

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

Ashley Langer[†] and Derek Lemoine[‡]

Preliminary

December 24, 2019

The event study methodology is widely used in economics and finance to estimate the market consequences of news. However, event studies cannot recover the full effects of events that the market anticipated might happen. 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. We estimate that the 2016 U.S. election outcome had only a 10% chance of occurring, and we find that most OPEC meetings fail to move oil markets in part because their outcomes are well-anticipated.

JEL: C58, G13, L71, Q43

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

*We are grateful to David Brown, David Card, Gautam Gowrisankaran, Todd Sorensen, and Tiemen Woutersen for discussions and feedback, 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 for excellent research assistance. We also thank seminar participants at HEC Montreal, TREE, University of Arizona, University of California Berkeley, and Yale.

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

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

1 Introduction

The event study methodology has been widely used to understand the consequences of news for firms' stock prices. These consequences illuminate the costs and benefits of policies as well as the value of mergers, patents, and much more. However, it is widely recognized that the event study methodology captures the full market valuation of news only when that news is a complete surprise. Such cases are rare. Most events are at least partially anticipated. In that case, event studies measure the effect of becoming sure about the news rather than of learning news that is completely new. The event study methodology can then provide only a lower bound on the consequences of news—and that bound can be arbitrarily loose.

We extend the event study methodology to capture the full effects of partially anticipated events. We show that empirical researchers can recover the priced-in probability of a realized event by running event study regressions in the prices of financial options. In contrast to previous work, our methods do not require a model of option prices (i.e., we do not impose parametric assumptions on the distribution of stock prices). We apply our methods to recover probabilities for several events that are of broad economic importance. First, we validate our new methods by estimating the probability of the Republican sweep of the 2016 U.S. election. Prediction markets, bookmakers, and polling-driven models imply probabilities of ranging from well below 0.15 to as high as 0.29. We recover a probability of around 0.1. Second, we estimate probabilities for the outcomes of meetings of the Organization of the Petroleum Exporting Countries (OPEC) in 2007–2016. The only meetings that generated surprising news were a pair of 2016 meetings that led to large cuts in production. In contrast, market actors judged the large production cuts of 2008 to be more likely than not, and therefore their effect was already priced in to the market.

Our method directly recovers the ex-ante probability of uncertain events, which is important for understanding the role of uncertainty in the economy generally. Beyond that, our extension of the event study methodology is important for increasing the use of event studies in policy evaluation. Event studies are widely used techniques for obtaining revealed preference estimates of the costs and benefits of policy. Event studies have informed our understanding of the costs of minimum wage laws (Card and Krueger, 1997; Bell and Machin, 2017), unionization (Ruback and Zimmerman, 1984; Bronars and Deere, 1990; Lee and Mas, 2012), dividend taxes (Auerbach and Hassett, 2007), antitrust suits (Bizjak and Coles, 1995; Bittlingmayer and Hazlett, 2000), and product liability suits (Prince and Rubin, 2002). Event studies have also proved valuable in political economy, highlighting the value of political connections (Fisman, 2001; Fisman et al., 2012), the consequences of party politics (Hughes,

2006; Jayachandran, 2006; Snowberg et al., 2011), and the relevance of theories of regulatory capture (Dann and James, 1982; James, 1983). And event studies are critical tools for evaluating monetary policy (Bernanke and Kuttner, 2005), health care policy (Al-Ississ and Miller, 2013), and environmental policy (Lange and Linn, 2008; Linn, 2010; Bushnell et al., 2013).

However, the inability to identify the full magnitude of event effects casts a cloud over many of these applications, as most of the cited studies clearly acknowledge. For instance, Card and Krueger (1997)[314] admit 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”. In reviewing the event study literature, 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.

We develop two new model-free techniques for estimating the market (risk-neutral) probability of realized events from widely traded financial options. The intuition underlying our first approach is relatively straightforward. Imagine that there is an option that 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, as with an out-of-the-money call option and an election outcome that makes high stock prices substantially more likely. On the day before the event, the value of this option is the probability of the event occurring times the value of the option if the event occurs. On the day after the event occurs, the value of this option is simply the value of the option given that the event has occurred. Given the standard event study assumption that nothing else has changed over the short window, then 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.¹

Our second method of estimating an event’s probability relies on a different identifying assumption. A variance swap rate reveals the market’s expected variance of stock prices over some horizon. The pre-event variance swap rate includes the vari-

¹More generally, options with extreme strike prices may not be worthless if the event fails to occur because there may still be a 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.

ance induced by the event's realization, but the post-event variance swap rate does not. Using this insight, we show that variance swap rates can identify the priced-in probability of the event as long as the date of the event is known in advance and the expected variance of the stock price process is not affected by the event's outcome. In this case, differencing the pre- and post-event variance swap rates eliminates the post-event variance but retains the variance induced by uncertainty about the event's realization. Recent work shows that variance swap rates can be synthesized from a linear combination of option prices under quite general assumptions (Martin, 2017). By running event studies on the set of options traded on a firm, researchers can estimate the variance swap rate that would have been implied by option prices if nothing but the event had occurred. Our second method therefore again relies on changes in option prices to identify the event's probability, but now using the full set of option strikes traded on a firm, undertaking a different type of calculation with them, and relying on a different identifying assumption.

We validate our methods by estimating the probability of the 2016 U.S. election. As described above, we recover probabilities consistent with the range of contemporary estimates. The two methods generate nearly identical results, despite relying on different identifying assumptions. They also move in ways consistent with theory: we demonstrate that the out-of-the-money option approach is strongly biased upward unless using only extreme strikes and that some bias remains unless restricting attention to firms with especially large event-day stock price movements, and the variance swap approach is also biased upwards unless 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; Ramelli et al., 2018; Wagner et al., 2018a,b). Our results imply that the full effect of the election is more than 10% larger than implied by standard event study estimates.

We demonstrate the broad applicability of our methods by considering 30 OPEC meetings and announcements from 2007 to 2016. Very few of these meetings appear to move oil prices in any notable way. However, we show that very few of these meetings produce news that was even moderately surprising. The few events that did produce surprising news did have large effects on oil markets. Our two methods again largely cohere, but special cases illustrate advantages of having two methods, as one or the other may not apply to a given event or may generate weak estimates.

The usefulness of priced-in event probabilities

Recovering the priced-in probabilities of events is important for many types of questions. First, these probabilities are critical for recovering the full effects of events,

which are in turn critical for estimating corporate tax avoidance, policy cost pass-through, the value of mergers, and the effects of government policies, among other applications.² Without the full valuation, event studies provide only limited information for cost-benefit analyses. Further, event studies have become a tool used by courts to estimate damages from insider trading and other illegal activities, but these damages are underestimated when events were partially anticipated (Cornell and Morgan, 1990).

Second, many studies are interested in whether events have small or large effects. For instance, researchers are interested in whether minimum wage policies (Card and Krueger, 1997, Chapter 10), layoffs (Hallock, 1998), and shareholder initiatives (Karpoff et al., 1996) meaningfully affect corporate profits. Each of these studies finds small effects. However, as these researchers recognize, these results cannot distinguish whether the true effects are indeed small or the events were simply well-anticipated.³ If our methods indicate that events were in fact surprising, then researchers may have greater confidence that the true effects are indeed small, but if our methods indicate that events were in fact well-anticipated, then researchers may be more hesitant to draw this conclusion.

Third, our methods can improve measures of policy uncertainty. For instance, Bianconi et al. (2019) proxy for trade policy uncertainty with the difference between the tariffs that would hold if the U.S. Congress did or did not grant Most Favoured Nation status to China at a given time. However, it is plausible that the probability of Congressional action is correlated with tariff level changes that would result from that action. In this case, they would mismeasure trade policy uncertainty. More broadly, some events affect firms in different ways as, for instance, when some firms will end up above a regulatory cutoff and other firms will end up below it (e.g., Meng, 2017) or as when different firms depend on different facets of a court's ruling. In such cases, the probability of the realized event will be firm-specific. Our methods allow researchers to estimate these firm-specific probabilities, which could then be used to test for the effects of policy uncertainty on decisions such as hiring and investment.

Fourth, some studies aggregate event effects in order to investigate the sign of an overall effect. For instance, Kogan et al. (2017) aggregate market reactions to different patent grants in order to test whether the net effect of innovation tends to be positive or negative. They use the overall frequency of successful patent applications

²These probabilities are also critical for using empirical results to test theory. In a corporate finance setting, Hennessy and Strebulaev (2020) show that the bias from testing theory-implied causal effects from observed short-run responses to shocks depends on the probability of the shocks.

³Cutler et al. (1989) show that even big news events tend to move the stock market by a relatively small amount, a result consistent with partial anticipation.

as a proxy for the priced-in probability of each patent being granted, but this probability is in fact likely to vary both across firms and over time within firms. Because the net effect of innovation likely also varies across firms, adjusting for patent-specific probabilities of success could disproportionately affect the value of patents with particular types of effects on stock prices and thus could plausibly change the sign of the aggregate effect.

Fifth, many researchers are interested in explaining variation in event effects. For instance, Farber and Hallock (2009) find that stock price reactions to job cut announcements have become less negative over time, and Bronars and Deere (1990) find that the effects of union elections have declined over time. However, if layoff announcements and union elections became better anticipated over time, then these findings may not have the economic significance attributed to them. Bronars and Deere (1990) also explore which types of firms are most affected by union elections. However, the economic interpretation of these results is sensitive to the possibility that anticipation of union elections varies with firm characteristics. Researchers have long noted that cross-sectional analyses could be severely biased—even to the point of estimating the wrong sign—when the market can forecast events based on the observable characteristics of interest (e.g., Lanen and Thompson, 1988; MacKinlay, 1997; Bhagat and Romano, 2002). Our methods allow future studies to control for the priced-in probability of the event.

Sixth, many researchers use 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. For instance, pre-election race ratings correctly call most elections that end up being close. When a stronger candidate or party can exert its influence on the margin, the few votes of separation will be more likely to favor that side. Our methods allow researchers to identify the elections that market participants viewed as effectively random.⁴ The outcomes of these elections can then be used as the randomized experiments that have been sought in elections that were close *ex post*.

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.⁵ If researchers show that some events are truly random, then these events’ risk-neutral probabilities tell us whether investors

⁴This 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.

⁵Risk-neutral probabilities are the probabilities needed to correct event study estimates.

expected consumption to be higher in the realized state or in the other possible state. For 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.

Previous approaches to correcting for partial anticipation

Much previous work has highlighted event studies' inability to correctly measure the full effect of an event when it is not a complete surprise, and researchers commonly acknowledge the problems posed by partial anticipation in their event study applications.⁶ The standard solution is to select events that the researcher judges to be relatively surprising. For instance, instead of investigating how stock prices change on the day that the minimum wage increases, Card and Krueger (1997, Chapter 10) use events in the policymaking process that are likely to contain more new information. However, such events are themselves rarely complete surprises, as these authors and others commonly acknowledge.⁷

Several researchers attempt to further reduce the effects of partial anticipation by extending the event window to include earlier time periods, hoping that the extended event window captures any news leaks that may have occurred prior to the documented event (e.g., Jayachandran, 2006; Auerbach and Hassett, 2007; Linn, 2010; Lee and Mas, 2012; Al-Ississ and Miller, 2013).⁸ In some cases, the event

⁶Some of the earliest event studies already recognize that partial anticipation can strongly attenuate estimated effects (e.g., Ball, 1972). Malatesta and Thompson (1985) distinguish the economic impact of an event and the announcement effect of an event. MacKinlay (1997) and Lamdin (2001), among others, emphasize that the problem of partially anticipated events may be especially severe in studies that seek to analyze the impact of regulations. Binder (1985) shows that the event study methodology has very little power to detect the effects of twenty major regulations because many of these regulations were partially (or even fully) anticipated.

⁷Dube et al. (2011) find that even top-secret coup authorizations leak to the markets. The coups themselves are then well-anticipated. 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. Some researchers do claim that their events were complete surprises (e.g., Bell and Machin, 2017). Our methods allow them to test such claims.

⁸Bernanke and Kuttner (2005) instead attempt to clean a continuous event (the choice of Federal funds rate) of its unsurprising components that are reflected in futures prices. This approach is not useful in most event study applications, as events such as elections or policy announcements have discrete outcomes and may lack the analogue of a futures contract on the outcome (but see below

window is years-long. However, the event study design requires that event-window effects be attributable to the event of interest, not to other news. This identification requirement becomes more demanding as the event window is extended. Further, the signal-to-noise ratio of event studies falls as the event window is extended, reducing 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, other researchers seek to recover that probability directly. The most widely applied way of recovering this probability is to use prediction market contracts as indicators of the event's prior probability (e.g., Roberts, 1990; Herron, 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).⁹ Prediction markets can be a valuable source of information when the proper contracts exist, but this method faces a significant hurdle in many applications: prediction market contracts are unavailable for many events of interest and can be quite thinly traded even when they are available. We develop methods that instead require the existence of liquid options markets for firms affected by the event. Options markets have been around longer and are more thickly traded than many prediction markets. Further, prediction markets' prices can fluctuate rapidly in response to real-time news, as in some of our empirical applications below. In these cases, it is not clear which moment's price corresponds to the probabilities priced-in by stock markets at the time they close. In contrast, the probabilities recovered from options markets should be the same probabilities priced-in by closing stock prices. Finally, if a different set of people, with different information, are trading in prediction markets, then these may not be the correct probabilities to use to adjust financial market results. Using options trades to adjust stock market estimates reduces this concern somewhat.

Several papers in the finance literature do attempt to infer the priced-in probability of events from the prices of financial options.¹⁰ These papers assume that options

regarding prediction markets).

⁹Working in a context without prediction markets, Fisman (2001) asks investment bankers how much the broader Indonesian stock market would have fallen if Suharto had died suddenly. He backs out the implied probability of Suharto's death from the change in the stock market that actually occurred upon Suharto's death.

¹⁰Some other finance literature is more loosely related. Several papers show that option prices anticipate the release of earnings announcements (e.g., Patell and Wolfson, 1979, 1981; Dubinsky et al., 2019) and of macroeconomic policy news (e.g., Ederington and Lee, 1996; Lee and Ryu, 2019), with implied volatilities falling upon the news being released. We show how to use the change in option prices to back out the probability of a policy event. Kelly et al. (2016) use

are priced according to specific parametric models and search for the event probability that reconciles observed option prices and theoretical option prices. In particular, Gemmill (1992) assumes that options are priced according to the Black (1976) model of lognormal futures prices, Barraclough et al. (2013) seem to assume that options are priced according to the Black and Scholes (1973) model of lognormal stock prices, Borochin and Golec (2016) assume that options are priced according to the Cox et al. (1979) binomial model, and Carvalho and Guimaraes (2018) assume that options are priced according to the Heston (1993) stochastic volatility model.¹¹ Each of these pricing models imposes specific parametric assumptions on the distribution of stock prices. However, these assumptions are likely to be violated in practice (and indeed conflict with each other). Well-known discrepancies between actual option prices and theoretical option prices generate “anomalies” such as implied volatility smiles and smirks. Such discrepancies will generally bias the estimated event probability: 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.¹²

Our new options-based methods do not impose parametric assumptions. Our

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) proposes a latent information model that extracts event probabilities from stock price movements. This model is appropriate only when the (temporary) lack of an event contains information, as is true for endogenous events such as corporate announcements. We instead focus on policy events that affect a cross-section of firms and are not endogenous to any one firm. Finally, Fisman and Zitzewitz (2019) and van Tassel (2016) use stock and option prices, respectively, to recover investors’ post-event beliefs. Here, we recover investors’ pre-event beliefs.

¹¹As an alternative to these approaches, one could imagine directly estimating the entire pre-event implied probability density function for stock prices following Breeden and Litzenberger (1978) and comparing the density at each peak in the distribution. However, this method relies on the second derivative of option prices, which is sensitive to small variations in option prices. Further, backing out a probability from such a distribution would still require assumptions about the conditional distributions being mixed together.

¹²Carvalho and Guimaraes (2018) highlight the potential biases that would be introduced by using the Black and Scholes (1973) pricing model in a market where prices do not exactly match the model’s predicted prices. They emphasize that the estimated event probability is identified primarily by variation in the prices of the same out-of-the-money options that tend to show discrepancies with respect to theoretical option pricing models. Concerned about this bias, Barraclough et al. (2013) develop a weighting scheme that relies less on the out-of-the-money options that tend to demonstrate implied volatility anomalies under the Black and Scholes (1973) model. We will see that these out-of-the-money options are in fact the ones that carry the most direct information about the priced-in probability of events.

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 only if the event occurs or (ii) that the expected post-event volatility does not depend on the event's realization. The latter is one of the assumptions imposed in Gemmill (1992) and Carvalho and Guimaraes (2018), so our second method can be viewed as a direct generalization of prior methods.¹³

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 the 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 proofs and extensions to the theoretical analysis.

2 Setting

We study the beliefs of a representative market investor. Let each state of the world at discrete time t be indexed (ω_t, k) , with a continuous component $\omega_t \in \mathcal{R}^N$ and a discrete component $k \in \{L, H\}$. $S(\omega_t, k)$ is the price of a firm's stock in state (ω_t, k) . We henceforth write S_t for the observed stock price and write S_t^L for $S(\omega_t, L)$ and S_t^H for $S(\omega_t, H)$.

At time τ , an event happens that reveals the state to be either H or L .¹⁴ For $t \geq \tau$, the researcher observes either $S_t = S_t^H$ or $S_t = S_t^L$, depending on the outcome of the event. Time t agents know ω_t , but prior to time τ , agents do not know whether $k = H$ or $k = L$.¹⁵ Let the time t representative agent assign risk-neutral probability p_t^H to $k = H$ and risk-neutral probability p_t^L to state $k = L$.¹⁶ For instance, consider

¹³In concurrent work, Grinblatt and Wan (2020) argue that one can in principle back out risk-neutral probabilities from the prices of options traded before the 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 states or the form of the event's effect.

¹⁴The event need not be binary. If, for instance, outcome H is realized, then we can aggregate all of the other possible outcomes into a single indicator L .

¹⁵Until Section 3.2, we do not specify whether agents know in advance that this information will be revealed at time τ .

¹⁶Absence 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

a presidential election between candidates H and L . Let $k = H$ correspond to the state in which next year's president is H and $k = L$ correspond to the state in which next year's president is L . The election outcome is revealed just before time τ . Prior to τ , agents do not know which candidate will win, assigning probabilities p_t^H and $p_t^L = 1 - p_t^H$ to each outcome. From τ onward, agents assign $p_t^H = 1$ if H won and assign $p_t^L = 1$ if L won.

2.1 The Bias in the Standard Event Study Methodology

Without loss of generality, assume that event H occurs at time τ . For ease of exposition, consider a case with only a single firm. An event study aims to recover $S_{\tau-1}^H$, the stock price just before the event if the event outcome were already known. Its identifying assumption is that the controls account for all elements of ω_τ that differ from $\omega_{\tau-1}$.¹⁷ In this case, the time τ return predicted from the controls captures the effects of changing $\omega_{\tau-1}$ to ω_τ and the excess time τ return relative to the predicted return (as captured by an event-day dummy) reflects the new information about k .

Researchers use event studies to calculate the event effect as $S_{\tau-1}^H - S_{\tau-1}$, but they would like to calculate $S_{\tau-1}^H - S_{\tau-1}^L$. By absence of arbitrage, the time $\tau - 1$ stock price must be:

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

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

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

The estimated event effect $S_{\tau-1}^H - S_{\tau-1}$ is less than the full event effect $S_{\tau-1}^H - S_{\tau-1}^L$. As $p_{\tau-1}^H \rightarrow 0$, researchers recover $S_{\tau-1}^H - S_{\tau-1}^L$ from $S_{\tau-1}^H - S_{\tau-1}$: outcome H was judged at time $\tau - 1$ to be extremely unlikely (or even impossible), so when outcome H nonetheless occurs, an event study provides the entire effect of outcome H relative to outcome L . For this reason, researchers have sought events that are surprises. However, for $p_{\tau-1}^H > 0$, an event study underestimates $S_{\tau-1}^H - S_{\tau-1}^L$ because $S_{\tau-1}$ already reflects the possibility of outcome H . This is the well-known problem of partially anticipated events. Moreover, as $p_{\tau-1}^H$ goes to 1, an event study measures an arbitrarily small fraction of the true event effect $S_{\tau-1}^H - S_{\tau-1}^L$. Event studies

risk into the probability weights by adjusting the “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).

¹⁷See Campbell et al. (1997, Chapter 4), MacKinlay (1997), and Kothari and Warner (2007), among others, for reviews of event study methods. The identifying assumption is weaker in event studies that have multiple firms.

can provide only a lower bound for the implications of an event in the absence of information about $p_{\tau-1}^H$.

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

We now describe two new approaches to recovering $p_{\tau-1}^H$ from running event studies in options prices. The first uses changes in out-of-the-money option prices and the second uses changes in a stock's expected variance.

3.1 Using the Change in Tail Probabilities

We begin with an intuitive discussion of how individual option prices can be used to recover the priced-in event probability before proceeding to the formal derivation. A call (put) option on stock S confers the right—but not the obligation—to buy (sell) the stock S at a defined “strike” price K on a defined expiration date T . Consider the pricing of a European-style call option around an event. The option's value derives from the chance that the underlying stock 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 allows the option to expire unexecuted. Thus, the date x value of the call option, $C_{x,T}(S_x, K)$, is the expected gap between the stock price and the strike conditional on the stock price being above the strike.

Imagine that event H is realized and that it increases the price of a firm's stock. Before the event's outcome is known, the option's value is a weighted average of the value conditional on the event occurring and the value conditional on the counterfactual outcome:

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

The weight on the actual outcome is the (risk-neutral) probability of the event. If the event's effect on the stock is large and the call option has a high strike, then the value of the call option if the counterfactual event outcome had occurred would have been small. In that case, we have:

$$C_{\tau-1,T}(S_{\tau-1}, K) \approx p_{\tau-1}^H C_{\tau-1,T}^H(S_{\tau-1}^H, K).$$

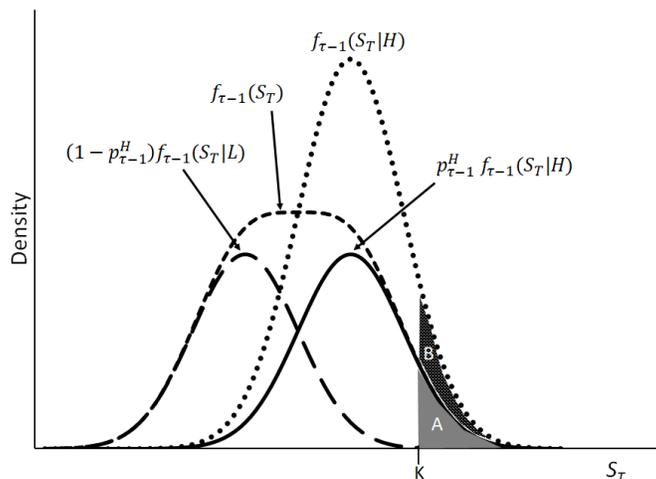


Figure 1: Illustration of how changes in option prices identify the prior probability $p_{\tau-1}^H$ when event H is realized at time τ .

Just after the event occurs, the value of the option is $C_{\tau,T}^H(S_\tau^H, K)$, which approximates $C_{\tau-1,T}^H(S_{\tau-1}^H, K)$ under the event study identification assumption that the only news at τ is the event realization.¹⁸ The ratio of the option price observed just before the event to the option price observed just after the event recovers the risk-neutral probability of the event.

Figure 1 presents a graphical version of this intuition. Conditional on the information available at time $\tau - 1$ (just before the event occurs), the risk-neutral distribution of prices for the stock at the expiration date T is $f_{\tau-1}(S_T)$, which is a mixture of the distribution conditional on event H occurring, $f_{\tau-1}(S_T|H)$, and the distribution conditional on the counterfactual event occurring, $f_{\tau-1}(S_T|L)$. The value of a call option with a strike K is the discounted expected value of $S_T - K$ conditional on the stock price being above K , in the area labeled A. Once the event occurs, the density of stock prices at the expiration date becomes $f_{\tau-1}(S_T|H)$ and the option's value integrates over the larger area $A + B$. If $f_{\tau-1}(S_T|L)$ contains very little mass above K , the jump in the option price between date $\tau - 1$ and τ identifies the extent to which distribution $f_{\tau-1}(S_T|H)$ was downweighted by the event probability $p_{\tau-1}^H$.

Formally, the value of the call option at date x that expires at T and has a strike

¹⁸In practice, we require the less restrictive event study identification assumption that allows controls to absorb non-event news.

of K is:¹⁹

$$C_{x,T}(S_x, K) = \frac{1}{R_{x,T}} \int_K^\infty (S_T - K) f_x(S_T|S_x) dS_T,$$

where we now explicitly condition the risk-neutral distribution of S_T on the time x stock price. $R_{x,T} \geq 1$ is the gross risk-free rate from time x to T . At time $\tau - 1$, the price of a call option with strike K and expiration $T > \tau - 1$ must satisfy:

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

where $f_{\tau-1}(\cdot|L)$ and $f_{\tau-1}(\cdot|H)$ condition on the realization of k .

Now imagine that at time τ the event reveals that H is the true state of the world. Consider how this event changes the option's price:

$$\begin{aligned} &C_{\tau-1,T}^H(S_{\tau-1}^H, K) - C_{\tau-1,T}(S_{\tau-1}, K) \\ &= \frac{1}{R_{\tau-1,T}} (1 - p_{\tau-1}^H) \int_K^\infty (S_T - K) [f_{\tau-1}(S_T|S_{\tau-1}^H, H) - f_{\tau-1}(S_T|S_{\tau-1}^L, L)] dS_T \\ &= (1 - p_{\tau-1}^H) C_{\tau-1,T}^H(S_{\tau-1}^H, K) - (1 - p_{\tau-1}^H) C_{\tau-1,T}^L(S_{\tau-1}^L, K). \end{aligned} \quad (3)$$

A standard arbitrage bound (e.g., Cochrane, 2005) requires $C_{\tau-1,T}^L(S_{\tau-1}^L, K) \geq 0$. Using this inequality in equation (3) implies the following estimator of the event probability, labeled \bar{p} :²⁰

$$p_{\tau-1}^H \leq \frac{C_{\tau-1,T}(S_{\tau-1}, K)}{C_{\tau-1,T}^H(S_{\tau-1}^H, K)} \triangleq \bar{p}. \quad (4)$$

The $p_{\tau-1}^H$ in inequality (4) is the same $p_{\tau-1}^H$ as in equation (2). Thus, we can use the observed changes in option prices to bound the risk-neutral probability of the

¹⁹We 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. The appendix 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 will focus on options with the shortest time to maturity.

²⁰The other two arbitrage bounds ($C_{\tau-1,T}^L(S_{\tau-1}^L, K) \geq S_{\tau-1}^L - K/R_{\tau-1,T}$ and $C_{\tau-1,T}^L(S_{\tau-1}^L, K) \leq S_{\tau-1}^L$) also imply upper bounds on $p_{\tau-1}^H$ when combined with equation (3). However, it is easy to show that these two upper bounds are each greater than 1 and thus are not relevant bounds.

realized outcome H and, from equation (2), thereby also bound the bias in the event study measure.

Now assume that there exists some \bar{S} such that $S(\omega_t, H) > \bar{S}$ implies $S(\omega_t, H) > S(\omega_t, L)$.²¹ And assume that $f_{\tau-1}(S_T|S_{\tau-1}^L, L)$ goes to zero as S_T becomes large. These two assumptions together imply that increasing K brings $C^L(S_{\tau-1}^L, K)$ to zero faster than it brings $C^H(S_{\tau-1}^H, K)$ to zero. The arbitrage bound $C_{\tau-1,T}^L(S_{\tau-1}^L, K) \geq 0$ then holds exactly as K becomes large, which means that \bar{p} converges to $p_{\tau-1}^H$ for some sufficiently large K . Because $C^H(S_{\tau-1}^H, K)$ and $C^L(S_{\tau-1}^L, K)$ are likely to be more distinct when the event has a larger effect on the stock price, the bound \bar{p} is likely to be tight for a broader set of strikes when the event moves the stock price by a large amount.

Figure 1 depicts just such a large K . For smaller K , the long-dashed distribution $f_{\tau-1}(S_T|L)$ may have nontrivial mass in region A. In this case, the change in the price of the option reflects both the rescaling by $p_{\tau-1}^H$ and the loss of this unobserved probability mass. The possibility of unobserved probability mass explains why equation (4) is an inequality rather than an equality.

The bound \bar{p} may be especially tight when $p_{\tau-1}^H$ is large, which is precisely the case in which standard event studies suffer arbitrarily large biases and therefore is the case in which a tight bound is most needed. From equation (3), the bias is

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

As $p_{\tau-1}^H$ grows, the bias vanishes, 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 mass in the shaded region.

Put options can also recover the event probability, which is useful when an event reduces the price of the underlying asset. Let event L now be the one realized at time τ . The time x price of a put option with expiration T and strike K is

$$P_{x,T}(S_x, K) = \frac{1}{R_{x,T}} \int_{-\infty}^K (K - S_T) f_x(S_T|S_x) dS_T.$$

²¹This 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 $S(\omega_t, H) > \sum_{i=1}^N [p_{\tau-1}^{L_i}/(1 - p_{\tau-1}^H)] S(\omega_t, L_i)$. The appendix relaxes the assumption that the realized outcome is extreme.

Using the arbitrage bound $P_{\tau-1,T}^H(S_{\tau-1}^H, K) \geq 0$, an analogous derivation to the foregoing yields:

$$p_{\tau-1}^L \leq \frac{P_{\tau-1,T}(S_{\tau-1}, K)}{P_{\tau-1,T}^L(S_{\tau-1}^L, K)}. \quad (6)$$

If there exists \underline{S} such that $S(\omega_t, L) < \underline{S}$ implies $S(\omega_t, L) < S(\omega_t, H)$ and if, in addition, $f_{\tau-1}(S_T | S_{\tau-1}^H, H) \rightarrow 0$ as S_T becomes small, then the arbitrage bound $P_{\tau-1,T}^H(S_{\tau-1}^H, K) \geq 0$ holds exactly as K becomes small. In that case, the right-hand side of inequality (6) converges to $p_{\tau-1}^L$ for some sufficiently small K .

However, much work has conjectured that out-of-the-money put options carry a premium because they offer protection against low-probability crashes.²² The possibility of such disasters is plausibly independent of the event outcome k . In that case, deep out-of-the-money put options may retain much of their value even if $k = H$. The right-hand side of inequality (6) then only loosely bounds $p_{\tau-1}^L$ and does not converge to $p_{\tau-1}^L$ as K becomes small. Further, if the value of the deepest out-of-the-money put options is primarily driven by disaster risk that is independent of k , then the bias from estimating $p_{\tau-1}^L$ via inequality (6) may actually increase as the strike price falls. It is no longer clear which strikes should provide the tightest bound. As a result, our proposed method may be most effective when the realized event increases a firm's stock price. In that case, researchers can estimate the bound \bar{p} from call options, for which the bias will often decrease monotonically in the observed strike prices and even vanish for sufficiently large strike prices.

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. That method placed an upper bound on $p_{\tau-1}^H$ and did not require advance knowledge of the date that the event would happen. We now derive a second method of estimating the priced-in probability, using changes in the expected variance of stock prices. This method does require advance knowledge of the date that the event will happen, but it

²²Since 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).

identifies the event probability under a different set of circumstances, now requiring that the post-event expected variance of the stock price not depend on the event's realization.

We again build intuition before formally deriving the approach. Figure 2 depicts the time $\tau - 1$ risk-neutral distribution of time τ stock prices decomposed into a compound lottery. Assume that market actors know that the event will occur at time τ . The first lottery is over the event outcome k and the second lottery is over the non-event news ω_τ . The variance of the first lottery is

$$\begin{aligned} \text{Var}_{\tau-1}[S_\tau|\omega_\tau] &= p_{\tau-1}^H (S_{\tau-1}^H - S_{\tau-1})^2 + (1 - p_{\tau-1}^H) (S_{\tau-1}^L - S_{\tau-1})^2 \\ &= \frac{p_{\tau-1}^H}{1 - p_{\tau-1}^H} (S_{\tau-1}^H - S_{\tau-1})^2, \end{aligned} \quad (7)$$

where the second line substitutes for $S_{\tau-1}^L$ from equation (1) because we again assume, without loss of generality, that event H is realized. The variance depends on both the probability of the event and the magnitude of the event effect. Rearranging, we find:

$$p_{\tau-1}^H = \frac{\text{Var}_{\tau-1}[S_\tau|\omega_\tau]}{\text{Var}_{\tau-1}[S_\tau|\omega_\tau] + (S_{\tau-1}^H - S_{\tau-1})^2}. \quad (8)$$

If we can estimate the variance of this first lottery, then we can infer $p_{\tau-1}^H$. The challenge is to estimate the variance of this first lottery.

Now consider the variance of the full compound lottery. Temporarily fixing $R_{\tau-1,\tau} = 1$, the appendix shows that

$$\text{Var}_{\tau-1}[S_\tau] = \text{Var}_{\tau-1}[S_\tau|\omega_\tau] + p_{\tau-1}^H \text{Var}_{\tau-1}[S_\tau^H] + (1 - p_{\tau-1}^H) \text{Var}_{\tau-1}[S_\tau^L]. \quad (9)$$

The first term on the right-hand side of equation (9) is the variance of the first, event lottery. The second and third terms capture the variance induced by the second lottery. Now imagine that a portfolio of options replicates $\text{Var}_{\tau-1}[S_\tau]$. In that case, we can construct a similar replicating portfolio for $\text{Var}_{\tau-1}[S_\tau^H]$ by applying standard event study techniques to each option in the replicating portfolio. Subtracting $\text{Var}_{\tau-1}[S_\tau^H]$ from each side of equation (9), substituting from equation (7), and rearranging, we have:

$$\frac{p_{\tau-1}^H}{1 - p_{\tau-1}^H} = \frac{\text{Var}_{\tau-1}[S_\tau] - \text{Var}_{\tau-1}[S_\tau^H]}{(S_{\tau-1}^H - S_{\tau-1})^2} + (1 - p_{\tau-1}^H) \frac{\overbrace{(\text{Var}_{\tau-1}[S_\tau^H] - \text{Var}_{\tau-1}[S_\tau^L])}^{\Delta \text{Var}}}{(S_{\tau-1}^H - S_{\tau-1})^2}. \quad (10)$$

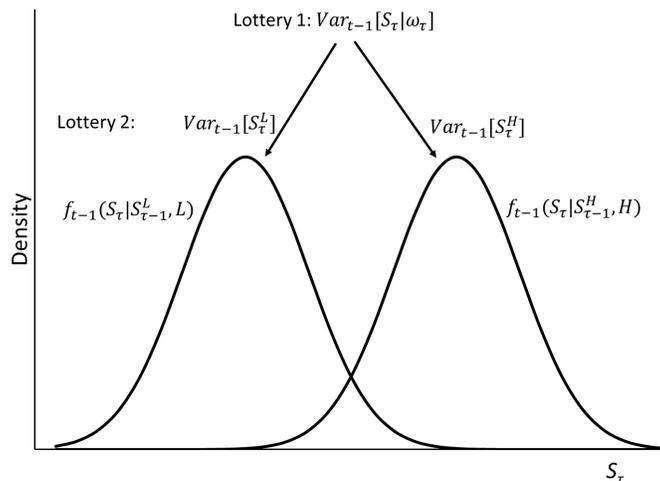


Figure 2: The variance of the time τ (post-event) stock price accounts for uncertainty about the realization of k and for the variance of the stock price conditional on each k .

Both $Var_{\tau-1}[S_\tau^L]$ and $p_{\tau-1}^H$ are unobserved. However, if the event's realization does not affect the variance of the second lottery, then $\Delta Var = 0$ and the second term vanishes. (In fact, several previous model-based approaches implicitly impose $\Delta Var = 0$ among their other assumptions (e.g., Gemmill, 1992; Carvalho and Guimaraes, 2018).) Rearranging, we then have:

$$p_{\tau-1}^H = \frac{Var_{\tau-1}[S_\tau] - Var_{\tau-1}[S_\tau^H]}{Var_{\tau-1}[S_\tau] - Var_{\tau-1}[S_\tau^H] + (S_{\tau-1}^H - S_{\tau-1})^2}. \quad (11)$$

We can estimate $p_{\tau-1}^H$ from the replicating portfolios for the variance from $\tau - 1$ to τ .²³ The difference $Var_{\tau-1}[S_\tau] - Var_{\tau-1}[S_\tau^H]$ cleans the variance of the compound lottery of the variance induced by ω_τ , leaving us with the variance of the first lottery in Figure 2. And from equation (8), that variance and the realized jump in stock prices together imply $p_{\tau-1}^H$.

²³Viewing $p_{\tau-1}^H$ as implicitly defined as a function of ΔVar , we have, using equation (2):

$$\left. \frac{dp_{\tau-1}^H}{d\Delta Var} \right|_{\Delta Var=0} = \frac{(1 - p_{\tau-1}^H)^3}{(S_{\tau-1}^H - S_{\tau-1})^2} = \frac{1 - p_{\tau-1}^H}{(S_{\tau-1}^H - S_{\tau-1}^L)^2} \geq 0.$$

If we estimate $p_{\tau-1}^H$ under the assumption that $\Delta Var = 0$, then the bias from small deviations in ΔVar is small when $p_{\tau-1}^H$ is large. We therefore again have an especially precise estimate in the case where, from equation (2), the full event effect is most sensitive to $p_{\tau-1}^H$.

Thus far, we have seen how we might estimate $p_{\tau-1}^H$ if we could construct a replicating portfolio for $Var_{\tau-1}[S_\tau]$. We are, in effect, seeking the single-day variance swap rate when the underlying asset's price can jump discretely.²⁴ However, the desired variance swap rate will rarely be directly observed in the market. Martin (2017) provides a critical result. He constructs the replicating portfolio for a related object, a “simple variance swap”. 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 .²⁵ Martin (2017) prices the simple variance swap under the assumptions of a constant interest rate, a constant dividend rate, and small timesteps, without assuming away the possibility of jumps.²⁶ Martin (2017) shows that

$$V_{\tau-1,T} = \frac{2R_{\tau-1,T}}{[\tilde{R}_{\tau-1,T} S_{\tau-1}]^2} \left\{ \int_0^{\tilde{R}_{\tau-1,T} S_{\tau-1}} P_{\tau-1,T}(S_{\tau-1}, K) dK + \int_{\tilde{R}_{\tau-1,T} S_{\tau-1}}^\infty C_{\tau-1,T}(S_{\tau-1}, K) dK \right\}. \quad (12)$$

We will also be interested in the variance strike $V_{\tau-1}^H$, which assumes that $k = H$ is known from time $\tau - 1$:

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

²⁴The 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 .

²⁵Note that $\tilde{R}_{\tau-1,y} S_{\tau-1}$ is the time $\tau - 1$ forward price of S_y .

²⁶The 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. We must here allow for the possibility of jumps. Jiang and Tian (2005) and Carr and Wu (2009) synthesize variance swaps in the presence of 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 and will use options with the shortest maturities.

Again using the results in Martin (2017), we have:

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

The following proposition relates $p_{\tau-1}^H$ to $V_{\tau-1,T}$ and $V_{\tau-1,T}^H$:

Proposition 1. *Define*

$$\tilde{V} \triangleq (S_{\tau-1})^2 V_{\tau-1,T} - (S_{\tau-1}^H)^2 V_{\tau-1,T}^H, \quad \tilde{p} \triangleq \frac{\tilde{V}}{\tilde{V} + [2\tilde{R}_{\tau-1,\tau} - 1] [S_{\tau-1}^H - S_{\tau-1}]^2}. \quad (13)$$

Then:

1. $p_{\tau-1}^H \rightarrow \tilde{p}$ 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 $p_{\tau-1}^H$ that can be thought of as a formal version of equation (11). The first result establishes that \tilde{p} becomes an arbitrarily good approximation to $p_{\tau-1}^H$ as $[S_{\tau-1}^H]^2 V_{\tau-1,T}^H - [S_{\tau-1}^L]^2 V_{\tau-1,T}^L \rightarrow 0$ or as $p_{\tau-1}^H \rightarrow 1$, and the second result establishes that \tilde{p} is then a lower (upper) bound on $p_{\tau-1}^H$ if the post-event variance is smaller (larger) following $k = L$ than following $k = H$. The intuition for the result tracks that already given for equation (10), 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 $\Delta Var \rightarrow 0$. 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, where \tilde{V} is a metric readily constructed from observed options prices and event study estimates. If $\tilde{V} > 0$, then the variance of the compound lottery is large relative to the variance of the lottery conditional on $k = H$. This is the standard case, which we implicitly assumed in discussing equation (10). In contrast, if $\tilde{V} < 0$, then the variance conditional on $k = H$ is at least as great as the variance of the compound lottery. From equation (9), the variance conditional on $k = H$ must therefore be substantially greater than the variance conditional on $k = L$. It is easy to see from the definition of \tilde{p} that it is an uninformative bound on $p_{\tau-1}^H$ in that case.

3.3 Comparing the Two Estimators

We have derived two estimators of the risk-neutral probability of an event. Both estimators are model-free, in contrast to prior literature (described in the introduction) that recovers event probabilities from option prices by assuming that stock prices evolve according to specific parametric processes. The first estimator (\bar{p}) is identified by the tail of the stock price distribution, and the second estimator (\tilde{p}) is identified by the expected variance of the stock price process. The first estimator requires that some options that are valuable when the realized event happens would have been nearly worthless if other events had happened, and the second estimator requires that the expected variance of the post-event stock price process 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. In particular, the integral approximation becomes poorer when the strike prices of the liquidly traded options 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. As a first example, \bar{p} performs best when the realized event is extreme. The appendix shows that the bound obtained from \bar{p} cannot become arbitrarily tight for "middle" events. In contrast, \tilde{p} does not depend on the realized event being extreme. As a second example, we described how \bar{p} may only loosely bound the probability of events that reduce the price of a stock because the prices of out-of-the-money put options may reflect disaster risks whose consequences are independent of the event realization. Because \tilde{p} is not solely identified by the tail of the stock price distribution, it is not as sensitive to this common chance of extreme stock price outcomes. We may therefore better estimate $p_{\tau-1}^H$ from \tilde{p} when an event reduces firms' value.

4 Empirical Approach

We now describe our empirical approach to estimating \bar{p} and \tilde{p} from observed option prices across a set of firms.²⁷ 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 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).²⁸ For a similar reason, we drop firms with a quarterly dividend yield greater than 2% over the estimation window (following 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). We therefore limit the sample to firms that do not have an earnings announcement in a 3-day window around the event and control for earnings announcements that occur elsewhere in the estimation window. We do not control for either the market index or its implied volatility because we analyze big events that may have affected that index. Controlling for the index could accidentally absorb the desired event effect. 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).

We obtain stock prices, quarterly dividends, and earnings dates from Compustat. We obtain 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.

We next describe additional, theoretically motivated restrictions designed to recover tight bounds on the priced-in event probability.

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 the elemental events differs by firm.

²⁸Options 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).

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

The objective is to estimate $p_{\tau-1}^H$ by obtaining a tight bound \bar{p} , where H again stands for the realized event. From equation (5), the bias $\bar{p} - p_{\tau-1}^H$ depends on $C_{\tau-1,T}^L(S_{\tau-1}^L, K)$. If we could identify options for which $C_{\tau-1,T}^L(S_{\tau-1}^L, K)$ were small, then we could empirically estimate the priced-in probability of an event from the following regression:²⁹

$$\ln \left(\frac{C_{iK(t-1)}}{C_{iKt}} \right) = \alpha_{iK} + \beta Event_t + \theta_{iK} X_{it} + \varepsilon_{iKt}, \quad (14)$$

where we change notation on the call option price, letting i index firms, K index strike prices, and t index trading dates. An analogous regression holds when we examine puts. $Event_t$ is a dummy variable for the event occurring on trading date t . X_{it} is a vector of controls, which includes dummies for the days before and after the event and dummies for a three-day window around an earnings announcement. In order to favor more liquid observations, we weight by the inverse of the relative bid-ask spread averaged over days t and $t - 1$. We assign a weight of zero if either day has a bid of zero. We use any trading days that are within 100 days before the event and 7 days before the option's expiration date and calculate standard errors that are robust to clustering by firm and by date.

We estimate \bar{p} by predicting $C_{iK(\tau-1)}/\hat{C}_{iK(\tau-1)}^H$. We predict $C_{iK(\tau-1)}/\hat{C}_{iK(\tau-1)}^H$ from $\hat{\beta}$ alone, 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. We thus have $\bar{p} = \exp(\hat{\beta})$. We do not let β vary across firms because the event probability should not vary across firms in our applications (see footnote 27).

Thus far we have assumed that the empirical researcher can identify those options for which $C_{\tau-1,T}^L(S_{\tau-1}^L, K)$ is small, but $C_{\tau-1,T}^L(S_{\tau-1}^L, K)$ is unobservable. We now describe how empirical researchers can estimate a tight bound on the event probability without knowledge of $C_{\tau-1,T}^L(S_{\tau-1}^L, K)$. To achieve this, we use theoretically motivated insights about how the bias from ignoring $C_{\tau-1,T}^L(S_{\tau-1}^L, K)$ 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. 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

²⁹We find that log-changes in option prices are approximately normally distributed.

disasters. Our preferred specifications therefore limit the sample to the deepest out-of-the-money liquid option for each firm.³⁰

Second, 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 given probability $p_{\tau-1}^H$, $S_{\tau-1}^L$ must be relatively small if $S_{\tau-1}^H - S_{\tau-1}$ is large, in which case $f_{\tau-1}(S_T|L)$ in Figure 1 may have little mass above K . We assess firms' sensitivity to the event via traditional event studies:

$$\ln(S_{it}/S_{i(t-1)}) = \gamma_{i1} + \gamma_{i2}Event_t + \gamma_{i3}X_{it} + \varepsilon_{it}, \quad (15)$$

where S_{it} is the closing stock price for firm i on trading date t , X_{it} is as before, and standard errors are robust to clustering by firm and by date. We use a 200-day estimation window (180 days before and 20 days after the event). Firms with statistically large stock price responses to the event will exhibit less bias, but there are fewer of them, resulting the typical bias-variance tradeoff. Our preferred specifications will limit the sample to firms with sufficiently high t-statistics on $\hat{\gamma}_{i2}$ that the event estimate does not move outside the confidence bounds for higher t-statistic cutoffs.

Finally, our preferred specifications use only call options because, as described in Section 3.1, put options may generate more bias than call options. Regression (15) tells us whether to use call or put options in regression (14): we should use call options for those firms with $\hat{\gamma}_{i2} > 0$ and put options for those firms with $\hat{\gamma}_{i2} < 0$. Our preferred specifications limit the sample to firms with $\hat{\gamma}_{i2} > 0$.

4.2 Estimating \tilde{p} from Synthesized Variance Swaps

Our second approach to estimating the priced-in probability of the uncertain event uses options at the full distribution of strikes for each firm. Let the probability inferred from firm i be \tilde{p}_i . Equation (13) shows that calculating \tilde{p}_i requires the underlying asset price on the day before the event ($S_{i(\tau-1)}$) and the counterfactual value of that asset if the event outcome were already known ($S_{i(\tau-1)}^H$). The former is observed in the data, and the latter is straightforward to recover from the event study regression (15).

³⁰We 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).

We also need the variance swap 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 (12), from Martin (2017), shows that calculating $V_{i(\tau-1,T)}$ requires integrating over the observed prices of put and call options. We use a daily version of the 3-month LIBOR rate to calculate $R_{\tau-1,T}$. We discretize the integral and calculate the forward price following the methodology used to construct the familiar VIX index (see Cboe, 2019).³¹ The forward price and $S_{i(\tau-1)}$ imply $\tilde{R}_{i(\tau-1,\tau)}$ (see footnote 25). To calculate $V_{i(\tau-1,T)}^H$, we use the counterfactual option prices predicted from a regression like (14), modified to allow the event effect (β_{iK}) to vary by firm and strike.³²

The VIX methodology drops options with a bid of zero as well as some others likely to be illiquid. We keep only firms that have at least three strikes surviving this restriction. And following Proposition 1, we drop firms for which $\tilde{V} < 0$.

We estimate \tilde{p} by taking a weighted average of the \tilde{p}_i . To weight firms, we integrate the inverse of the day $\tau - 1$ relative bid-ask spread over the strikes used in calculating $V_{i(\tau-1,T)}$. A firm receives a high weight if it has liquid options covering a wide range of strikes. We calculate the standard error of \tilde{p} via the delta method, using a covariance matrix robust to clustering by firm and by date.

Finally, any given \tilde{p}_i reflects $p_{\tau-1}^H$ only if firm i was in fact affected by the event. Our preferred specifications therefore again limit the sample to those firms with large t-statistics on $\hat{\gamma}_{i2}$ from regression (15). In contrast to when we estimate \bar{p} as described in Section 4.1, our preferred specifications do not restrict the sign of $\hat{\gamma}_{i2}$.

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, the regular meetings of the Organization of the Petroleum Exporting Countries, uses our new methods to explore the evolution of event uncertainty over the past decade.

³¹We modify the VIX algorithm only to require nonzero bids on both day $\tau - 1$ and day τ .

³²As 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 and option prices, we use a second-order Taylor expansion that adjusts for the standard error of the estimated event effect.

5.1 The 2016 U.S. Presidential Election

On Tuesday November 8, 2016, Donald Trump was elected President of the United States and his Republican party captured both houses of Congress. This extreme outcome was widely surprising.³³ The country learned this outcome in between the close of markets on November 8 and their opening on November 9. On the morning of November 8, prediction markets Betfair and PredictIt gave Trump a 20% and 22% chance of winning, respectively. The bookmakers' Paddy Power and Ladbroke's odds implied that Trump had a 22% and 24% chance of winning, respectively. 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.³⁴ PredictWise gave the Republicans a 33% chance of controlling the Senate and around a 94% chance of controlling the House of Representatives. The New York Times' Upshot and FiveThirtyEight both gave Republicans around a 50% chance of controlling the Senate. PredictIt gave the Republicans a 41% chance of controlling the Senate and a 16% chance of controlling the presidency and both chambers of Congress, implying a 73% chance of winning the Senate conditional on a Trump victory. Our method will recover the joint probability of Trump winning the Presidency and Republicans controlling both the House of Representatives and the Senate. This realized outcome was clearly extreme and it was clearly unlikely, but its exact probability was debatable.

Our methods require that some number of firms be sensitive to the election outcome. The 2016 election clearly passes this test. Figure 3 plots the distribution of t-statistics on firms' event coefficients estimated from regression (15). Most firms' stocks jumped in extremely unusual fashion, with most gaining value. The large number of firms affected by the same event is an ideal setting in which to apply our methods.

Before turning to estimation results, Figure 4 plots the observed ratio of the option price the day of the election to the option price the day after the election ($C_{iK(\tau-1)}/C_{iK\tau}$) against the t-statistic on that firm's stock price event study t-statistic for the same day, $\hat{\gamma}_{i2}$. We plot only the lowest well-traded put option for firms with negative t-statistics (negative stock price jumps on the election news) and

³³This election is particularly attractive 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.

³⁴Whereas 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.

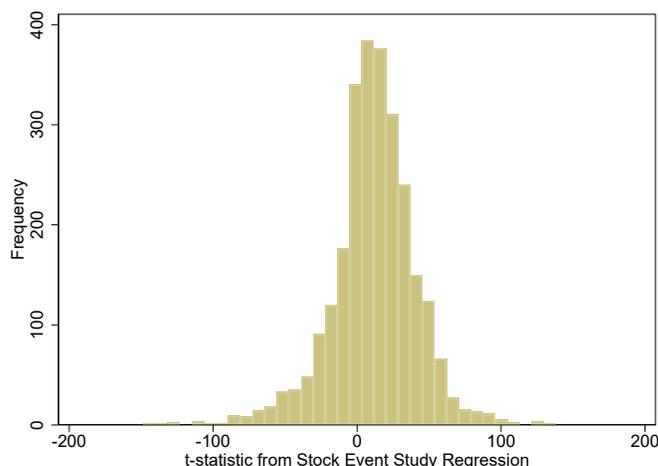


Figure 3: Distribution of statistical significance of firms' responses to the 2016 U.S. election.

the highest well-traded call option for firms with positive t-statistics. The size of the circle represents the average of the inverse spread weight for that option on the event day. Even without regression correction or any smoothing, it is clear that the ratio of option prices begins to converge to a low implied probability for the highest and lowest t-statistics. Further, for put options, the implied probability is somewhat higher than for call options, potentially reflecting persistent bias from disaster risk.

Table 1 reports the estimated probability of the realized 2016 election outcome. 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. These options' extreme strikes make them less vulnerable to upward bias, and we indeed see the estimated probability fall to 0.64. The third column aims to further reduce bias by restricting attention to specifications with call options: it drops firms for which $\hat{\gamma}_{i2} \leq 0$. 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 unaffected by the election (and thus unconnected to $p_{\tau-1}^H$) or not strongly affected by the election (and thus biased upwards). The estimated probability falls substantially, to 0.39. The remain-

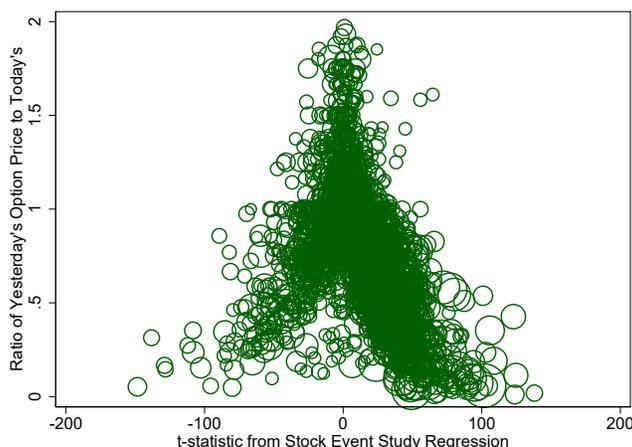


Figure 4: Raw Option Price Ratios Against Firm Event Study t-statistics

From the out-of-the-money approach, each firm's raw jump in option prices (truncated vertically at 2) against the t-statistic from the firm's event study in stock prices. Circle sizes show the event-day inverse average spread ratio weights.

ing columns tighten this restriction further, with the estimated probability falling all the way to 0.12 when the t-statistic threshold is 100.³⁵ This estimated probability is well within the range of pre-election probabilities surveyed above. Moreover, the estimated probability appears to converge for high t-statistic cutoffs.

Figure 5 shows graphically how the estimate of \bar{p} and its confidence interval converge as the regression is run only on firms with higher $\hat{\gamma}_{i2}$ t-statistic cutoffs. This figure makes the bias-variance trade-off clear: the estimated probability does not change substantially above a t-statistic cutoff of about 80, and the standard error starts expanding substantially after that point.

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} .

Figure 6 shows how the firm-by-firm estimate of \tilde{p} changes with the firm's event study t-statistic, where the circle size shows the weight on each firm. Again, the relationship between firm-exposure to the event and the estimated probability is

³⁵If 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 Probabilities for 2016 Presidential Election

	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	1.0316	0.6448	0.6025	0.3866	0.2206	0.1320	0.1191
Standard Error	(0.0031)	(0.0055)	(0.0075)	(0.0064)	(0.0068)	(0.0063)	(0.0236)
# Firms	2,629	2,563	1,835	764	188	43	10
# Options	79,885	2,563	1,835	764	188	43	10
# Option-Days	2,979,681	106,776	77,569	32,057	7,624	1,897	406
Adjusted R^2	0.0171	0.1197	0.1361	0.2773	0.3397	0.4063	0.4985
Variance Swap Approach:							
Probability	0.6182			0.2637	0.1094	0.0817	0.0454
Standard Error				(0.0081)	(0.0096)	(0.0034)	(0.0023)
# Firms	617			196	65	21	4
# Options	7,578			2,329	796	233	59
# Firm-Days	122,783			39,004	12,935	4,179	796
# Option-Days	269,079			83,900	26,609	8,229	1,793

Standard restrictions weights by the average of the inverse bid-ask ratio, 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. For the Variance Swap Approach, each column shows the estimated probability, \tilde{p} , if the absolute value of the event study t-statistic is greater than the cutoff. Standard errors are clustered at both the firm and trade-day level.

clear: the estimated probability converges as the effect of the event on the firm becomes large.

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 $p_{\tau-1}^H$: upon dropping these event-less firms, the estimated \tilde{p} falls to 0.26. This estimate is closer to the converged \bar{p} estimate than was the analogous column of the top panel, which is consistent with the theory predicting that increasing the t-statistic beyond an initial level reduces bias in \bar{p} but is not as critical for \tilde{p} . In fact, tightening the t-statistic cutoff to only 50 recovers a probability consistent with the

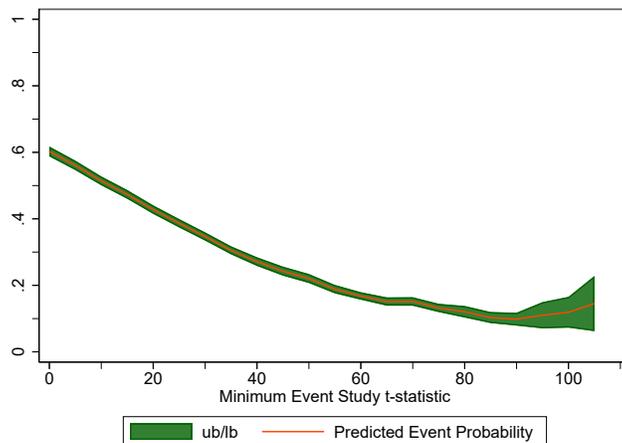


Figure 5: Out-of-the-Money Estimates For Increasing t-statistic Cutoffs

For the 2016 U.S. election and the out-of-the-money approach, each firm’s estimated event probability using different cutoffs in the t-statistic from firms’ event studies in stock prices, with 95% confidence intervals.

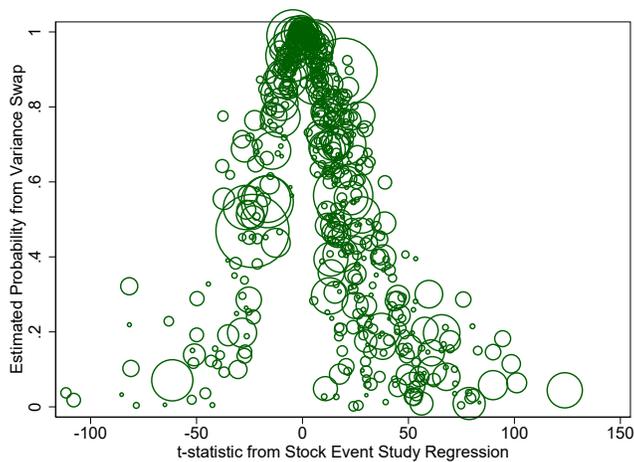


Figure 6: Firm-level Variance Swap Estimates Against t-statistics

For the 2016 U.S. election and the variance swap approach, firm-level estimates of the event probability against the firm’s t-statistic on the event study in stock prices. Circle sizes represent the firm’s weight.

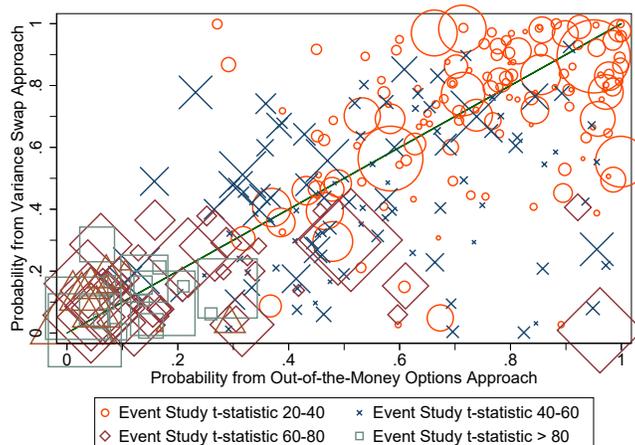


Figure 7: Comparing Firm-Level Versions of the Two Approaches

For the 2016 U.S. election, compares \tilde{p}_i to a firm-by-firm estimate of \bar{p} . Shape sizes represent the weight on \tilde{p}_i

converged \bar{p} .³⁶

Figure 7 compares \tilde{p}_i to the firm-by-firm estimate of \bar{p} .³⁷ Different shapes represent different ranges of the firms' stock price events study t-statistic, and the size of the shapes represent the weight on \tilde{p}_i . Most firms are near the 45-degree line, indicating a high degree of consistency between the two methods despite their different assumptions, their different data (via their different sets of strikes), and their different calculations. In cases where the two estimates differ on whether the probability is small or not, the variance swap approach typically reports the smaller probability, likely because these firms' extreme strikes are not sufficiently extreme to be free of bias. Both approaches recover low probabilities for firms that are substantially affected by the event. Overall, the two approaches' preferred results are driven by obtaining similar estimates from highly weighted firms, despite the different identifying assumptions.

Our new, model-free methods of estimating the risk-neutral probability of an event appear to generate estimates that are compatible with each other and move

³⁶The 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 \bar{p} because this method uses all liquid strikes rather than only a firm's most extreme liquid strike.

³⁷We truncate the horizontal axis at 1. Extending it reveals a number of firms with \tilde{p}_i near 1 and \bar{p}_i much larger than 1.

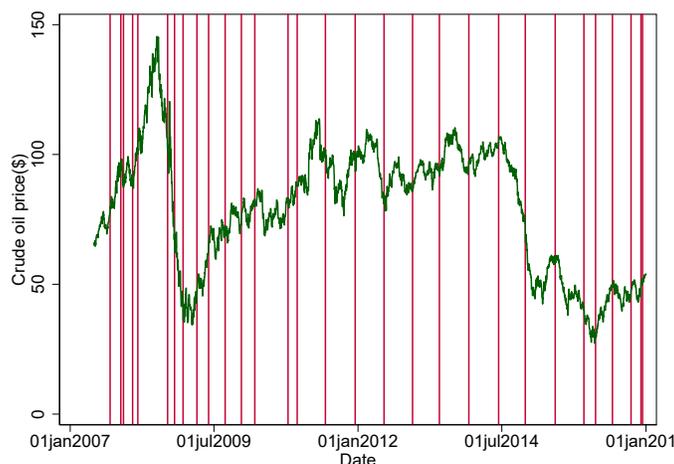


Figure 8: Time series of oil prices, from the front-month WTI contract, with vertical lines marking OPEC meetings and key announcements.

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; Ramelli et al., 2018; Wagner et al., 2018a,b). This election was especially surprising, but our results nonetheless suggest that the event study estimates should be inflated by just over 10% (multiplied by $1/(1 - \hat{p})$) to recover the full effect of the election.

5.2 OPEC Meetings

The Organization of the Petroleum Exporting Countries (OPEC) is a cartel that aims to achieve higher prices by restricting oil supply. OPEC meets at least twice a year to assess its production quotas, with special meetings as warranted. These meetings are typically of high interest, as the price of oil has wide-ranging implications throughout the global economy. We study the 30 OPEC meetings (and related events) from 2007–2016, summarized in Figure 8.

Most meetings result in no change in production quotas, and press reports suggest that most of these outcomes are largely expected. But some meetings do generate news. In a somewhat contentious meeting, OPEC overcame concerns about a slowing global economy to increase production in September 2007. However, by autumn of

2008, a slowing global economy had cut into demand for oil and OPEC was primarily concerned with propping up prices. The members failed to agree on a policy in September but then agreed to production cuts in an October emergency meeting and again in a regular December meeting. The 2008 production quotas would remain unchanged for eight years. By December 2015 and June 2016, disagreement on production quotas was severe enough to keep OPEC from agreeing on changes to their quotas. In September 2016 the members managed to reach a preliminary agreement to cut oil production, which was eventually finalized in a November meeting. Media reports suggest that these latter two meetings were far from forgone conclusions.

The top panel of Figure 9 plots the coefficients from event-study estimates of OPEC meetings' consequences on oil prices.³⁸ Oil prices are generally volatile, and most of these meetings do not stand out from that background noise. The 2007 increase in production had virtually no effect on oil prices, and consistent with contemporary news accounts, oil prices actually fell on the days OPEC announced its 2008 production cuts. These 2007 and 2008 meetings highlight that OPEC's decisions are endogenous to the oil price, in which case they may not affect oil prices in the obvious fashion and may be fairly well anticipated. Oil prices fell significantly after the December 2011 meeting, but this meeting highlights the dangers of contamination from other events: the OPEC meeting generated little news of note, but the European debt and currency crisis took a sharp turn for the worse on that same day as a deal reached at a recent emergency summit began to unravel. Finally, the preliminary and final deals to cut production in late 2016 illustrate that OPEC can indeed affect oil markets. Consistent with intuition, prices jumped significantly upon news of reduced supply.

We construct a time series of probabilities for the realized outcome at OPEC meetings. Most OPEC meetings generated very little news of note. In this context, we cannot select firms based on large t-statistics because we would then estimate an OPEC meeting's probability from firms that happened to have some sort of significant event on that day, whether or not it was related to OPEC. The resulting estimates would be biased towards low probabilities. Instead, we identify a set of firms that should be exposed to OPEC news and estimate \bar{p} and \tilde{p} for this set of firms over time. We use the firms in two GICS sub-industries: oil and gas drilling (10101010), and oil and gas exploration and production (10102020). These firms are directly exposed to the price of oil, and we verify that they respond in a consistent fashion to OPEC's 2016 production cuts.

The bottom panel of Figure 9 plots the estimated \bar{p} (green) and \tilde{p} (orange) for

³⁸We regress futures prices (as log returns) on OPEC event dummies, using an 60-day estimation window.

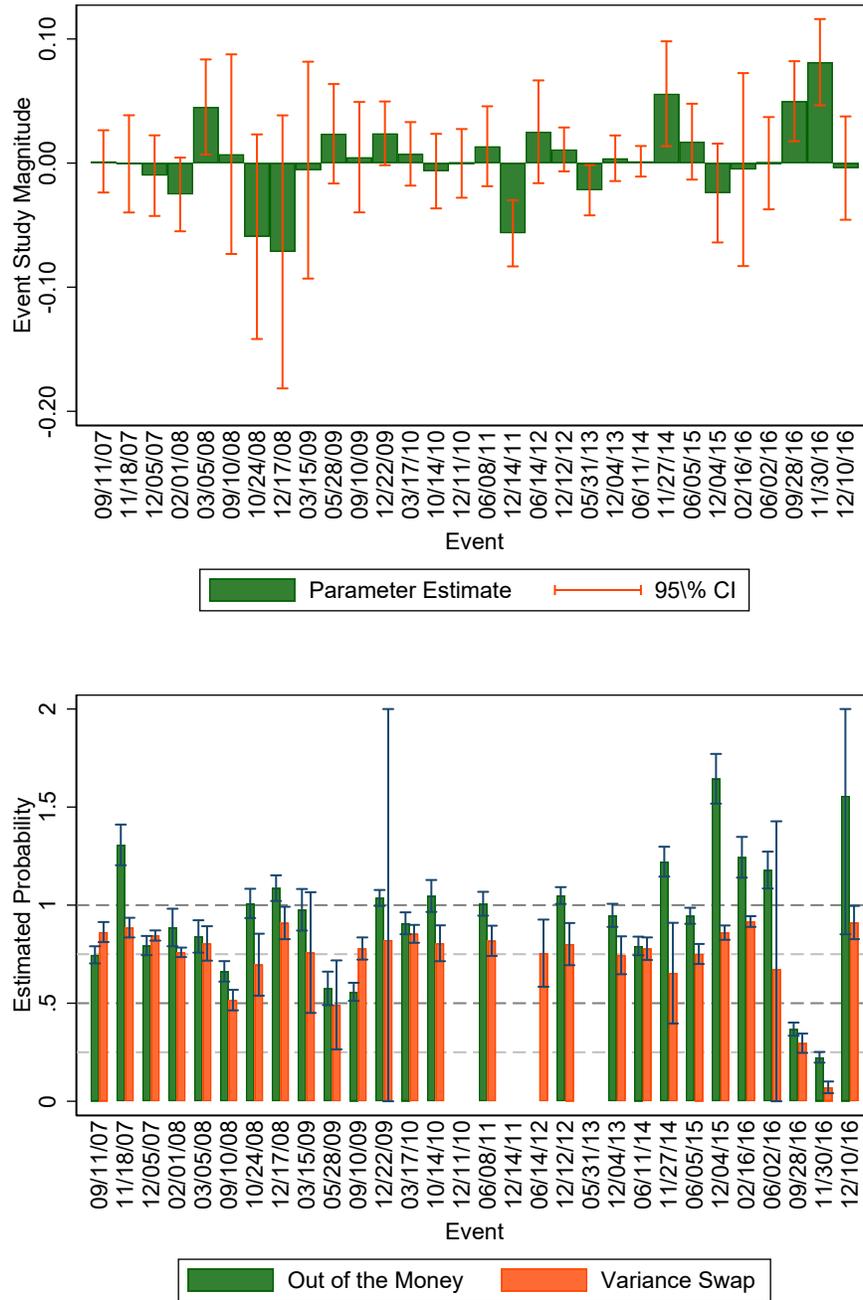


Figure 9: Top: Coefficients and 95% confidence intervals from event studies in the front-month crude oil (WTI) future contract. Bottom: Estimated \bar{p} (green) and \tilde{p} (orange), with 95% confidence intervals (95% confidence intervals are truncated at 0 and 2).

each meeting. These results use our preferred specifications, imposing the theoretical restrictions from the first three columns of Table 1 but no restriction on the magnitude of the t-statistic (see above). The two methods' estimates move in broadly similar ways from event to event, and they generate broadly reasonable probabilities. Most of the meetings' outcomes were indeed well-anticipated, with probabilities between 0.75 and 1. The first of the fall 2008 production cuts was somewhat surprising, with a probability between 0.5 and 0.7. The most surprising events were the two agreements to cut production in the fall of 2016: the preliminary agreement had a probability around 0.3, and the final agreement had a probability less than 0.25.

A few meetings illustrate the value of having two methods. First, the June 2012 meeting had no firms with positive event effects on their stock prices, preventing us from estimating \bar{p} using our preferred specification. We therefore only have a \tilde{p} for this event. Second, many of the estimates for \bar{p} are much larger than 1, reflecting the potential for high upward bias in \bar{p} when events have small effects. Nonetheless, these events have reasonable \tilde{p} . Third, two of the \tilde{p} estimates have very large standard errors, but the corresponding \bar{p} are much more precisely estimated.

Overall, these results highlight the importance of estimating event probabilities. A researcher who ran an event study on the full series of OPEC meetings would obtain results like those in the left panel of Figure 9 and might conclude that OPEC does not have a significant effect on oil markets. However, we here see that the results of OPEC meetings are largely anticipated, so we should not be surprised that they fail to move oil markets. Further, we see that the 2016 production cuts were far more surprising than the 2008 production cuts. It is therefore reasonable that news of the 2008 production cuts was not sufficient to halt the ongoing decline in the price of oil.

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 approaches require no parametric model of stock prices. Both approaches boil down to running event studies in option prices to complement conventional event studies in stock prices. We have demonstrated 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 bookmakers, prediction markets, and polling-driven models, and we estimate probabilities for OPEC meetings that vary in ways consistent with narrative evidence.

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. Second, we recover the probabilities of an event shortly 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. One could in principle construct such a time series using our methods, but doing so would challenge the identifying assumptions underpinning event study regressions in option prices.

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 narrative 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 that 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.
- Ball, Ray (1972) “Changes in accounting techniques and stock prices,” *Journal of Accounting Research*, Vol. 10, pp. 1–38.
- Barracough, 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.
- Bernanke, Ben S. and Kenneth N. Kuttner (2005) “What explains the stock market’s reaction to Federal Reserve policy?” *The Journal of Finance*, Vol. 60, No. 3, pp. 1221–1257.
- 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 (2019) “Trade policy uncertainty and stock returns,” SSRN Scholarly Paper ID 3340700, Social Science Research Network, Rochester, NY.
- 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.
- Bittlingmayer, George and Thomas W Hazlett (2000) “DOS Kapital: Has antitrust action against Microsoft created value in the computer industry?” *Journal of Financial Economics*, Vol. 55, No. 3, pp. 329–359.
- Bizjak, John M. and Jeffrey L. Coles (1995) “The effect of private antitrust litigation on the stock-market valuation of the firm,” *The American Economic Review*, Vol. 85, No. 3, pp. 436–461.

- Björk, Tomas (2004) *Arbitrage Theory in Continuous Time*, Oxford; New York: Oxford University Press, 2nd edition.
- Black, Fischer (1976) “The pricing of commodity contracts,” *Journal of Financial Economics*, Vol. 3, No. 1, pp. 167–179.
- Black, Fischer and Myron Scholes (1973) “The pricing of options and corporate liabilities,” *Journal of Political Economy*, Vol. 81, No. 3, pp. 637–654.
- 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.
- Breeden, Douglas T. and Robert H. Litzenberger (1978) “Prices of state-contingent claims implicit in option prices,” *The Journal of Business*, Vol. 51, No. 4, pp. 621–651.
- 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.
- 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 (1997) *Myth and Measurement*, Princeton, NJ: Princeton University Press, reprint edition edition.
- 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.
- Cox, John C., Stephen A. Ross, and Mark Rubinstein (1979) “Option pricing: A simplified approach,” *Journal of Financial Economics*, Vol. 7, No. 3, pp. 229–263.
- Cutler, David M., James M. Poterba, and Lawrence H. Summers (1989) “What moves stock prices?” *Journal of Portfolio Management; New York*, Vol. 15, No. 3, p. 4.
- Dann, Larry Y. and Christopher M. James (1982) “An analysis of the impact of deposit rate ceilings on the market values of thrift institutions,” *The Journal of Finance*, Vol. 37, No. 5, pp. 1259–1275.
- 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.
- Ederington, Louis H. and Jae Ha Lee (1996) “The creation and resolution of market uncertainty: The impact of information releases on implied volatility,” *The Journal of Financial and Quantitative Analysis*, Vol. 31, No. 4, pp. 513–539.
- 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, David, Raymond J. Fisman, Julia Galef, Rakesh Khurana, and Yongxiang Wang (2012) “Estimating the value of connections to Vice-President Cheney,” *The B.E. Journal of Economic Analysis & Policy*, Vol. 12, No. 3.
- Fisman, Raymond (2001) “Estimating the value of political connections,” *The American Economic Review*, Vol. 91, No. 4, pp. 1095–1102.
- 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.
- Hallock, Kevin F. (1998) “Layoffs, top executive pay, and firm performance,” *The American Economic Review*, Vol. 88, No. 4, pp. 711–723.
- Hennessy, Christopher A and Ilya A Strebulaev (2020) “Beyond random assignment: Credible inference of causal effects in dynamic economies,” *Journal of Finance*, Vol. forthcoming.
- 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.
- Heston, Steven L. (1993) “A closed-form solution for options with stochastic volatility with applications to bond and currency options,” *The Review of Financial Studies*, Vol. 6, No. 2, pp. 327–343.
- Hughes, Paul (2006) “Do EPA defendants prefer Republicans? Evidence from the 2000 election,” *Economic Inquiry*, Vol. 44, No. 3, pp. 579–585.
- 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.

- James, Christopher (1983) “An analysis of intra-industry differences in the effect of regulation: The case of deposit rate ceilings,” *Journal of Monetary Economics*, Vol. 12, No. 3, pp. 417–432.
- 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.
- 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.
- 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.
- Lee, Jieun and Doojin Ryu (2019) “The impacts of public news announcements on intraday implied volatility dynamics,” *Journal of Futures Markets*, Vol. forthcoming.
- Lemoine, Derek (2017) “Green expectations: Current effects of anticipated carbon pricing,” *Review of Economics and Statistics*, Vol. 99, No. 3, pp. 499–513.
- 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.
- MacKinlay, A. Craig (1997) “Event studies in economics and finance,” *Journal of Economic Literature*, Vol. 35, No. 1, pp. 13–39.
- 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.
- 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.
- Prince, David W. and Paul H. Rubin (2002) “The effects of product liability litigation on the value of firms,” *American Law and Economics Review*, Vol. 4, No. 1, pp. 44–87.
- Ramelli, Stefano, Alexander F Wagner, Richard J Zeckhauser, and Alexandre Ziegler (2018) “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.
- Ruback, Richard S. and Martin B. Zimmerman (1984) “Unionization and profitability: Evidence from the capital market,” *Journal of Political Economy*, Vol. 92, No. 6, pp. 1134–1157.
- 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.

van Tassel, Peter (2016) “Merger options and risk arbitrage,” Technical Report 761, Federal Reserve Bank of New York.

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.

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.

Appendix

A Extensions to Section 3.1

A.1 Theory with 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 k . 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.

A.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 $k \in \{L, M, H\}$ such that $S(\omega_t, M) > \bar{S}$ implies $S(\omega_t, H) > S(\omega_t, M) > S(\omega_t, L)$. Assume that $k = M$ is realized.³⁹ 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

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

³⁹If either $k = L$ or $k = H$ were realized, then the analysis in the main text holds, because we can combine $k = M$ with whichever other value for k was not realized. In addition, partitioning the event space into three possible values is not restrictive: if, for instance, there were $k \in \{L_1, L_2, M, H\}$ 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 $k = M$ 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:

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

where $\neg M$ means that $k \in \{L, H\}$. Of course $\Gamma_{\tau-1,T}^{\neg M}(S_{\tau-1}^{\neg M}, K_1, K_2) \geq 0$. 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.⁴⁰ The bound becomes tighter as $\Gamma_{\tau-1,T}^{\neg M}(S_{\tau-1}^{\neg M}, K_1, K_2)$ 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

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

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 .

B Derivations for the Variance Swap Analysis

B.1 Equation (9)

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

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

⁴⁰Intuitively, 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.

B.2 Proof of Proposition 1

Using the assumption that k 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:

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

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 (B-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:

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

This is the analogue of equation (10), 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 (B-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 (B-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 $\tilde{p} > 1$.