{V} or, in one case, not enough liquidly traded strikes for the counterfactual calculation.

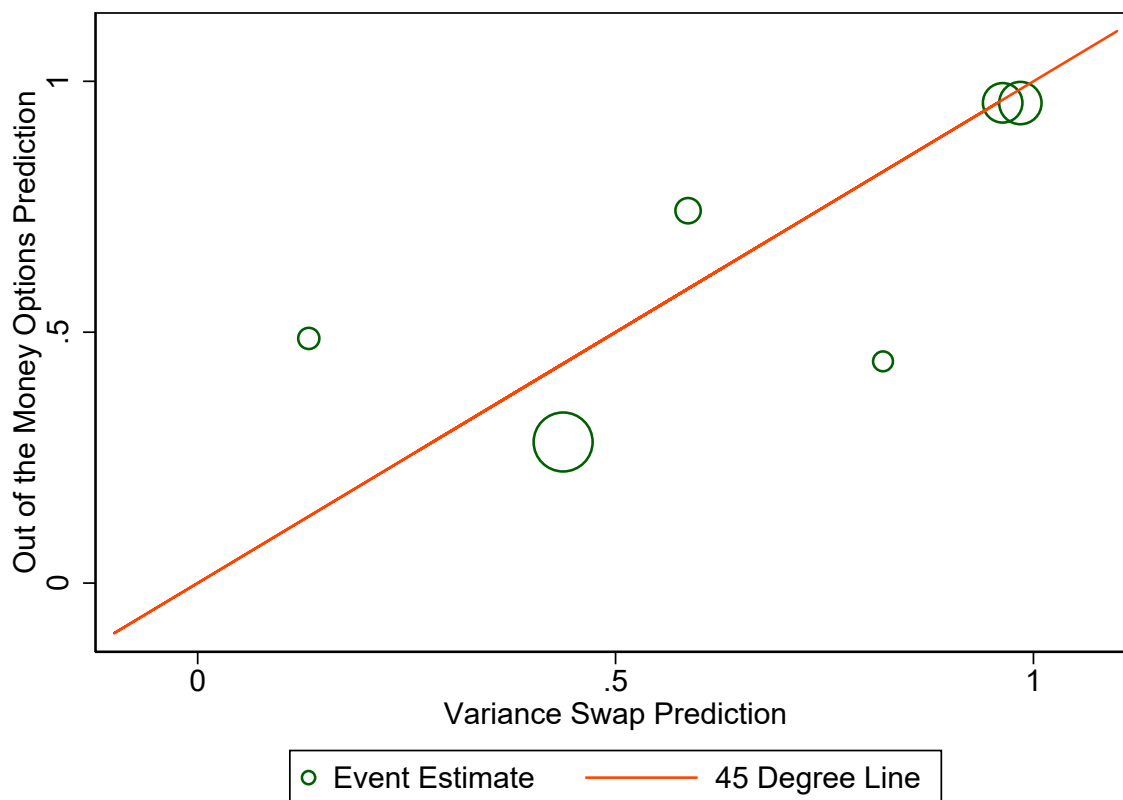


Figure B-4: Comparison Between \bar{p} and \tilde{p} Estimates

Each point is one of the 6 events for which we have an estimate of the OPEC meeting announcement probability using the variance swap approach (\tilde{p}). The vertical axis shows the estimate of \bar{p} using the out-of-the-money options approach, while the horizontal axis shows the estimate of \tilde{p} for the same event. We cannot reject that the estimates are equal for most events.

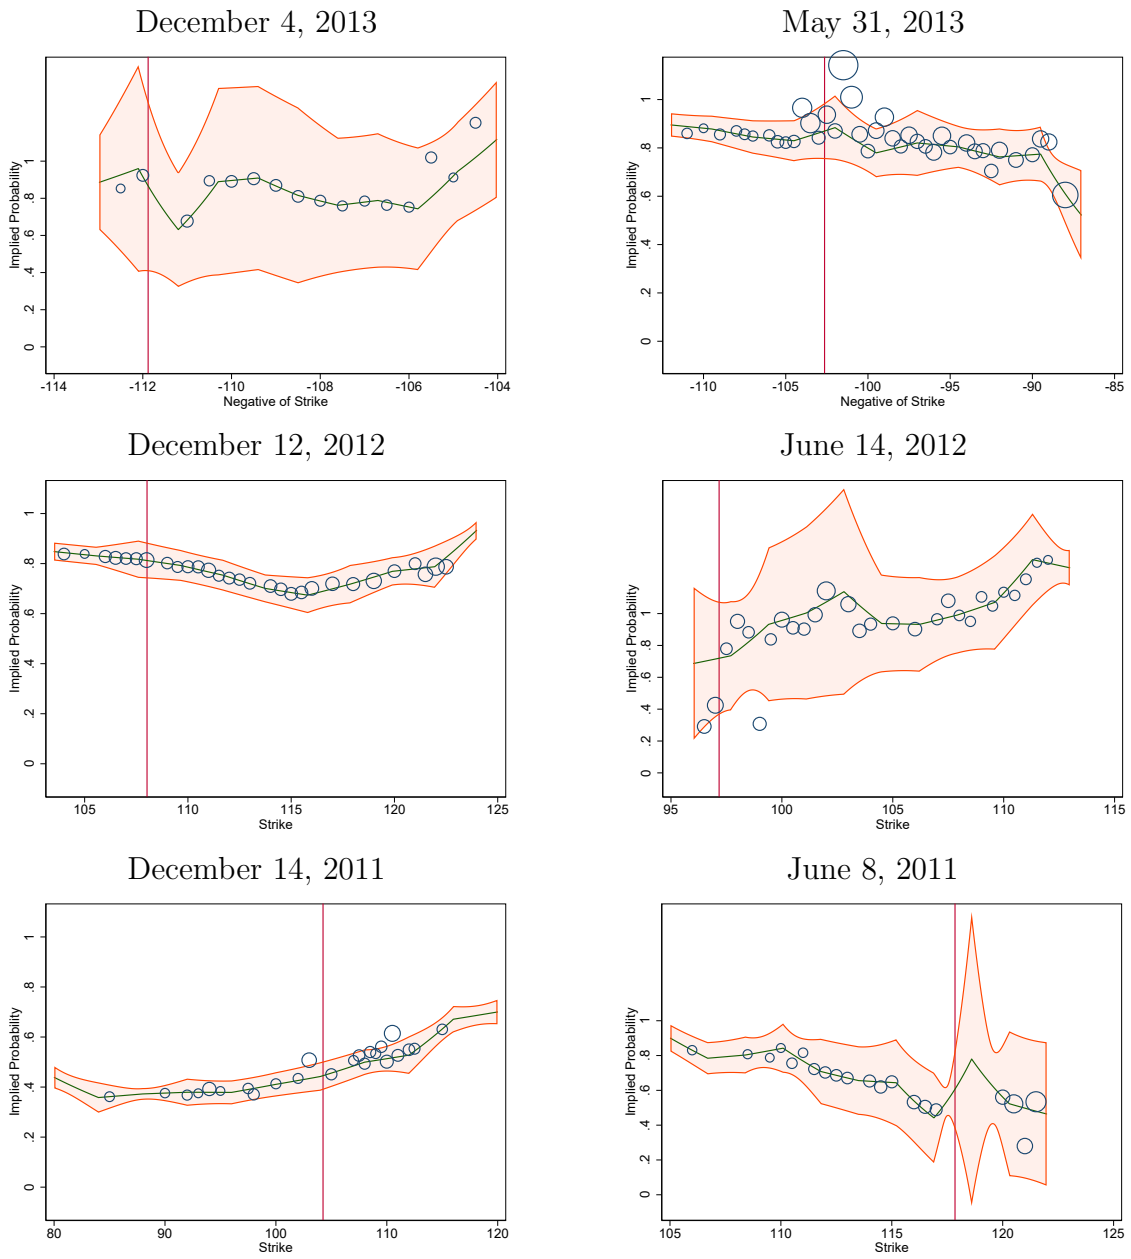


Figure B-5: Spline fits for additional OPEC meeting dates in 2011-2013.

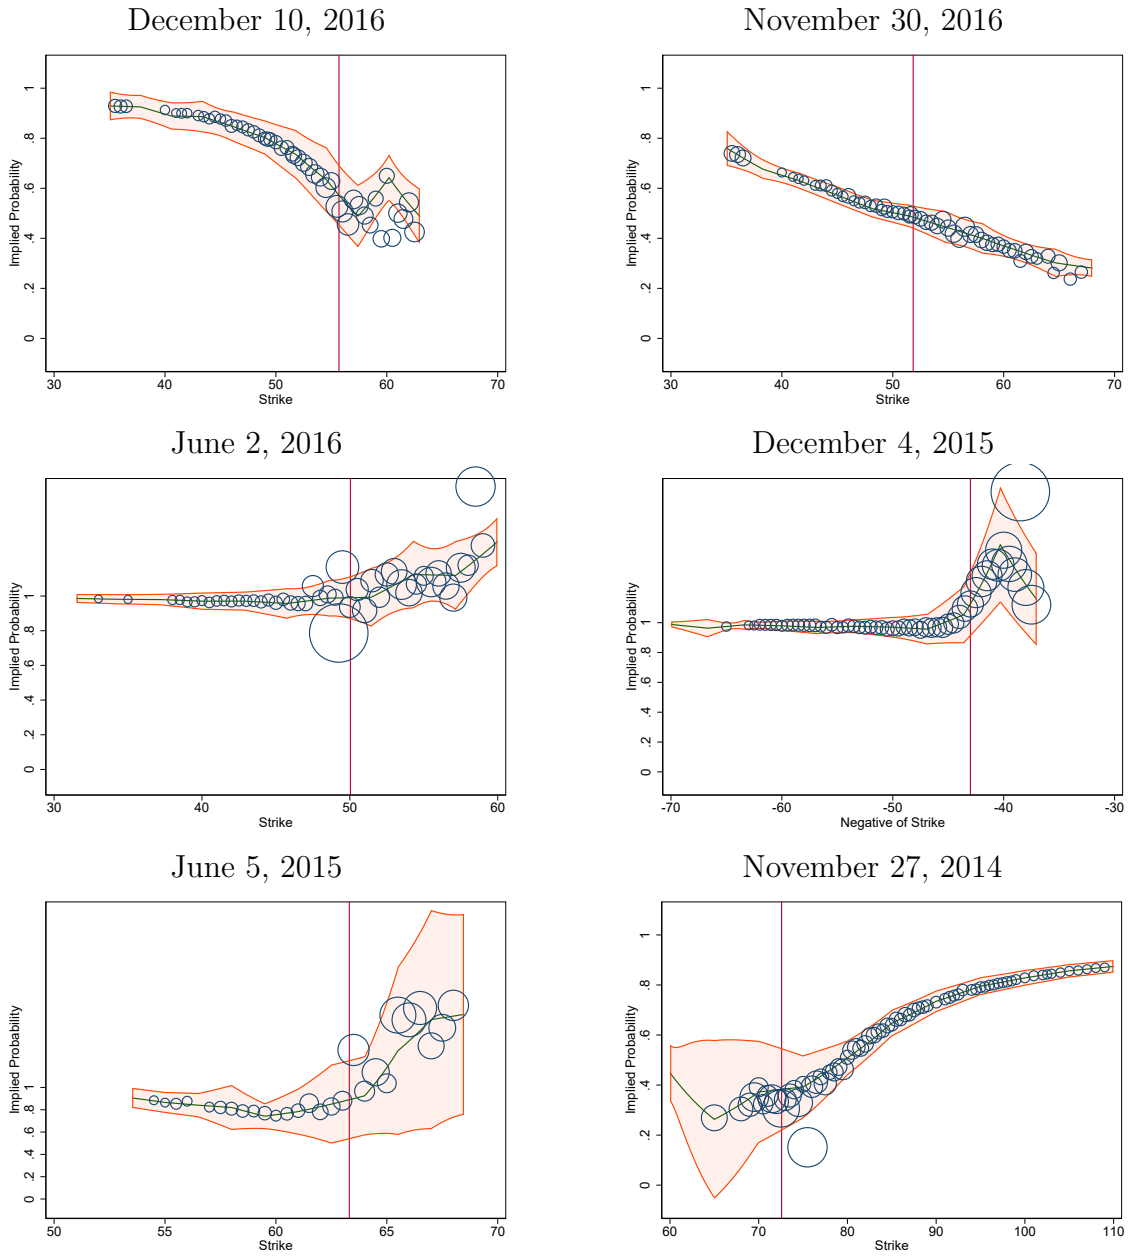


Figure B-6: Spline fits for additional OPEC meeting dates in 2014-2016.

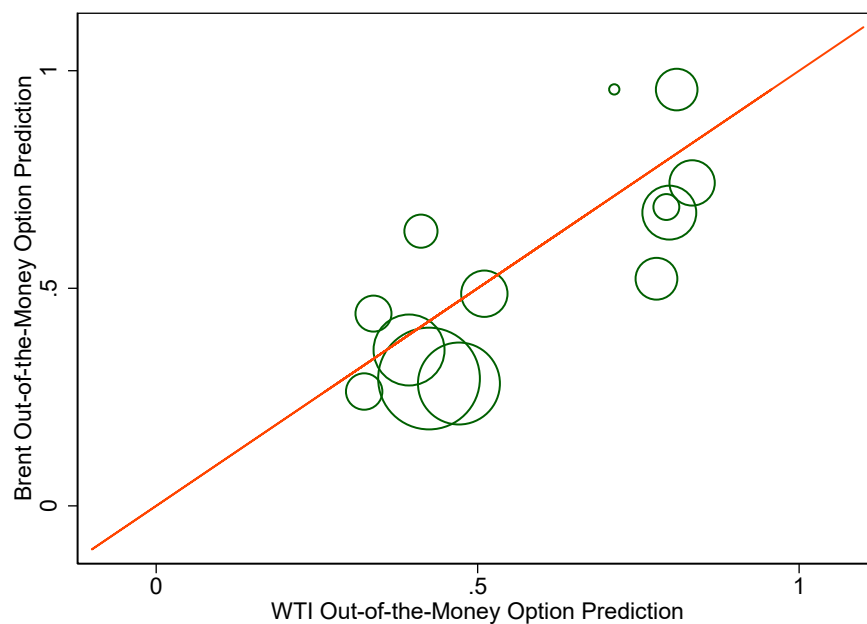


Figure B-7: Comparison Between Brent and WTI Estimates of \bar{p}

Probability estimated via the out-of-the-money option approach from options on Brent crude futures (vertical axis) and options on WTI futures (horizontal axis). Circle sizes are proportional to the inverse of the average standard error across the two data sets, so larger circles have smaller average standard errors. The orange line is the 45° line, where both data sets yield the same estimate.

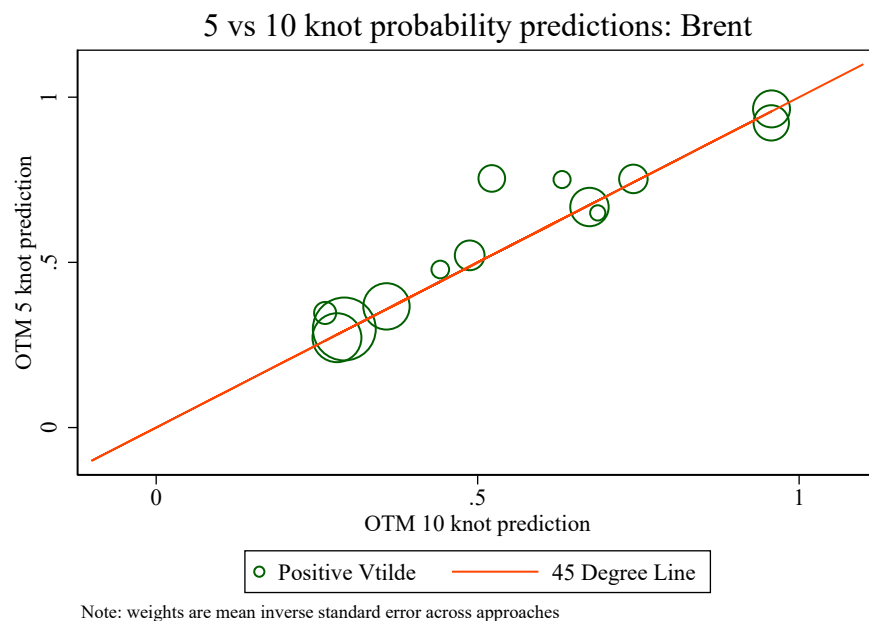


Figure B-8: Comparison Between 5-knot and 10-knot Estimates of \bar{p}

Probability estimated via the out-of-the-money option approach from options on Brent crude futures using 5 knot points (vertical axis) and 10 knot points (horizontal axis). Circle sizes are proportional to the inverse of the average standard error across the two data sets, so larger circles have smaller average standard errors. The orange line is the 45° line, where both knot definitions yield the same estimate.

C Extensions to Section 3.1

C.1 American Options

We have hitherto assumed that options are European-style options; however, the options in the data tend to be American-style options, which allow for early exercise. This appendix extends the theory of Section 3.1 to American-style options.

Melick and Thomas (1997) and Beber and Brandt (2006) express the price of an American-style option as a convex combination of upper and lower bounds that are tied to the price of a European option. Consider the price of an American-style call option (the analysis of puts will be similar), denoted with a tilde. Drawing on results from Chaudhury and Wei (1994), the option's price is

$$\tilde{C}_{x,T}(S_x, K) = \lambda R_{x,T} C_{x,T}(S_x, K) + (1 - \lambda) \max \{C_{x,T}(S_x, K), E_x[S_T] - K\},$$

for some $\lambda \in [0, 1]$. We can ignore the case with $C_{x,T}(S_x, K) < E_x[S_T] - K$: our nonparametric bound on $p_{\tau-1}^H$ is very loose for such in-the-money options, which is why we ignored such options in the empirical applications. For the options of interest, we therefore have:

$$\tilde{C}_{x,T}(S_x, K) = [\lambda R_{x,T} + (1 - \lambda)] C_{x,T}(S_x, K).$$

Now observe that

$$\frac{\tilde{C}_{\tau-1,T}(S_{\tau-1}, K)}{\tilde{C}_{\tau-1,T}^H(S_{\tau-1}^H, K)} = \frac{[\lambda R_{\tau-1,T} + (1 - \lambda)] C_{\tau-1,T}(S_{\tau-1}, K)}{[\lambda^H R_{\tau-1,T} + (1 - \lambda^H)] C_{\tau-1,T}^H(S_{\tau-1}^H, K)},$$

where we allow the weight λ to vary with e . As either $\lambda^H \rightarrow \lambda$ or $R_{\tau-1,T} \rightarrow 1$, we have:

$$\frac{\tilde{C}_{\tau-1,T}(S_{\tau-1}, K)}{\tilde{C}_{\tau-1,T}^H(S_{\tau-1}^H, K)} \rightarrow \frac{C_{\tau-1,T}(S_{\tau-1}, K)}{C_{\tau-1,T}^H(S_{\tau-1}^H, K)} = \bar{p},$$

where the right-hand side is the upper bound on $p_{\tau-1}^H$ derived in the main text. In these cases, it does not matter whether we estimate the upper bound on $p_{\tau-1}^H$ using American-style or European-style options. In general, we have:

$$\begin{aligned} \frac{\tilde{C}_{\tau-1,T}(S_{\tau-1}, K)}{\tilde{C}_{\tau-1,T}^H(S_{\tau-1}^H, K)} &\in \left[\frac{1}{R_{\tau-1,T}} \frac{C_{\tau-1,T}(S_{\tau-1}, K)}{C_{\tau-1,T}^H(S_{\tau-1}^H, K)}, R_{\tau-1,T} \frac{C_{\tau-1,T}(S_{\tau-1}, K)}{C_{\tau-1,T}^H(S_{\tau-1}^H, K)} \right] \\ &= \left[\frac{1}{R_{\tau-1,T}} \bar{p}, R_{\tau-1,T} \bar{p} \right]. \end{aligned}$$

The maximum possible error from estimating \bar{p} from American-style options is controlled by $R_{\tau-1,T} - 1$. In the empirical application, we focus on options with near expiration dates (smaller T) in order to limit the possible magnitude of this error.

C.2 When the Realized Event Was Not Extreme

We now consider how to obtain a tighter bound when the realized event is not extreme. Assume that we can partition the event space into $e \in \{L, M, H\}$ such that $S(\omega_t, M) > \bar{S}$ implies $S(\omega_t, H) > S(\omega_t, M) > S(\omega_t, L)$. Assume that $e = M$ is realized.⁶ We seek $p_{\tau-1}^M$.

At time $\tau - 1$, the price of a call option with strike K and expiration $T > \tau - 1$ must satisfy:

$$\begin{aligned} & C_{\tau-1,T}(S_{\tau-1}, K) \\ &= \frac{1}{R_{\tau-1,T}} \int_K^\infty (S_T - K) \left[p_{\tau-1}^L f_{\tau-1}(S_T | S_{\tau-1}^L, L) + p_{\tau-1}^M f_{\tau-1}(S_T | S_{\tau-1}^M, M) + p_{\tau-1}^H f_{\tau-1}(S_T | S_{\tau-1}^H, H) \right] dS_T. \end{aligned}$$

Consider buying a call option with strike K_1 and selling a call option with strike $K_2 > K_1$. Label this portfolio $\Gamma_{\tau-1,T}(S_{\tau-1}, K_1, K_2)$. The value of this portfolio is

$$\begin{aligned} & \Gamma_{\tau-1,T}(S_{\tau-1}, K_1, K_2) \\ & \triangleq C_{\tau-1,T}(S_{\tau-1}, K_1) - C_{\tau-1,T}(S_{\tau-1}, K_2) \\ &= \frac{1}{R_{\tau-1,T}} \int_{K_1}^{K_2} (S_T - K_1) \left[p_{\tau-1}^L f_{\tau-1}(S_T | S_{\tau-1}^L, L) + p_{\tau-1}^M f_{\tau-1}(S_T | S_{\tau-1}^M, M) + p_{\tau-1}^H f_{\tau-1}(S_T | S_{\tau-1}^H, H) \right] dS_T \\ & \quad + \frac{K_2 - K_1}{R_{\tau-1,T}} \int_{K_2}^\infty \left[p_{\tau-1}^L f_{\tau-1}(S_T | S_{\tau-1}^L, L) + p_{\tau-1}^M f_{\tau-1}(S_T | S_{\tau-1}^M, M) + p_{\tau-1}^H f_{\tau-1}(S_T | S_{\tau-1}^H, H) \right] dS_T. \end{aligned}$$

⁶If either $e = L$ or $e = H$ were realized, then the analysis in the main text holds, because we can combine $e = M$ with whichever other value for e was not realized. In addition, partitioning the event space into three possible values is not restrictive: if, for instance, there were $e \in \{L_1, L_2, M, H\}$ such that $S(\omega_t, M) > \bar{S}$ implies $S(\omega_t, H) > S(\omega_t, M) > S(\omega_t, L_1), S(\omega_t, L_2)$ and $e = M$ were realized, then we could combine L_1 and L_2 into a single indicator L .

Consider how the realization of the event changes the value of this portfolio:

$$\begin{aligned}
& \Gamma_{\tau-1,T}^M(S_{\tau-1}^M, K_1, K_2) - \Gamma_{\tau-1,T}(S_{\tau-1}, K_1, K_2) \\
&= (1 - p_{\tau-1}^M) \frac{1}{R_{\tau-1,T}} \int_{K_1}^{K_2} (S_T - K_1) [f_{\tau-1}(S_T | S_{\tau-1}^M, M) - f_{\tau-1}(S_T | S_{\tau-1}^M, \neg M)] \, dS_T \\
&\quad + (1 - p_{\tau-1}^M) \frac{K_2 - K_1}{R_{\tau-1,T}} \int_{K_2}^{\infty} [f_{\tau-1}(S_T | S_{\tau-1}^M, M) - f_{\tau-1}(S_T | S_{\tau-1}^M, \neg M)] \, dS_T \\
&= (1 - p_{\tau-1}^M) \Gamma_{\tau-1,T}^M(S_{\tau-1}^H, K_1, K_2) - (1 - p_{\tau-1}^M) \Gamma_{\tau-1,T}^{\neg M}(S_{\tau-1}^M, K_1, K_2),
\end{aligned}$$

where $\neg M$ means that $e \in \{L, H\}$. Of course $\Gamma_{\tau-1,T}^{\neg M}(S_{\tau-1}^M, K_1, K_2) \geq 0$. We then have:

$$p_{\tau-1}^M \leq \frac{\Gamma_{\tau-1,T}(S_{\tau-1}, K_1, K_2)}{\Gamma_{\tau-1,T}^M(S_{\tau-1}^M, K_1, K_2)}.$$

We again have an upper bound on the desired risk-neutral probability.⁷ The bound becomes tighter as $\Gamma_{\tau-1,T}^{\neg M}(S_{\tau-1}^M, K_1, K_2)$ becomes small, which occurs when $f_{\tau-1}(S_T | S_{\tau-1}^L, L) \rightarrow 0$ as S_T increases beyond K_1 . However, whereas the bound could become arbitrarily tight in the main text's case, the tightness of the bound is here limited by the fact that

$$\Gamma_{\tau-1,T}^{\neg M}(S_{\tau-1}^M, K_1, K_2) \geq \frac{K_2 - K_1}{R_{\tau-1,T}} \int_{K_2}^{\infty} [f_{\tau-1}(S_T | S_{\tau-1}^L, L) + f_{\tau-1}(S_T | S_{\tau-1}^H, H)] \, dS_T.$$

Intuitively, there is always probability mass from the distribution conditional on H present in the interval between K_1 and K_2 . The closer together are K_2 and K_1 , the greater the potential for the bound to be arbitrarily tight. In general, the upper bound on $p_{\tau-1}^M$ becomes tighter when neither event L nor event H gives much chance of S_T ending up between K_1 and K_2 .

D Derivations for the Variance Swap Analysis

D.1 Equation (7)

Noting that $R_{\tau-1,\tau} = 1$ implies $S_{\tau-1} = E_{\tau-1}[S_{\tau}]$, we have:

$$\text{Var}_{\tau-1}[S_{\tau}] = E_{\tau-1} [(S_{\tau})^2] - [S_{\tau-1}]^2.$$

⁷Intuitively, area A in Figure 1 is bounded on the left by K_1 and on the right by K_2 , instead of stretching all the way to infinity. The bound on $p_{\tau-1}^M$ becomes tighter when the distributions conditional on H and L do not have much mass between K_1 and K_2 .

The assumption that time τ of the event is known then implies

$$\begin{aligned} \text{Var}_{\tau-1}[S_\tau] &= p_{\tau-1}^H E_{\tau-1} [(S_\tau^H)^2] + (1 - p_{\tau-1}^H) E_{\tau-1} [(S_\tau^L)^2] - [S_{\tau-1}]^2 \\ &= p_{\tau-1}^H \text{Var}_{\tau-1} [S_\tau^H] + (1 - p_{\tau-1}^H) \text{Var}_{\tau-1} [S_\tau^L] \\ &\quad + p_{\tau-1}^H (S_{\tau-1}^H)^2 + (1 - p_{\tau-1}^H) (S_{\tau-1}^L)^2 - [S_{\tau-1}]^2 \\ &= \frac{p_{\tau-1}^H}{1 - p_{\tau-1}^H} (S_{\tau-1}^H - S_{\tau-1})^2 + p_{\tau-1}^H \text{Var}_{\tau-1} [S_\tau^H] + (1 - p_{\tau-1}^H) \text{Var}_{\tau-1} [S_\tau^L], \end{aligned}$$

where the last equality substitutes for $S_{\tau-1}^L$ from equation (1) and simplifies.

D.2 Proof of Proposition 1

Using the assumption that e will be known by time τ , we have:

$$\begin{aligned} V_{\tau-1,T} &= E_{\tau-1} \left[\left(\frac{S_\tau - S_{\tau-1}}{S_{\tau-1}} \right)^2 \right] + (1 - p_{\tau-1}^H) E_{\tau-1} \left[\sum_{j=1}^{T-\tau} \left(\frac{S_{\tau+j}^L - S_{\tau+j-1}^L}{\tilde{R}_{\tau-1,\tau+j-1} S_{\tau-1}} \right)^2 \right] \\ &\quad + p_{\tau-1}^H E_{\tau-1} \left[\sum_{j=1}^{T-\tau} \left(\frac{S_{\tau+j}^H - S_{\tau+j-1}^H}{\tilde{R}_{\tau-1,\tau+j-1} S_{\tau-1}} \right)^2 \right]. \end{aligned}$$

Therefore:

$$\begin{aligned} V_{\tau-1,T} - \left(\frac{S_{\tau-1}^H}{S_{\tau-1}} \right)^2 V_{\tau-1,T} &= E_{\tau-1} \left[\left(\frac{S_\tau - S_{\tau-1}}{S_{\tau-1}} \right)^2 - \left(\frac{S_\tau^H - S_{\tau-1}^H}{S_{\tau-1}} \right)^2 \right] \\ &\quad + (1 - p_{\tau-1}^H) E_{\tau-1} \left[\sum_{j=1}^{T-\tau} \left\{ \left(\frac{S_{\tau+j}^L - S_{\tau+j-1}^L}{\tilde{R}_{\tau-1,\tau+j-1} S_{\tau-1}} \right)^2 - \left(\frac{S_{\tau+j}^H - S_{\tau+j-1}^H}{\tilde{R}_{\tau-1,\tau+j-1} S_{\tau-1}} \right)^2 \right\} \right]. \end{aligned} \tag{D-1}$$

Analyze the first term on the right-hand side. Because $E_{\tau-1}[S_\tau] = \tilde{R}_{\tau-1,\tau} S_{\tau-1}$ and $E_{\tau-1}[S_\tau^H] = \tilde{R}_{\tau-1,\tau} S_{\tau-1}^H$, we have

$$\begin{aligned} &E_{\tau-1} \left[\left(\frac{S_\tau - S_{\tau-1}}{S_{\tau-1}} \right)^2 - \left(\frac{S_\tau^H - S_{\tau-1}^H}{S_{\tau-1}} \right)^2 \right] \\ &= \frac{1}{[S_{\tau-1}]^2} \left\{ E_{\tau-1} [(S_\tau)^2] + [S_{\tau-1}]^2 - 2\tilde{R}_{\tau-1,\tau} [S_{\tau-1}]^2 - E_{\tau-1} [(S_\tau^H)^2] - [S_{\tau-1}^H]^2 + 2\tilde{R}_{\tau-1,\tau} [S_{\tau-1}^H]^2 \right\}. \end{aligned}$$

Adding and subtracting $(1 - p_{\tau-1}^H)E_{\tau-1}[(S_{\tau}^L - S_{\tau-1}^L)^2 - (S_{\tau}^H - S_{\tau-1}^H)^2]$ and simplifying, we then obtain:

$$\begin{aligned} & E_{\tau-1} \left[\left(\frac{S_{\tau} - S_{\tau-1}}{S_{\tau-1}} \right)^2 - \left(\frac{S_{\tau}^H - S_{\tau-1}^H}{S_{\tau-1}} \right)^2 \right] \\ &= \frac{1}{[S_{\tau-1}]^2} \left\{ (1 - p_{\tau-1}^H) (E_{\tau-1}[(S_{\tau}^L - S_{\tau-1}^L)^2] - E_{\tau-1}[(S_{\tau}^H - S_{\tau-1}^H)^2]) \right. \\ & \quad \left. + (2\tilde{R}_{\tau-1,\tau} - 1) (p_{\tau-1}^H [S_{\tau-1}^H]^2 - [S_{\tau-1}]^2 + (1 - p_{\tau-1}^H) [S_{\tau-1}^L]^2) \right\}. \end{aligned}$$

Substituting into equation (D-1) and combining with the summation, we have:

$$\begin{aligned} V_{\tau-1,T} - \left(\frac{S_{\tau-1}^H}{S_{\tau-1}} \right)^2 V_{\tau-1,T}^H &= \frac{2\tilde{R}_{\tau-1,\tau} - 1}{[S_{\tau-1}]^2} \left\{ p_{\tau-1}^H [S_{\tau-1}^H]^2 - [S_{\tau-1}]^2 + (1 - p_{\tau-1}^H) [S_{\tau-1}^L]^2 \right\} \\ & \quad + (1 - p_{\tau-1}^H) E_{\tau-1} \left[\sum_{j=0}^{T-\tau} \left\{ \left(\frac{S_{\tau+j}^L - S_{\tau+j-1}^L}{\tilde{R}_{\tau-1,\tau+j-1} S_{\tau-1}} \right)^2 - \left(\frac{S_{\tau+j}^H - S_{\tau+j-1}^H}{\tilde{R}_{\tau-1,\tau+j-1} S_{\tau-1}} \right)^2 \right\} \right], \end{aligned}$$

which in turn implies:

$$\begin{aligned} [S_{\tau-1}]^2 V_{\tau-1,T} - [S_{\tau-1}^H]^2 V_{\tau-1,T}^H &= (2\tilde{R}_{\tau-1,\tau} - 1) (p_{\tau-1}^H [S_{\tau-1}^H]^2 - [S_{\tau-1}]^2 + (1 - p_{\tau-1}^H) [S_{\tau-1}^L]^2) \\ & \quad + (1 - p_{\tau-1}^H) ([S_{\tau-1}^L]^2 V_{\tau-1,T}^L - [S_{\tau-1}^H]^2 V_{\tau-1,T}^H). \end{aligned}$$

Substituting for $S_{\tau-1}^L$ from equation (1) and rearranging, we obtain:

$$\frac{p_{\tau-1}^H}{1 - p_{\tau-1}^H} = \frac{[S_{\tau-1}]^2 V_{\tau-1,T} - [S_{\tau-1}^H]^2 V_{\tau-1,T}^H}{(2\tilde{R}_{\tau-1,\tau} - 1) [S_{\tau-1}^H - S_{\tau-1}]^2} + (1 - p_{\tau-1}^H) \frac{[S_{\tau-1}^H]^2 V_{\tau-1,T} - [S_{\tau-1}^L]^2 V_{\tau-1,T}^L}{(2\tilde{R}_{\tau-1,\tau} - 1) [S_{\tau-1}^H - S_{\tau-1}]^2}. \quad (\text{D-2})$$

This is the analogue of equation (8), adapted for the possibility that $\tilde{R}_{\tau-1,T} > 1$ and for the use of simple variance swaps. Denote the unobserved term $[S_{\tau-1}^H]^2 V_{\tau-1,T}^H - [S_{\tau-1}^L]^2 V_{\tau-1,T}^L$ by x . The first part of the proposition follows taking a first-order Taylor approximation around $x = 0$, with the derivative of $p_{\tau-1}^H$ with respect to $([S_{\tau-1}^H]^2 V_{\tau-1,T}^H - [S_{\tau-1}^L]^2 V_{\tau-1,T}^L)$ following from applying the the implicit function theorem to equation (D-2):

$$p_{\tau-1}^H = \tilde{p} + \frac{[1 - p_{\tau-1}^H] [(S_{\tau-1}^H)^2 V_{\tau-1,T}^H - (S_{\tau-1}^L)^2 V_{\tau-1,T}^L]}{[2\tilde{R}_{\tau-1,\tau} - 1] [S_{\tau-1}^H - S_{\tau-1}^L]^2} + O \left([(S_{\tau-1}^H)^2 V_{\tau-1,T}^H - (S_{\tau-1}^L)^2 V_{\tau-1,T}^L]^2 \right),$$

using $[S_{\tau-1}^H - S_{\tau-1}]^2 = (1 - p_{\tau-1}^H)^2 [S_{\tau-1}^H - S_{\tau-1}^L]^2$. The second part of the proposition follows from solving for $p_{\tau-1}^H$ in equation (D-2) with assumptions on the relationship between $[S_{\tau-1}^L]^2 V_{\tau-1,T}^L$ and $[S_{\tau-1}^H]^2 V_{\tau-1,T}^H$. The third part of the proposition follows from observing that $\tilde{V} < 0$ implies that either $\tilde{p} < 0$ or $\tilde{p} > 1$.