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

Ashley Langer† and Derek Lemoine‡

PRELIMINARY AND INCOMPLETE.

April 26, 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 overcome this widely known limitation by developing two nonparametric methods for recovering the market's priced-in probability of events from the prices of widely traded financial options. These methods involve running event studies on options prices to complement the standard event study in stock prices. We demonstrate the power of our new methods through applications to prominent events in health care regulation: we recover probabilities for two elections, a court case, and a legislative event that are consistent with available evidence such as polling data and prediction markets. These probabilities suggest that a conventional event study would draw incorrect conclusions about which events were most consequential for firms.

JEL: C58, G13, I11, I18

Keywords: event study, anticipation, policy uncertainty, options, derivatives, health care, Affordable Care Act

*We are grateful to 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.

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

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

1 Introduction

The event study methodology measures the consequences of news for firms' stock prices. 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 here 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 impose parametric assumptions on the distribution of stock prices. We apply our methods to recover probabilities for two U.S. elections that have been the subject of recent event studies, finding probabilities that are consistent with polling data and remarkably in line with prediction markets. We also recover probabilities for a court case and a legislative event, two types of events that have challenged the event study methodology because they are often well-anticipated. Our estimates suggest that these events were more likely that not and thus that traditional event studies would dramatically underestimate the full effect of these events.

Our extension of the event study methodology is important for increasing its relevance for policy. 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), 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 their review, MacKinlay (1997, 37) lament 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 only gradually.

We develop two new nonparametric techniques for estimating the market 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 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. If 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 the event's probability relies on a different identifying assumption. A variance swap rate reveals the market's expected volatility of stock prices over some horizon. The pre-event variance swap rate includes the volatility 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 volatility 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 volatility but retains the volatility induced by uncertainty about the event's realization. Recent results show that variance swap rates can be synthesized from a linear combination of option prices under quite general assumptions. This second method therefore again relies on changes in option prices to identify the

[1] In practice, options with extreme strike prices may not be worthless if the event fails to occur because there may still be a small chance of reaching an extreme stock price. In this case, our estimated probability is an upper bound on the market's priced-in probability of the event occurring. We can use that upper bound to obtain an upper bound on an event's full effect to complement the conventional lower bound obtained by assuming the event was a complete surprise. We provide evidence that theoretically motivated restrictions on the set of options used in estimation lead to a tight bound in our applications.

event's probability, but now using the full set of option strikes traded on a firm and undertaking a different type of calculation with them. 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.

We demonstrate the accuracy and broad applicability of our new methods by estimating the probability of several events that were important to recent U.S. health care policy. We show that our two methods recover probabilities consistent with each other and consistent with outside evidence from polling data and prediction markets. First, we extend the event study of Al-Ississ and Miller (2013), who use a March 2010 Massachusetts special Senate election to estimate the costs of the Affordable Care Act (ACA, "Obamacare"). We estimate that the probability of this event was 0.54, a number consistent with polling trends and prediction markets leading up to the election. The estimates in Al-Ississ and Miller (2013) therefore miss over half of the event's effect.

Second, we estimate the effect of the November 2016 U.S. election on health care firms. In contrast to the 2010 Senate special election, the 2016 election of Donald Trump to the presidency and the Republican sweep of Congress were considered unlikely. We estimate a probability of 0.22 from either of our methods. This probability is in the middle of prominent pre-election polling-driven estimates and is virtually identical to the probabilities implied by several prediction markets' contracts and by several brokers' odds. The almost exact match to outside estimates for this widely discussed event is strong evidence that our methods can pick up event probabilities from option prices.

Finally, we estimate the probability of the Supreme Court ruling in June 2012 that upheld the individual mandate to purchase health insurance and the probability of the rejection in July 2017 of the "skinny repeal" of the ACA, best remembered for Senator John McCain's "thumbs down" vote. We estimate probabilities for these events of 0.76 and 0.81, respectively, making them substantially more likely than our election estimates. The probability of the Supreme Court ruling is consistent with previous literature that used a parametric model of option prices (Borochin and Golec, 2016) and with investment analysts' suggestions but inconsistent with prediction markets, which implied a much lower probability. Given the coherence of the other estimates and the fact that they were more accurate ex post, we might wonder whether about the accuracy of prediction markets for events in which all actors have access to the same minimal information. The Senate vote highlights an additional challenge facing prediction markets: because they fluctuated rapidly within the days leading up to the vote as news events unfolded, it would be impossible to know which probability within a broad range of possible probabilities corresponds to the probability that had been priced-in by markets' close the day before the event.

We use our estimated event probabilities to recover the full magnitude of the events' effects on stock prices. Because we find that the 2016 election results were fairly surprising, the estimates from a standard event study should be inflated by only 28.5% to calculate the full effect of electing President Trump and a Republican Congress. However, since the Massachusetts special election outcome was more likely, its event study estimates need to be inflated by 116% to recover the full event effect. And since the June 2012 Supreme Court ruling and the "skinny repeal" of the ACA outcomes were quite likely, their event study estimates need to be inflated by 314% and 437%, respectively, to recover their full effects. Adjusting for the probability of the event can therefore make an enormous difference to event study estimates of the ACA. Further, whereas a traditional event study would conclude that the 2016 presidential election was the most important of the four events for health care firms, our corrected estimates show that the 2012 Supreme Court decision upholding the individual mandate was in fact a substantially larger event for health care facility companies (i.e., hospitals) and that the "skinny repeal" vote was particularly important for pharmaceutical companies. Without these probabilities, a conventional event study would draw incorrect conclusions about the most important drivers of health care policy.

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. 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 effects at all. 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) affect corporate profits at all. These studies test the joint hypothesis that the particular event does affect profits and that the event was sufficiently surprising. Each study finds effects that are either small or not significantly different from zero. However, as these researchers recognize, we cannot tell whether these results arise because the decisions themselves only have small effects (perhaps even no effect) or because the events were well-anticipated.[2] If our methods indicate that such events were in fact surprising, then researchers may have greater confidence that the decision indeed has only small effects. If our methods indicate that such events were well-anticipated, then researchers may be more hesitant to conclude that the true effect is small.

Third, 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 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 within firms. Because the net effect of innovation probably 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.

Fourth, many researchers are interested in explaining the 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 also 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 the 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.

Fifth, 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

[2] Cutler 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 effectively random.[3] The outcomes of these elections can then be used as the randomized experiments that have been sought in elections that were close ex post.

Sixth, 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.[4] If researchers show that some events are truly random, then these events' risk-neutral probabilities tell us whether investors 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 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.

Finally, 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 this difference in tariffs. 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.

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.[5] The standard solution is to select events that the researcher judges to be

[3] This 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.

[4] Risk-neutral probabilities are the probabilities needed to correct event study estimates.

[5] Some 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),

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.[6]

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; Linn, 2010; Lee and Mas, 2012; Al-Ississ and Miller, 2013).[7] In some cases, the event 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).[8] Prediction markets can be a valuable source of information when the proper contracts exist,

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.

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

[7]Bernanke 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 regarding prediction markets).

[8]Working 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.

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 financial markets at the time they close. Recovering probabilities directly from options markets avoids this difficulty.

Several papers in the finance literature also attempt to infer the priced-in probability of events from the prices of financial options.[9] These papers assume that options are priced according to specific parametric models and search for the event probability that reconciles observed option prices and theoretical option prices. 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.[10] However, each pricing model imposes specific parametric assumptions on the distribution of stock prices that may be violated in practice—and the assump-

[9] Some 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 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 methods that allow the date to be unknown in advance. Finally, 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 will instead focus on policy events that affect a cross-section of firms and are not endogenous to any firm.

[10] As 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.

tions imposed by each of these papers conflict with the assumptions imposed in the others. 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.[11] Our new options-based methods do not impose parametric assumptions. Our methods for estimating the event probability require only the absence of arbitrage and either (i) that some out-of-the-money options have value only if the event occurs or (ii) that the expected post-event volatility does not depend on the event's realization.

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 shows the results of our empirical application to ACA events. Section 6 concludes. The appendix contains proofs, extensions to the theoretical analysis, and additional empirical results.

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, $k \in \{L, H\}$.[12] $S(\omega_t, k)$ 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 $S(\omega_t, H)$. 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.

[11] Carvalho 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.

[12] As we will see, the 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.

At time τ, an event happens that reveals the state to be either H or L. Time t agents know ω_t, but prior to time τ, agents do not know whether $k = H$ or $k = L$.[13] 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 $k = L$.[14] For instance, consider an election for president 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 instead assign $p_t^L = 1$ if L won.

2.1 The Bias in the Standard Event Study Methodology

Before continuing to our approach to estimating p_t^H, we show why the event study approach cannot estimate the full market value of an event except in very specific contexts. Without loss of generality, assume that event H occurs at time τ. For ease of exposition, consider a case with only a single firm. The event study's identifying assumption is that the controls account for all elements of ω_τ that differ from $\omega_{\tau-1}$.[15] 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 reflects the new information about k. The event study can therefore recover $S_{\tau-1}^H$ and calculates the event effect as $S_{\tau-1}^H - S_{\tau-1}$.

However, the researcher would like to calculate $S_{\tau-1}^H - S_{\tau-1}^L$. By absence of arbitrage, the time $\tau - 1$ stock price must be:

$$S_{\tau-1} = p_{\tau-1}^L S_{\tau-1}^L + p_{\tau-1}^H S_{\tau-1}^H. \tag{1}$$

We then have:

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

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

[14] A risk-neutral measure exists in the absence of arbitrage and is unique in complete markets. The risk-neutral measure can be interpreted as embedding 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).

[15] See 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.

The estimated event effect $S^H_{\tau-1} - S_{\tau-1}$ is less than the full event effect $S^H_{\tau-1} - S^L_{\tau-1}$. As $p^H_{\tau-1} \to 0$, the event study does recover $S^H_{\tau-1} - S^L_{\tau-1}$ from $S^H_{\tau-1} - S_{\tau-1}$: outcome H was judged at time $\tau - 1$ to be extremely unlikely (or even impossible), so when outcome H nonetheless occurs, the researcher sees the entire effect of outcome H relative to outcome L. For this reason, researchers have sought events that are complete surprises. However, for $p^H_{\tau-1} > 0$, the event study underestimates $S^H_{\tau-1} - S^L_{\tau-1}$ because $S_{\tau-1}$ already reflects the possibility of outcome H. This is the well-known problem of partially anticipated events. Moreover, as $p^H_{\tau-1}$ goes to 1, the event study measures an arbitrarily small fraction of the true event effect $S^H_{\tau-1} - S^L_{\tau-1}$. The bias in event studies' estimates can therefore become arbitrarily large. Event studies do recover a lower bound for the implications of an event, but they fail to provide an upper bound unless the priced-in probability of the event can be bounded from above.

3 Two Nonparametric Approaches to Recovering the Event Probability from Options Data

We now describe two new approaches to recovering $p^H_{\tau-1}$ from running event studies in options prices.

3.1 As Implied By Out-of-the-Money Options

A call (put) option on the 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.[16] Consider a call option on S. If it happens that $S_T > K$, then the holder of the call option will buy the stock at K and sell it for S_T, netting a profit of $S_T - K$. But if $S_T < K$, the holder of the call option will prefer to buy the stock on the market rather than buy it at K. In that case, the option's holder will allow the option to expire without exercising it, netting a payoff of 0. Therefore, at some time x prior to T, the price $C_{x,T}(K)$ of the call option should satisfy

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

[16] For the present purposes, we ignore the distinction between "European" and "American" options, which depends on the ability 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.

where $f_x(S_T|S_x)$ is the risk-neutral distribution of S_T under the time x information set (with the conditioning making dependence on S_x explicit) and $R_{x,T} \geq 1$ is the gross risk-free rate from time x to T. The price of the option is the expected payoff conditional on the price being above K. At time $\tau - 1$, the price of a call option with strike K and expiration $T > \tau - 1$ must satisfy:

$$\begin{aligned}
C_{\tau-1,T}(S_{\tau-1}, K) =& \frac{1}{R_{\tau-1,T}} \int_K^\infty (S_T - K)\, f_{\tau-1}(S_T|S_{\tau-1})\, \mathrm{d}S_T \\
=& \frac{1}{R_{\tau-1,T}} \int_K^\infty (S_T - K)\, \left[p^L_{\tau-1} f_{\tau-1}(S_T|S^L_{\tau-1}, L) + p^H_{\tau-1}\, f_{\tau-1}(S_T|S^H_{\tau-1}, H)\right] \mathrm{d}S_T \\
=& p^L_{\tau-1} C^L_{\tau-1,T}(S^L_{\tau-1}, K) + p^H_{\tau-1} C^H_{\tau-1,T}(S^H_{\tau-1}, K),
\end{aligned}$$

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

Now imagine that H happens to be realized at time τ. Consider how this event changes the option's price:

$$\begin{aligned}
& C^H_{\tau-1,T}(S^H_{\tau-1}, K) - C_{\tau-1,T}(S_{\tau-1}, K) \\
=& \frac{1}{R_{\tau-1,T}} (1 - p^H_{\tau-1}) \int_K^\infty (S_T - K)\, \left[f_{\tau-1}(S_T|S^H_{\tau-1}, H) - f_{\tau-1}(S_T|S^L_{\tau-1}, L)\right]\, \mathrm{d}S_T \\
=& (1 - p^H_{\tau-1})\, C^H_{\tau-1,T}(S^H_{\tau-1}, K) - (1 - p^H_{\tau-1})\, C^L_{\tau-1,T}(S^L_{\tau-1}, K).
\end{aligned} \tag{3}$$

A standard arbitrage bound (e.g., Cochrane, 2005) requires $C^L_{\tau-1,T}(S^L_{\tau-1}, K) \geq 0$. Using this inequality in equation (3) implies:[17]

$$p^H_{\tau-1} \leq \frac{C_{\tau-1,T}(S_{\tau-1}, K)}{C^H_{\tau-1,T}(S^H_{\tau-1}, K)} \triangleq \bar{p}. \tag{4}$$

The $p^H_{\tau-1}$ in inequality (4) is the same $p^H_{\tau-1}$ as in equation (2). Thus, we can use the observed changes in option prices to bound the risk-neutral probability of the realized outcome H and thereby also bound the bias in the event study measure (2).

Now assume that there exists some $\bar{S}$ such that $S(\omega_t, H) > \bar{S}$ implies $S(\omega_t, H) > S(\omega_t, L)$.[18] Also assume that $f_{\tau-1}(S_T|S^L_{\tau-1}, L)$ goes to zero as S_T becomes large.

[17] The other two arbitrage bounds ($C^L_{\tau-1,T}(S^L_{\tau-1}, K) \geq S^L_{\tau-1} - K/R_{\tau-1,T}$ and $C^L_{\tau-1,T}(S^L_{\tau-1}, K) \leq S^L_{\tau-1}$) also imply upper bounds on $p^H_{\tau-1}$ 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.

[18] This 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 $\{L_1, ..., L_N\}$, then this assumption requires that $S(\omega_t, H) > \bar{S}$ implies $S(\omega_t, H) > \sum_{i=1}^N [p^{L_i}_{\tau-1}/(1 - p^H_{\tau-1})] S(\omega_t, L_i)$. The appendix relaxes the assumption that the realized outcome is extreme.

These two assumptions together imply that increasing K brings $C^L(S^L_{\tau-1}, K)$ to zero faster than it brings $C^H(S^H_{\tau-1}, K)$ to zero. The arbitrage bound $C^L_{\tau-1,T}(S^L_{\tau-1}, K) \geq 0$ then holds exactly as K becomes large, which means that $\bar{p}$ converges to $p^H_{\tau-1}$ for some sufficiently large K. Because $C^H(S^H_{\tau-1}, K)$ and $C^L(S^L_{\tau-1}, K)$ 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 illustrates the intuition underlying inequality (4). The solid and long-dashed curves depict the two elements of the distribution of S_T viewed from the time $\tau-1$ information set: the risk-neutral distribution conditional on H weighted by the probability of H (solid), and the risk-neutral distribution conditional on L weighted by the probability of L (long-dashed). The short-dashed curve aggregates these two weighted distributions to give the unconditional risk-neutral distribution of S_T at the time $\tau-1$ information set. At time $\tau-1$, an option with strike K is priced as the conditional expectation of $S_T - K$ in the tail of the dashed distribution, as indicated by the shaded area A. For the sufficiently large K plotted here, this tail does not include any significant piece of the distribution conditioned on L. The dotted curve gives the time $\tau-1$ risk-neutral distribution conditional on outcome H being realized. The left tail of this distribution differs from the short-dashed distribution because it no longer mixes in the distribution for L, and the right tail of this distribution differs from the short-dashed distribution because it is no longer scaled down to reflect $p^H_{\tau-1} < 1$. The time τ price of the option reflects the expectation of $S_T - K$ in the combination of areas A and B. For the depicted K, the change in this option's price from $C_{\tau-1,T}(S_{\tau-1}, K)$ to $C^H_{\tau-1,T}(S^H_{\tau-1}, K)$ simply reflects the rescaling from $p^H_{\tau-1} < 1$ to $p^H_{\tau-1} = 1$.[19]

The bound $\bar{p}$ may be especially tight when $p^H_{\tau-1}$ is large. From equation (3), the bias is

$$\bar{p} - p^H_{\tau-1} = (1 - p^H_{\tau-1}) \frac{C^L_{\tau-1,T}(S^L_{\tau-1}, K)}{C^H_{\tau-1,T}(S^H_{\tau-1}, K)}. \tag{5}$$

As $p^H_{\tau-1}$ grows, the bias vanishes, both because $1 - p^H_{\tau-1}$ shrinks and because $S^L_{\tau-1}$ shrinks (for given observables $S_{\tau-1}$ and $S^H_{\tau-1}$). In Figure 1, large $p^H_{\tau-1}$ corresponds

[19] This case with large K is the case in which the relation in (4) holds with approximate equality because the arbitrage bound $C^L_{\tau-1,T}(S^L_{\tau-1}, K) \geq 0$ holds with approximate equality. For smaller K, the long-dashed distribution may have nontrivial mass in region A. In this case, the change in the price of the option reflects both the rescaling by $p^H_{\tau-1}$ and the loss of this unobserved probability mass. The possibility of unobserved probability mass explains why (4) is an inequality rather than an equality.

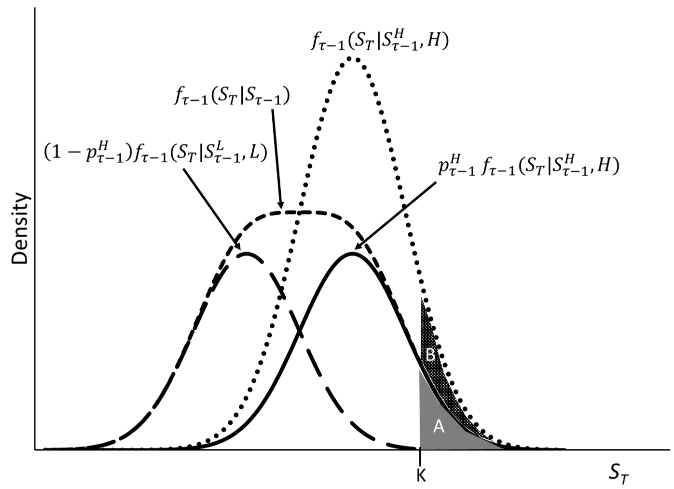

Figure 1: Illustration of how changes in option prices identify the prior probability $p^H_{\tau-1}$ when event H is realized at time τ.

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. Importantly, the case with large $p^H_{\tau-1}$ is precisely the case in which standard event studies suffer arbitrarily large biases (equation (2)). We therefore estimate an especially tight bound on $p^H_{\tau-1}$ precisely when a tight bound is most needed.

Put options can also recover the event probability, which is useful when an event reduces the price of the underlying asset. For ease of exposition, 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

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

An econometrician can estimate $P^L_{\tau-1,T}(S^L_{\tau-1}, K)$ through a standard event study design. A derivation analogous to the previous case yields:

$$\begin{aligned}
&P^L_{\tau-1,T}(S^L_{\tau-1}, K) - P_{\tau-1,T}(S_{\tau-1}, K) \\
=&\frac{1}{R_{\tau-1,T}}(1 - p^L_{\tau-1}) \int_{K}^{\infty} (S_T - K)\ \left[f_{\tau-1}(S_T|S^L_{\tau-1}, L) - f_{\tau-1}(S_T|S^H_{\tau-1}, H)\right]\ \mathrm{d}S_T \\
=&(1 - p^L_{\tau-1})\, P^L_{\tau-1,T}(S^L_{\tau-1}, K) - (1 - p^L_{\tau-1})\, P^H_{\tau-1,T}(S^H_{\tau-1}, K). \qquad (6)
\end{aligned}$$

The relevant arbitrage bound becomes $P^H_{\tau-1,T}(S^H_{\tau-1}, K) \geq 0$. Using this in equation (6) implies:

$$p^L_{\tau-1} \leq \frac{P_{\tau-1,T}(S_{\tau-1}, K)}{P^L_{\tau-1,T}(S^L_{\tau-1}, K)}. \tag{7}$$

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, $f_{\tau-1}(S_T|S^H_{\tau-1}, H) \to 0$ as S_T becomes small, then the arbitrage bound $P^H_{\tau-1,T}(S^H_{\tau-1}, K) \geq 0$ holds exactly as K becomes small. In that case, the right-hand side of inequality (7) converges to $p^L_{\tau-1}$ 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.[20] The possibility of such disasters is plausibly independent of 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 (7) then only loosely bounds $p^L_{\tau-1}$ and does not converge to $p^L_{\tau-1}$ 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 $p^L_{\tau-1}$ via inequality (7) may actually increase as the strike price falls. It is then 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 strike price and even vanish for some sufficiently large strike prices.

3.2 As Implied By Variance Swaps

The previous method of estimating the priced-in event probability $p^H_{\tau-1}$ relied on changes in the tail of the distribution of S_T, as reflected in option prices. That method placed an upper bound on $p^H_{\tau-1}$ 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 variance swap rates that can be synthesized by the set of traded options. This method does require advance knowledge of the date that the event will happen, but it identifies the event probability under a different set

[20] Since 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 Ursa, 2012).

of circumstances, now requiring that the post-event expected volatility of the stock price not depend on the event's realization.

Assume that market actors know at time $\tau - 1$ that the event will happen at time τ (i.e., they know that they will learn whether $k = H$ or $k = L$ by time τ). Consider the variance of time τ payoffs, as viewed from the time $\tau - 1$ information set. As in previous sections, an econometrician observes $S_{\tau-1}$ and can construct $S^H_{\tau-1}$ from an event study. Temporarily fix $R_{\tau-1,\tau} = 1$. The appendix shows that:

$$Var_{\tau-1}[S_\tau] = \frac{p^H_{\tau-1}}{1 - p^H_{\tau-1}} (S^H_{\tau-1} - S_{\tau-1})^2 + p^H_{\tau-1} Var_{\tau-1}\left[S^H_\tau\right] + (1 - p^H_{\tau-1}) Var_{\tau-1}\left[S^L_\tau\right]. \tag{8}$$

Figure 2 illustrates the intuition. We can think of the change from time $\tau - 1$ to time τ as being determined by the outcome of a compound lottery. A first lottery determines whether $k = L$ or $k = H$, which induces variance captured by the first term on the right-hand side of equation (8). This variance depends on uncertainty about k and on the stock price's sensitivity to the realization of k. With k realized, a second lottery then determines the observed stock price. This lottery captures all other information revealed between times $\tau - 1$ and τ. The second and third terms on the right-hand side of equation (8) capture the variance induced by this second lottery, allowing it to depend on the realization of k.

Imagine that a portfolio of options replicates $Var_{\tau-1}[S_\tau]$, in which case an econometrician can construct a similar replicating portfolio for $Var_{\tau-1}[S^H_\tau]$ by applying event study methods to option prices. Subtracting $Var_{\tau-1}[S^H_\tau]$ from each side of equation (8) and rearranging, we have:

$$\frac{p^H_{\tau-1}}{1 - p^H_{\tau-1}} = \frac{Var_{\tau-1}[S_\tau] - Var_{\tau-1}[S^H_\tau]}{(S^H_{\tau-1} - S_{\tau-1})^2} + (1 - p^H_{\tau-1}) \frac{\overbrace{\left(Var_{\tau-1}\left[S^H_\tau\right] - Var_{\tau-1}\left[S^L_\tau\right]\right)}^{\Delta Var}}{(S^H_{\tau-1} - S_{\tau-1})^2}. \tag{9}$$

Both $Var_{\tau-1}[S^L_\tau]$ and $p^H_{\tau-1}$ 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. Rearranging, we then have:

$$p^H_{\tau-1} = \frac{Var_{\tau-1}[S_\tau] - Var_{\tau-1}[S^H_\tau]}{Var_{\tau-1}[S_\tau] - Var_{\tau-1}[S^H_\tau] + (S^H_{\tau-1} - S_{\tau-1})^2}.$$

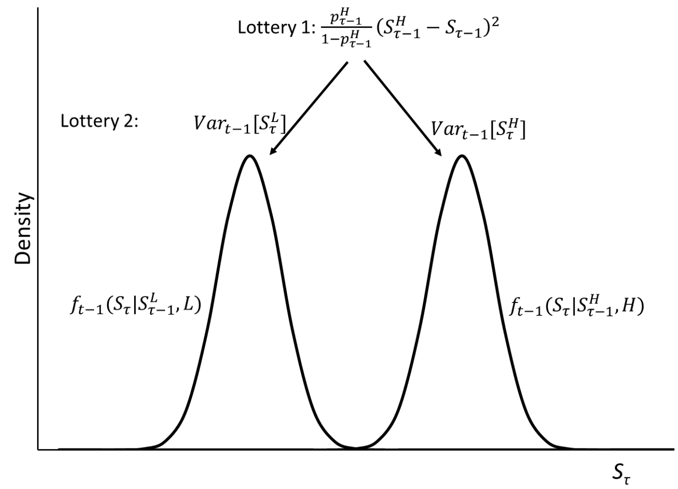

Figure 2: The variance of the time τ (post-event) stock price accounts for uncertainty about the realization of k and for the variance of the stock price conditional on each k, from equation (8).

We can estimate $p^H_{\tau-1}$ from the replicating portfolios for the variance from $\tau-1$ to τ.[21]

Thus far, we have seen how we might estimate $p^H_{\tau-1}$ 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.[22] However, the desired variance swap rates may rarely be 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

[21]Viewing $p^H_{\tau-1}$ as implicitly defined as a function of ΔVar, we have, using equation (2):

$$\left.\frac{\mathrm{d}p^H_{\tau-1}}{\mathrm{d}\Delta Var}\right|_{\Delta Var=0}=\frac{(1-p^H_{\tau-1})^3}{(S^H_{\tau-1}-S_{\tau-1})^2}=\frac{1-p^H_{\tau-1}}{(S^H_{\tau-1}-S^L_{\tau-1})^2}\geq 0.$$

If we estimate $p^H_{\tau-1}$ under the assumption that $\Delta Var=0$, then the bias from small deviations in ΔVar is small when $p^H_{\tau-1}$ 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 $p^H_{\tau-1}$.

[22]The 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.

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.[23] 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.[24] Martin (2017) shows that

$$V_{\tau-1,T} = \frac{2R_{\tau-1,T}}{[\tilde{R}_{\tau-1,T}S_{\tau-1}]^2}\left\{\int_0^{\tilde{R}_{\tau-1,T}S_{\tau-1}} P_{\tau-1,T}(S_{\tau-1},K)\,\mathrm{d}K + \int_{\tilde{R}_{\tau-1,T}S_{\tau-1}}^{\infty} C_{\tau-1,T}(S_{\tau-1},K)\,\mathrm{d}K\right\}. \quad (10)$$

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

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

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

$$V^H_{\tau-1,T} = \frac{2R_{\tau-1,T}}{[\tilde{R}_{\tau-1,T}S^H_{\tau-1}]^2}\left\{\int_0^{\tilde{R}_{\tau-1,T}S^H_{\tau-1}} P^H_{\tau-1,T}(S^H_{\tau-1},K)\,\mathrm{d}K + \int_{\tilde{R}_{\tau-1,T}S^H_{\tau-1}}^{\infty} C^H_{\tau-1,T}(S^H_{\tau-1},K)\,\mathrm{d}K\right\}.$$

The following proposition relates $p^H_{\tau-1}$ to $V_{\tau-1,T}$ and $V^H_{\tau-1,T}$:

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

[24] The 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 rates and will use options with the shortest maturities.

Proposition 1. *Define*

$$\tilde{V} \triangleq (S_{\tau-1})^2 V_{\tau-1,T} - \left(S^H_{\tau-1}\right)^2 V^H_{\tau-1,T}, \quad \tilde{p} \triangleq \frac{\tilde{V}}{\tilde{V} + [2\tilde{R}_{\tau-1,\tau} - 1]\left[S^H_{\tau-1} - S_{\tau-1}\right]^2}. \quad (11)$$

Then:

1. $p^H_{\tau-1} \to \tilde{p}$ *as either* $[S^H_{\tau-1}]^2 V^H_{\tau-1,T} - [S^L_{\tau-1}]^2 V^L_{\tau-1,T} \to 0$ *or* $p^H_{\tau-1} \to 1$.
2. *If* $\tilde{V} > 0$*, then* $p^H_{\tau-1} \geq \tilde{p}$ *if and only if* $[S^L_{\tau-1}]^2 V^L_{\tau-1,T} \leq [S^H_{\tau-1}]^2 V^H_{\tau-1,T}$.
3. *If* $\tilde{V} < 0$*, then* $p^H_{\tau-1} < \tilde{p}$ *if and only if* $\tilde{p} > 0$.

Proof. See appendix. □

The proposition defines an estimator $\tilde{p}$ of $p^H_{\tau-1}$. The first result establishes that $\tilde{p}$ becomes an arbitrarily good approximation to $p^H_{\tau-1}$ as $[S^H_{\tau-1}]^2 V^H_{\tau-1,T} - [S^L_{\tau-1}]^2 V^L_{\tau-1,T} \to 0$ or as $p^H_{\tau-1} \to 1$. The intuition for the result tracks that already given for equation (9), with $[S^H_{\tau-1}]^2 V^H_{\tau-1,T} - [S^L_{\tau-1}]^2 V^L_{\tau-1,T} \to 0$ serving as the analogue of $\Delta Var \to 0$.

The second and third results sign the bias from estimating $p^H_{\tau-1}$ via $\tilde{p}$. The sign of $\tilde{V}$ plays a critical role, where $\tilde{V}$ is a metric that the econometrician can construct from options data 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 (9). In contrast, if $\tilde{V} < 0$, then the variance conditional on $k = H$ is greater than the variance of the compound lottery. The variance conditional on $k = H$ must therefore be substantially greater than the variance conditional on $k = L$.[25] The first part of the proposition therefore implies that the case with $\tilde{V} < 0$ is one in which $\tilde{p}$ is not a tight bound on $p^H_{\tau-1}$. The variance swap approach is therefore most useful when $\tilde{V} > 0$. From the second result of the proposition, $\tilde{p}$ is then a lower (upper) bound on $p^H_{\tau-1}$ if the post-event variance is smaller (larger) following $k = L$ than following $k = H$.

[25]Technically, if $\tilde{V}$ is only a bit less than zero (specifically, if $\tilde{V} \in (-[2\tilde{R}_{\tau-1,\tau} - 1][S^H_{\tau-1} - S_{\tau-1}]^2, 0)$), then $\tilde{p}$ is not informative about $p^H_{\tau-1}$. However, if $\tilde{V}$ is much less than zero (specifically, if $\tilde{V} < -[2\tilde{R}_{\tau-1,\tau} - 1][S^H_{\tau-1} - S_{\tau-1}]^2$), then the compound lottery must be mixing in the lottery conditional on L with some non-negligible probability.

3.3 Comparing the Two Estimators

We have derived two estimators of the risk-neutral probability of an event. Both estimators are fully nonparametric, in contrast to prior literature that uses options data to recover event probabilities by assuming that stock prices evolve according to specific parametric processes (described in the introduction). 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 volatility of the stock 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 volatility of post-event stock prices is not sensitive to the realization of the event. We now discuss the tradeoffs in the choice of estimator.

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, the econometrician 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 risk whose consequences are independent of the event realization. Because $\tilde{p}$ is not 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 $p^H_{\tau-1}$ from $\tilde{p}$ when firms are harmed by an event.

4 Empirical Approach

We now describe our empirical approach to estimating $\bar{p}$ and $\tilde{p}$ from observed option prices. 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.

In order to avoid biases from firms with non-event-related news, our preferred specifications limit the sample to firms that do not have an earnings announcement in the event window and control for earnings announcements that occur elsewhere in the estimation window.[26] We do not control for the market index because we analyze some big events that may have affected that index (so that controlling for the index could accidentally absorb some of the desired event effect). Following, among others, Dubinsky et al. (2019), our preferred specifications also drop firms whose stock price falls below \$5 at any point in either the estimation or event windows and drop firms with a quarterly dividend yield greater than 2% over the estimation window. We next describe additional, theoretically motivated restrictions designed to recover tight bounds on the priced-in event probability.

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

The objective is to estimate $p^H_{\tau-1}$ by obtaining a tight bound $\bar{p}$, where H again stands for the realized event. From equation (5), the bias $\bar{p} - p^H_{\tau-1}$ depends on $C^L_{\tau-1,T}(S^L_{\tau-1}, K)$. If we could identify options for which $C^L_{\tau-1,T}(S^L_{\tau-1}, K)$ were small, then we could empirically estimate the priced-in probability of an event from the following regression:[27]

$$ln\left(\frac{C_{iKe(t-1)}}{C_{iKet}}\right) = \alpha_{iKe} + \beta Event_t + \theta_{iKe} X_{iKet} + \varepsilon_{iKet}, \tag{12}$$

where we change notation on the call option price, letting i index firms, K index strike prices, e index expiration dates, and t index trading dates. An analogous regression holds in the case where we should instead examine prices of put options. $Event_t$ is a dummy variable for the event occurring on trading date t. The vector X_{iKet} includes the controls. We then estimate $\bar{p}$ by predicting $C_{iKe(\tau-1)}/\hat{C}^H_{iKe(\tau-1)}$. We do not include $\hat{\alpha}$ or $\hat{\theta}$ when predicting $C_{iKe(\tau-1)}/\hat{C}^H_{iKe(\tau-1)}$ because we do not want to include underlying trends or observable shocks in $\hat{C}^H_{iKe(\tau-1)}$. We therefore predict $\ln(C_{iKe(\tau-1)}/\hat{C}^H_{iKe(\tau-1)})$ from $\hat{\beta}$ alone, comparing the option price on the day

[26]Previous work has shown that options prices respond to earnings announcements (e.g., Patell and Wolfson, 1979, 1981; Dubinsky et al., 2019). We control for earnings announcements by using a dummy variable for each day in a three-day window centered around the announcement.

[27]We use a log on the left-hand side because we find that log-changes in option prices are approximately normally distributed.

before the event to what that option price would have been if the event outcome had been known but nothing else had changed. We thus have $\bar{p} = \exp(\hat{\beta})$.[28]

Thus far we have assumed that the empirical researcher can identify those options for which $C^L_{\tau-1,T}(S^L_{\tau-1}, K)$ is small, but in fact $C^L_{\tau-1,T}(S^L_{\tau-1}, K)$ is unobservable. We now describe how empirical researchers can estimate a tight bound on the event probability without knowledge of $C^L_{\tau-1,T}(S^L_{\tau-1}, K)$. To achieve this, we use theoretically motivated insights about how the bias from ignoring $C^L_{\tau-1,T}(S^L_{\tau-1}, K)$ 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 depend on the risk of disasters. Our preferred specifications therefore limit the sample to the deepest out-of-the-money option for each firm-expiration pair that has nonzero volume on each day surrounding the event and also has positive bids on each day surrounding the event. To address concerns about the accuracy of option prices at strikes far from the money, our preferred specifications also weight equation (12) by the inverse of the average relative bid-ask spread for each option, with the average taken over days τ and $\tau - 1$.[29]

In addition, our preferred specifications limit the sample for the regression in equation (12) to the nearest major expiration date.[30] We do this for two reasons. First, we developed the theory in Section 3.1 for European-style options, but many options are American-style. The appendix shows that the potential error from misrepresenting American-style options as European-style options vanishes as the time to maturity decreases. Second, the value of deep out-of-the-money options generally increases in their time to maturity because more extreme stock prices become more likely over longer intervals. Therefore, for a given set of strikes, options with nearer expiration dates are likely to have smaller $C^L_{\tau-1,T}(S^L_{\tau-1}, K)$.

Next, conditional on the above considerations, options on stocks with relatively large stock price responses to the event will generate tighter bounds on the event

[28]We do not let β vary across firms because the event probability should not vary across firms in our applications.

[29]Thus the weight on each observation is the inverse of $\frac{1/2(A_t - B_t) + 1/2(A_{t-1} - B_{t-1})}{1/2(D_{t+1} + D_t)}$, where A_t is the ask price for the option on trading date t, B_t is the bid price for the option on trading date t, and $D_t \triangleq \frac{1}{2}(A_t + B_t)$. Relative bid-ask spread is a standard measure of liquidity. For instance, see Madhavan (2000) and Vayanos and Wang (2013).

[30]Options in our data overwhelmingly expire on the third Friday of the month. There are some options that expire on other dates within the month, but we focus our analysis on the major expiration dates because the other expiration dates are less liquid.

probability because their stock price distributions will be more affected by conditioning on the observed outcome versus the counterfactual outcome. We therefore also run a traditional event study regression for each firm:

$$ln(S_{it}/S_{i(t-1)}) = \gamma_{i1} + \gamma_{i2} Event_t + \gamma_{i3} X_{it} + \varepsilon_{it}, \tag{13}$$

where S_{it} is the underlying asset price for firm i on trading date t (so the left-hand side is the log daily return on the asset) and X_{it} again includes the earnings report controls. Our preferred specifications focus on firms that have large event-day jumps in asset prices. We measure "large" by the t-statistic on $\hat{\gamma}_{i2}$ in order to identify firms whose event-day stock price jumps are large relative to their underlying volatility.

Finally, our preferred specifications use call options and not put options because, as described in Section 3.1, put options may generate more bias than call options. Regression (13) tells us whether to estimate $\bar{p}$ from call options or put options: 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 $\hat{\gamma}_{i2} > 0$.

These considerations give us guidance about how to construct the estimation sample for regression (12). When estimating the priced-in probability of an event, we will be particularly interested in estimates of β from a sample of call options that are well-traded and deep out-of-the-money, expire at the nearest major expiration date, and are written on firms that have large, positive t-statistics on $\hat{\gamma}_{i2}$.

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 and expiration date. Equation (11) shows that to calculate $\tilde{p}$ we need information on the underlying asset price on the day before the event ($S_{\tau-1}$) as well as the counterfactual value of that asset if the event outcome were already known ($S^H_{\tau-1}$). $S_{\tau-1}$ is observed in the data and $S^H_{\tau-1}$ is straightforward to recover from the conventional event study in equation (13).

We also need the variance swap rate on the day before the event ($V_{\tau-1,T}$) and the counterfactual variance swap rate if the event's outcome were already known ($V^H_{\tau-1,T}$). First consider $V_{\tau-1,T}$. Equation (10), from Martin (2017), shows that $V_{\tau-1,T}$ can be calculated by integrating over the observed prices of put and call options. We use a daily version of the 3-month LIBOR rate to calculate $\tilde{R}_{\tau-1,\tau}$ and $R_{\tau-1,T}$. We replace $\tilde{R}_{\tau-1,T} S_{\tau-1}$ with the forward price, which is the strike at which the observed prices of call and put options are equalized. We discretize the integral following the methodology used in constructing the VIX index.[31]

We calculate the counterfactual variance swap rate $V^H_{\tau-1,T}$ in a nearly identical way. The difference is that, in place of the observed prices of options, we use the counterfactual options prices predicted from the event study in options. We predict these counterfactual option prices from equation (12), except that now we allow the intercepts, α_{iKe}, and the event effects, β_{iKe}, to vary by firm, expiration, and strike.

As before, our preferred specifications restrict the sample in order to emphasize the variance swaps that are most likely to be informative about the true market probability of the event. First, some variance swaps are based on less liquid options than others. The VIX methodology does drop sufficiently illiquid options when calculating the integral. Further, we weight firms by the inverse of the relative bid-ask spread, averaged over the strikes available for that combination of firm and expiration date.

Second, the identifying assumption is that the expected variance of the stock prices conditional on the realized event outcome is similar to the expected variance conditional on the counterfactual event outcome. As discussed at the end of Section 3.2, one way that this assumption would clearly be violated is if $\tilde{V} < 0$. Our preferred specifications therefore limit the sample to only those firms and expiration dates for which $\tilde{V} > 0$.

Third, our preferred specifications again limit the sample to the nearest major expiration date. This restriction is motivated by a desire to increase the signal-to-noise ratio of the day-ahead variance swap relative to the T-day-ahead variance swap synthesized from observed option prices. It also accounts for observing American-style options yet applying results from Martin (2017) derived for European-style options. The error from observing American-style options is further reduced by dropping firms with high dividend rates.

Finally, the estimate of $\tilde{p}$ is going to be better if the asset is particularly affected by the event, because this will mean that the amount of variance that is resolved by the event is large relative to the general variance in asset prices. Our preferred specifications therefore again limit the sample to those firms with large t-statistics on $\hat{\gamma}_{i2}$ from equation (13). The t-statistic takes into account both the size of the stock price jump on the day of the event and the typical variation in a stock's price over the estimation window. In contrast to Section 4.1, our preferred specifications no longer restrict the sign of $\hat{\gamma}_{i2}$.

[31] `http://www.cboe.com/products/vix-index-volatility/vix-options-and-futures/vix-index/the-vix-index-calculation`.

5 Empirical Application: Recovering the Probability of Health Care Events

We now both test and demonstrate our methods by applying them to four important events in recent health care policy. The Patient Protection and Affordable Care Act (ACA, or "Obamacare") was a cornerstone policy of the Obama presidency. But the path to the ACA in its current form was complicated and had numerous points at which it could have been altered or scrapped. Al-Ississ and Miller (2013) study one such event in the history of the ACA: the special election of Republican Scott Brown to the U.S. Senate from Massachusetts in January of 2010. They focus on identifying which health care sectors were helped or harmed by the ACA. They account for the increasing likelihood of Brown's victory over the week leading up to the vote by including more pre-event days in their event window. We will instead directly estimate the probability of a Brown victory and use it to rescale conventional event study estimates.

We examine four events that were considered critical for the ACA: the January 19, 2010 Brown election from Al-Ississ and Miller (2013), the June 28, 2012 Supreme Court decision that upheld the individual mandate, the November 8, 2016 Presidential election,[32] and the July 28, 2017 vote in the Senate on the "skinny repeal" of the ACA that is best known for Senator John McCain's "thumbs down" vote. This series of events was critical to the evolution of the ACA and to the durability of the ACA. Given that the content of the ACA and its implications were potentially changing over time, it is interesting to know whether industries were affected similarly by the different events and which events were especially important. In all cases, we study the sectors and firms identified by Al-Ississ and Miller (2013). Importantly, as discussed below, these events varied widely in the amount of information about their likelihood that was available beforehand and in estimates of that likelihood. These diverse applications therefore test whether the probabilities recovered by our methods vary in a reasonable way across events.

5.1 Data

We use standard data sources for our analysis. Stock prices, dividends, and earnings dates come from Compustat, and options data come from OptionMetrics. We construct closing prices by averaging the closing bid and ask prices. Our estimation

[32] The November 2016 election elected President Donald Trump and also gave Republicans control of both houses of Congress. Because the former result may have been extremely unlikely to occur without the latter result, we refer to the event as the presidential election.

window begins 80 days before the event and ends either 60 days after the event or 7 trading days before the option expires, whichever comes first.[33]

When calculating the value of a variance swap, we need to determine which options to use for each firm and expiration date. We follow the approach used to calculate the VIX: starting at the money and using calls (puts) as we increase (decrease) strikes away from the money, we include every option with a strictly positive bid until we encounter two options in a row with zero bids. We do not include any options that are further from the money than this last option, whether or not they have strictly positive bids.

5.2 Traditional Event Studies

We begin by conducting traditional event studies for each of the four events. In each case, we follow Al-Ississ and Miller (2013) and look at the effect of the event on each of three healthcare industries: facilities (e.g., hospitals), managed care firms (e.g., insurance companies), and pharmaceutical companies. For each of our event studies, we estimate versions of equation (13) where we allow the event dummy to vary only over these three industries and where we examine effects on the 2-day return $ln(S_t/S_{t-2})$ to account for the possibility that information about the event outcome "leaks" on the day of the event. Standard errors are clustered at the firm level to account for arbitrary autocorrelation in stock returns.

Table 1 reports the results of these event studies. Keep in mind that the Brown and Trump elections increased barriers to the passage or existence of the ACA, whereas the Supreme Court decision upholding the individual mandate (SCOTUS) and the McCain Senate vote (McCain) increased the likelihood that the ACA would continue. That said, each of these events had slightly different implications for the ACA. The Brown election occurred before any bill had been passed, and so the ACA as it was thought of then looked somewhat different than it did by the time of President Trump's election. We simply compare the relative impact of the events on each of the industries because our emphasis here is not on the vagaries of each event's effect on various firms but on our ability to recover the priced-in probability of the event.

For the Trump election, we find large and statistically significant jumps in the value of facility and pharmaceutical company stocks at the election and a smaller and statistically insignificant change in the value of managed care stocks. Facility

[33]Following Beber and Brandt (2006) and Kelly et al. (2016), we exclude options within seven days of maturity because of concerns about their liquidity. The CBOE also drops these options when calculating the VIX.

stocks lost an average of nearly 6.4 percent of their value and pharmaceutical stocks gained 4.9 percent of their value at the election. These are by far the largest jumps in facility and pharmaceutical stock values of the four events that we look at.

Following Al-Ississ and Miller (2013), the event study results for the Brown election are generally smaller and less statistically significant. Managed care and pharmaceutical stocks gained 2.2% and 1% of their value on average, although the effects are only statistically significant at the 10% level.

For the Supreme Court decision, only facility stock prices are statistically significantly affected. However, that effect is large, with facility stocks returning 4.2% on average. This is consistent with the fact that the Supreme Court decision upheld the individual mandate to purchase health insurance, which should both increase demand for health care and increase the portion of demand that is reimbursed by insurance companies. The small response from managed care firms is consistent with their response to the Trump election, and the small response of pharmaceutical firms may combine the negative effect of the ACA seen in the Trump election with the positive effect of increased insurance coverage.

Finally, the Senate vote to *not* rescind the ACA had a negative effect for both facilities (1.2%) and pharmaceuticals (1.7%). Interestingly, the latter effect is consistent with the effect of electing President Trump but the former is not. The difference could reflect the presence of a concrete repeal plan for the ACA. The effect on managed care firms was statistically insignificant, which is consistent with the effect of electing President Trump.

These results suggest that the Trump election was the most important event for health care stocks, but the Trump election was also an especially unlikely event. We compare the full effects of the events once we recover the priced-in event probabilities.

5.3 Event Probabilities Estimated From Out-of-the-Money Options

We now run regression (12) to estimate $\bar{p}$, the market's priced-in probability of each event. As in the conventional stock event study, we compare option prices on the day before the event to the day after the event in order to avoid having information about the event outcome leaking on the event day and contaminating the results. We present the estimated probabilities (with 95% confidence intervals) for each event in Tables 2 through 5, introducing restrictions on the sample as we move from left to right within each table.

We begin with the most extensively polled event, the 2016 U.S. election (Table 2). The first column imposes the fewest restrictions. It selects neither the firms nor the options that are likely to generate the tightest bounds: it runs the regression for all facilities, managed care firms, and pharmaceutical firms, using all call options for firms with a positive event study effect of the Trump election and all put options for firms with a negative event study effect. Unsurprisingly, this wide range of options (which includes even in-the-money options) yields a very loose upper bound on the event probability: we estimate a $\bar{p}$ of 0.91, well above any reasonable polling data from before the election. In column 2, we reweight observations by the inverse of the relative bid-ask spread. Downweighting the observations that may be relatively thinly traded (and therefore potentially mispriced) reduces the estimated $\bar{p}$ substantially, to 0.78. Column 3 further focuses on well-traded options by removing any option with a bid of zero or with zero volume, and column 4 applies standard restrictions from the empirical finance literature. The estimated $\bar{p}$ is still 0.78.

The remaining columns introduce restrictions designed to yield a tighter bound on $p^H_{\tau-1}$. Column 5 restricts the sample to the options with the nearest major expiration date. This restriction reduces the estimated probability of a Trump win to 0.69. Column 6 restricts the sample to the deepest out-of-the-money options that are sufficiently liquid. As predicted in Section 3.1, this restriction dramatically tightens the bound, reducing the estimated probability to 0.228. Column 7 drops all put options in the sample, motivated by the argument in Section 3.1 that put options are likely to generate a weaker bound. This restriction actually leads to a slightly higher estimate of the event probability in this application, but we will see that this restriction does reduce the estimated probability in the other applications. Finally, columns 8 and 9 restrict the sample to only those firms with statistically large event-day stock price jumps. Column 8 restricts the sample to firms with event-study t-statistics above 1.96, and column 9 restricts the sample to firms with event study t-statistics above 5. In either case, the estimated probability of a Trump election with this fully restricted estimation sample is 0.222. Because it uses all of the restrictions motivated in Section 4, this probability is our preferred estimate.

Compare this estimate to prominent polling-driven estimates from before the election. Fivethirtyeight.com put the probability of a Trump victory around 0.28 on the day before the election. The New York Times put the probability of a Trump victory between 0.01 and 0.15 in the days leading up to the election. Our preferred estimate is between the probabilities calculated by fivethirtyeight.com and the New York Times.[34]

[34]Note that we estimate a risk-neutral probability rather than the physical probability of interest to journalists. The difference between the two could go in either direction. The risk-neutral probability is the one needed to correct event study estimates. Prediction markets should reflect risk-neutral probabilities. Also, note that the probability of a Trump election along with a Republi-

Now compare our estimate to the probabilities implied by prediction market contracts.[35] On the morning of the election, Betfair gave Trump a 22% chance of winning and PredictIt gave him a 22% chance of winning. The bookmaker Paddy Power's odds implied that Trump had a 22% chance of winning, and the bookmaker Ladbroke's odds gave him a 24% chance of winning. These probabilities are remarkably consistent with our estimate of 22%.

Table 3 repeats this estimation exercise for the Brown election studied by Al-Ississ and Miller (2013). Here the estimate of Brown winning the election when using the full sample is quite large, with the 95% confidence interval even reaching above 1. However, as we limit the data to options that are less likely to be biased upwards by value in the event of a Brown loss, we again see that the estimates fall substantially and even appear to converge. The biggest decline in the probability estimate again comes from restricting our analysis to only each firm's deepest out-of-the-money, liquid options. However, this event also shows larger declines in the estimate when we remove put options and when we restrict attention to firms that had significant movements in the traditional event study. Given that the Brown election may have been a significant event for fewer firms than was the Trump election, it makes sense that these last restrictions now tighten our bound to a greater degree.

The most important difference between the Brown election and the Trump election is that the preferred probability estimate ends up in the neighborhood of 0.54 rather than the 0.22 found for the Trump election. The higher estimated probability for the Brown election is in keeping with the data that was available before the election. The polls had been trending toward Brown in the last week before the election. A poll only six days before the election had his opponent ahead, but another poll only two days before the election predicted a tie. Brown eventually did win, by 4.7 percentage points. The election was held on the Tuesday after Martin Luther King Jr. day, so the last trading day on the market was the Friday before the election. Al-Ississ and Miller (2013) report that an Intrade prediction market contract placed the probability of Brown winning at 0.30 on Friday January 15 but suggested a probability of 0.77 by the evening of Monday January 18. Our estimated probability should be thought of as the risk-neutral probability of Brown winning as of the Friday before the election. Our estimated probability of 0.54 therefore appears to lead the change in Intrade.com's odds that would occur over the coming weekend.

can sweep of the Senate and House may have been smaller—but probably not much smaller—than the probability of a Trump victory.

[35]The probabilities given below are as reported at `https://www.livemint.com/Politics/Aw7vJFLufs5QN2DHvcO26M/Hillary-Clinton-far-ahead-of-Donald-Trump-on-betting-exchang.html`.

Tables 4 and 5 present the estimated probabilities for the 2012 Supreme Court ruling and for the 2017 Senate vote on the skinny repeal. In both cases, the priced-in probability is very high in the full data set but declines substantially with the sample restrictions, to 0.759 for the Supreme Court decision and 0.814 for the Senate vote. Note, however, that the 95% confidence intervals for both of these events with the restricted sample are large, even including a probability of 1. Our estimates suggest that the realized outcomes were more likely than not, but our preferred specifications do have the drawback of not allowing for particularly precise estimates because they limit the number of firms in the sample.

Using a parametric model of option prices, Borochin and Golec (2016) estimate the probability of the Supreme Court ruling as 0.68, roughly in line with our estimates despite the significant differences in methodology and assumptions between the two studies. They report that the Intrade probability was only 0.30; however, they also report that investment analysts judged the probability to be 0.50 or greater. They hypothesize that the Supreme Court may be especially opaque to outsiders. Given the coherence between the estimates produced by the two very different econometric approaches and between those econometric estimates and analysts' estimates, and given that these estimates do perform better than the prediction market contracts ex post, we may reconsider the performance of prediction market contracts in similar settings with only limited public information.[36]

Using prediction markets to correct event study estimates also faces another hurdle: when prediction market prices are changing rapidly, it is unclear which point in time to use for the desired probability. We have already discussed that the prediction market price was changing dramatically in the lead-up to Scott Brown's election. The lead-up to the Senate ACA repeal vote was even more dramatic. PredictIt gave the Senate health care bill an 81% chance of failing as of midnight Eastern on July 25, a 57% chance of failing as of midnight Eastern on July 26, and only a 15% chance of failing as of midnight Eastern on July 27, less than two hours before the vote.[37] The probability fluctuated a lot in between these times, as news events unfolded at a rapid rate. For instance, the bill seemed doomed after a press conference held around 5 PM by two Republican Senators, but the Speaker of the House of Representatives offered a compromise around 8 PM that substantially improved the bill's chances.[38]

[36]Even if prediction markets were to generate probabilities that are more "correct" than those from the market, the market probabilities are the ones needed to correct event study estimates.

[37]The prediction market contract is for the Senate passing the bill by July 31. The reported probabilities are from `https://www.predictit.org/markets/detail/3413/Will-the-Senate-pass-the-Better-Care-Reconciliation-Act-by-July-31`. The time that defines the prediction market's close is from personal communication.

[38]`https://www.rollcall.com/politics/highlights-senates-long-vote-rama-night`

To correct event study estimates, we require the probability that had been priced-in by financial markets as of their close on July 27, which is 4 PM Eastern for many markets of interest. Our estimates suggest that financial markets had priced in an 81% chance of the Senate vote failing by their close on July 27. This probability is not inconsistent with that afternoon's fluctuations in news but would be difficult to isolate from prediction market movements.

5.4 Event Probabilities Estimated From Variance Swaps

Our second approach to recovering the priced-in event probability uses the full range of strikes for each firm. Table 6 presents the probability estimates and standard errors for each of the four events. For each event we begin with the estimate of the event probability from an unrestricted set of firms and then restrict the sign of $\tilde{V}$, the expiration date of the options, and the event study t-statistic. All estimates are weighted by the inverse of the average relative bid-ask spread, as described in Section 4.2. Once we implement the full restrictions, the event probabilities are very similar to the estimates from Section 5.3. For the Trump election, the estimated probability of 0.215 is extremely close to the 0.222 estimated previously. For the Brown election, the 0.58 estimated here is slightly higher than the 0.537 estimated before, but the two estimates are still well within a single standard error of each other. For both of these events the 95% confidence intervals (here presented as standard errors) are somewhat larger here than in the out-of-the-money option approach.

The second panel of Table 6 shows the estimates for the Supreme Court and Senate events using this approach. The biggest difference from Section 5.3 appears in the Supreme Court event. Here we estimate a probability of 0.568 that the Supreme Court would uphold the ACA, whereas we estimated a probability of 0.759 in Section 5.3. However, this is the event in which a key assumption underlying the variance swap approach is unlikely to hold: the precise date of the Supreme Court ruling was probably not known in advance. In that case, the event is resolving uncertainty about the court's ruling as well as about the event date. This creates two problems. First, the counterfactual event outcome now combines the chances of other court rulings on this date with a chance of the ruling happening at a later date. Because near-term stock prices would be especially volatile in the case of a later ruling (due to the fact that the ruling would cause a jump once it occurred), the assumption that the post-event volatility of stock prices does not depend on the event is violated. By Proposition 1, we should then be estimating an upper bound on $p^H_{\tau-1}$. Second, the relevant event probability is now the joint probability of the ruling happening on that date and the probability of the ruling being what it was. We therefore have to inflate the estimated event probability to account for uncertainty about the date that the ruling would occur. Under the reasonable assumption that the ruling would have had to occur either that day or the next day (a Friday), this logic suggests doubling the estimated probability. Combining these conclusions, the estimated probability of 0.57 may therefore be an upper bound on a probability that should itself be doubled.[39]

Finally, the estimated probability of the skinny repeal of the ACA failing is now 0.847, whereas it was 0.814 in Section 5.3. While these are quite close, the 95% confidence interval is substantially smaller under the present approach. The smaller confidence interval reflects that we are using the full distribution of strikes rather than only using one option per firm. However, the variance swap approach yields larger standard errors for the other events. We are still working on developing proper standard errors for the variance swap method, so these results may change.

5.5 Comparison of Different Events' Effects

With the estimated priced-in probabilities of each event in hand, we now use equation (2) to rescale the original event study estimates to get a sense of the full value of the event to each industry's firms (Table 7). The results lead to striking reorderings in the relative importance of the different events. In particular, whereas Trump's election appeared to be the most important event for both facilities and pharmaceuticals in the traditional event studies, the rescaling demonstrates that this was largely because Trump's victory was relatively unexpected (equivalently, stocks had largely already priced in a victory by Hillary Clinton). In particular, the Supreme Court ruling now appears to be the most important event for many firms. This finding is consistent with the widespread belief that the ACA could stand or fall on this ruling and that the individual mandate would have a substantial impact on the overall demand for health care, which would be particularly important for health care facilities.

Additionally, Table 7 shows that the failure of the "skinny repeal" (labeled "McCain") was especially important for health care firms, even though the original event study magnitudes are relatively small. In fact, this event had the largest adjusted effect on pharmaceutical stocks, reducing their value by over 9% on average. This again displays the importance of understanding the probability that the market placed on an event before it happened, particularly for events that were judged to be relatively likely and therefore had their effects largely priced-in ahead of time.

[39] In ongoing work, we will explore ways of accounting for uncertainty about the event date in the variance swap approach.

6 Conclusions

We have developed two new approaches to estimating event probabilities from options markets. In contrast to prior literature, our approaches impose no parametric assumptions about the distribution of stock prices. Both approaches boil down to running event studies in option prices to complement the standard event study 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 remarkably consistent with probabilities implied by bookmakers and prediction markets and in the middle of polling-driven predictions. Further, we estimate a very different probability for a second election, and this probability is again consistent with polling data and prediction markets. Finally, we also apply our methods to two types of events that have historically challenged event studies because they are often well-anticipated: a court ruling and a legislative vote. In those cases, we again derive reasonable probabilities, in one case closer to the probabilities posited by investment analysts than to the probabilities implied by prediction markets and in the other case consistent with the broad range of probabilities implied by prediction market fluctuations near the event.

Our new methods come with at least 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. Our estimated are useful in these cases, but they are only part of the adjustment required to correct the conventional event study estimates. Second, some event studies consider whether there was an effect at all, but they have been unable to discriminate between the true effect being small and the effect being well-anticipated. Our methods can in principle recover probabilities in cases with very small observed event effects, but in practice the signal-to-noise ratio in option prices might be low in such cases.

Future researchers could combine our estimated probabilities of these health care events with variation in the events' details to study the market valuation of particular aspects of the ACA. Further, our estimated probabilities allow these market effects to be rescaled for use in cost-benefit analyses of health regulations. In the introduction, we described many more applications of economic research that could benefit from estimated market probabilities of events. We have seen that it can be easy to recover these probabilities. We encourage the development of this agendas hitherto hindered by ignorance of market probabilities.

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.

Ball, Ray (1972) "Changes in accounting techniques and stock prices," *Journal of Accounting Research*, Vol. 10, pp. 1–38.

Barraclough, Kathryn, David T. Robinson, Tom Smith, and Robert E. Whaley (2013) "Using option prices to infer overpayments and synergies in M&A transactions," *Review of Financial Studies*, Vol. 26, No. 3, pp. 695–722.

Barro, Robert J. (2006) "Rare disasters and asset markets in the twentieth century," *The Quarterly Journal of Economics*, Vol. 121, No. 3, pp. 823–866.

Barro, Robert J. and Jos F. Ursa (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 I: Technique and corporate litigation," *American Law and Economics Review*, Vol. 4, No. 1, pp. 141–168.

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 elelctions 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.

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. forthcoming.

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.

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.

Hallock, Kevin F. (1998) "Layoffs, top executive pay, and firm performance," *The American Economic Review*, Vol. 88, No. 4, pp. 711–723.

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. 45, 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. 23, 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.

Madhavan, Ananth (2000) "Market microstructure: A survey," *Journal of Financial Markets*, Vol. 3, No. 3, pp. 205–258.

Malatesta, Paul H. and Rex Thompson (1985) "Partially anticipated events: A model of stock price reactions with an application to corporate acquisitions," *Journal of Financial Economics*, Vol. 14, No. 2, pp. 237–250.

Martin, Ian (2017) "What is the expected return on the market?" *The Quarterly Journal of Economics*, Vol. 132, No. 1, pp. 367–433.

Melick, William R. and Charles P. Thomas (1997) "Recovering an asset's implied PDF from option prices: An application to crude oil during the Gulf crisis," *Journal of Financial and Quantitative Analysis*, Vol. 32, No. 1, pp. 91–115.

Meng, Kyle C. (2017) "Using a free permit rule to forecast the marginal abatement cost of proposed climate policy," *American Economic Review*, Vol. 107, No. 3, pp. 748–784.

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.

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.

Vayanos, Dimitri and Jiang Wang (2013) "Market liquidity—Theory and empirical evidence," in George M. Constantinides, Milton Harris, and Rene M. Stulz eds. *Handbook of the Economics of Finance*, Vol. 2: Elsevier, pp. 1289–1361.

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.

Table 1: Event Study Magnitude Estimates

Event	Trump	Brown	SCOTUS	McCain
Event if Facilities	-0.0638	-0.0107	0.0420	-0.0123
	(0.0265)	(0.0125)	(0.0083)	(0.0056)
Event if Managed Care	-0.0182	0.0225	0.0107	-0.0050
	(0.0265)	(0.0125)	(0.0083)	(0.0056)
Event if Pharmaceutical	0.0491	0.0095	0.0053	-0.0174
	(0.0088)	(0.0053)	(0.0056)	(0.0062)
R^2	0.1284	0.0525	0.0842	0.0641
N	6,005	9,398	7,357	7,321

Table 2: Trump Event health Industry Probability Estimates

	1	2	3	4	5	6	7	8	9
Event Probability	0.9077	0.7783	0.7901	0.7776	0.7110	0.2227	0.2253	0.2218	0.2218
95% CI Lower Bound	0.90	0.77	0.79	0.77	0.70	0.19	0.18	0.18	0.18
95% CI Upper Bound	0.92	0.78	0.79	0.78	0.72	0.26	0.28	0.27	0.27
Spread Ratio Weights	N	Y	Y	Y	Y	Y	Y	Y	Y
Bid > 0 Only	N	N	Y	Y	Y	Y	Y	Y	Y
Volume > 0 Only	N	N	Y	Y	Y	Y	Y	Y	Y
Earnings Announcements	Y	Y	Y	N	N	N	N	N	N
High Dividend Stocks	Y	Y	Y	N	N	N	N	N	N
Nearest Exp. Only	N	N	N	N	Y	Y	Y	Y	Y
Extreme Strikes Only	N	N	N	N	N	Y	Y	Y	Y
Event > 0 Only	N	N	N	N	N	N	Y	Y	Y
Event t-stat > 1.96	N	N	N	N	N	N	N	Y	Y
Event t-stat > 5	N	N	N	N	N	N	N	N	Y
R^2	0.0130	0.0220	0.0244	0.0278	0.0451	0.2861	0.2713	0.2717	0.2717
N	792,496	792,496	530,802	419,364	75,780	1,393	988	979	979

Table 3: Brown Event health Industry Probability Estimates

	1	2	3	4	5	6	7	8	9
Event Probability	1.0222	0.9520	0.9523	0.9502	0.9436	0.6951	0.6114	0.5983	0.5365
95% CI Lower Bound	1.01	0.94	0.94	0.94	0.93	0.64	0.55	0.54	0.48
95% CI Upper Bound	1.04	0.96	0.96	0.96	0.96	0.76	0.68	0.66	0.60
Spread Ratio Weights	N	Y	Y	Y	Y	Y	Y	Y	Y
Bid > 0 Only	N	N	Y	Y	Y	Y	Y	Y	Y
Volume > 0 Only	N	N	Y	Y	Y	Y	Y	Y	Y
Earnings Announcements	Y	Y	Y	N	N	N	N	N	N
High Dividend Stocks	Y	Y	Y	N	N	N	N	N	N
Nearest Exp. Only	N	N	N	N	Y	Y	Y	Y	Y
Extreme Strikes Only	N	N	N	N	N	Y	Y	Y	Y
Event > 0 Only	N	N	N	N	N	N	Y	Y	Y
Event t-stat > 1.96	N	N	N	N	N	N	N	Y	Y
Event t-stat > 5	N	N	N	N	N	N	N	N	Y
R^2	0.0059	0.0116	0.0132	0.0132	0.0192	0.1701	0.1669	0.1720	0.2322
N	162,774	162,774	115,248	109,690	46,370	1,923	1,187	1,129	837

Table 4: SCOTUS Event health Industry Probability Estimates

	1	2	3	4	5	6	7	8	9
Event Probability	1.0130	1.0258	1.0288	1.0292	1.0298	0.8524	0.8114	0.7796	0.7587
95% CI Lower Bound	1.00	1.02	1.02	1.02	1.02	0.69	0.63	0.60	0.56
95% CI Upper Bound	1.02	1.03	1.04	1.04	1.04	1.05	1.04	1.01	1.02
Spread Ratio Weights	N	Y	Y	Y	Y	Y	Y	Y	Y
Bid > 0 Only	N	N	Y	Y	Y	Y	Y	Y	Y
Volume > 0 Only	N	N	Y	Y	Y	Y	Y	Y	Y
Earnings Announcements	Y	Y	Y	N	N	N	N	N	N
High Dividend Stocks	Y	Y	Y	N	N	N	N	N	N
Nearest Exp. Only	N	N	N	N	Y	Y	Y	Y	Y
Extreme Strikes Only	N	N	N	N	N	Y	Y	Y	Y
Event > 0 Only	N	N	N	N	N	N	Y	Y	Y
Event t-stat > 1.96	N	N	N	N	N	N	N	Y	Y
Event t-stat > 5	N	N	N	N	N	N	N	N	Y
R^2	0.0087	0.0296	0.0285	0.0280	0.0257	0.1816	0.1818	0.1946	0.1971
N	337,674	337,674	220,510	211,701	72,820	1,687	1,223	1,085	866

Table 5: McCain Event health Industry Probability Estimates

	1	2	3	4	5	6	7	8	9
Event Probability	1.0820	1.0020	0.9987	1.0003	0.9965	0.8341	0.8259	0.8206	0.8139
95% CI Lower Bound	1.07	1.00	0.99	1.00	0.99	0.71	0.68	0.67	0.65
95% CI Upper Bound	1.09	1.01	1.00	1.00	1.01	0.98	1.01	1.01	1.02
Spread Ratio Weights	N	Y	Y	Y	Y	Y	Y	Y	Y
Bid > 0 Only	N	N	Y	Y	Y	Y	Y	Y	Y
Volume > 0 Only	N	N	Y	Y	Y	Y	Y	Y	Y
Earnings Announcements	Y	Y	Y	N	N	N	N	N	N
High Dividend Stocks	Y	Y	Y	N	N	N	N	N	N
Nearest Exp. Only	N	N	N	N	Y	Y	Y	Y	Y
Extreme Strikes Only	N	N	N	N	N	Y	Y	Y	Y
Event > 0 Only	N	N	N	N	N	N	Y	Y	Y
Event t-stat > 1.96	N	N	N	N	N	N	N	Y	Y
Event t-stat > 5	N	N	N	N	N	N	N	N	Y
R^2	0.0132	0.0706	0.0539	0.0566	0.0636	0.3769	0.3958	0.4004	0.4010
N	741,059	741,059	486,327	434,800	86,825	1,624	1,213	1,172	981

Table 6: Probability Estimates from the Variance Swap Approach

Event	Trump				Brown		
Probability	0.7547	0.3262	0.3146	0.2153	0.5476	0.8007	0.7721	0.5829
	(0.6909)	(0.0321)	(0.0734)	(0.0508)	(0.2393)	(0.0379)	(0.0521)	(0.0740)
Positive $\tilde{V}$	N	Y	Y	Y	N	Y	Y	Y
Nearest Expiration	N	N	Y	Y	N	N	Y	Y
Event Study t-stat> 5	N	N	N	Y	N	N	N	Y
Number of Firms	169	88	21	19	73	50	27	14
Event		SCOTUS			McCain			
Probability	0.6296	0.6898	0.5912	0.5676	0.9948	0.9588	0.8562	0.8472
	(0.1385)	(0.0266)	(0.0480)	(0.0522)	(0.0429)	(0.0099)	(0.0378)	(0.0494)
Positive $\tilde{V}$	N	Y	Y	Y	N	Y	Y	Y
Nearest Expiration	N	N	Y	Y	N	N	Y	Y
Event Study t-stat> 5	N	N	N	Y	N	N	N	Y
Number of Firms	147	111	43	28	196	149	30	22

Table 7: Comparison of Full Event Effect Across Events

Event	Trump	Brown	SCOTUS	McCain
Estimated Event Study Magnitudes				
Event if Facilities	-0.0638	-0.0107	0.0420	-0.0123
Event if Managed Care	-0.0182	0.0225	0.0107	-0.0050
Event if Pharmaceutical	0.0491	0.0095	0.0053	-0.0174
Predicted Probability	0.2218	0.5365	0.7587	0.8139
Adjusted Effect				
Facilities	-0.0820	-0.0231	0.1741	-0.0661
Managed Care	-0.0234	0.0485	0.0443	-0.0269
Pharmaceutical	0.0631	0.0205	0.0220	-0.0935

Table 8: Trump Event energy Industry Probability Estimates

	1	2	3	4	5	6	7	8	9
Event Probability	0.9580	0.9958	0.9970	0.9988	0.5786	0.5794	0.5308	0.4822	0.4497
95% CI Lower Bound	0.95	0.99	0.99	0.99	0.51	0.51	0.46	0.42	0.38
95% CI Upper Bound	0.96	1.00	1.00	1.01	0.65	0.65	0.61	0.56	0.53
Spread Ratio Weights	N	Y	Y	Y	Y	Y	Y	Y	Y
Earnings Announcements	Y	Y	N	N	N	N	N	N	N
High Dividend Stocks	Y	Y	N	N	N	N	N	N	N
Nearest Exp. Only	N	N	N	Y	Y	Y	Y	Y	Y
Extreme Strikes Only	N	N	N	N	Y	Y	Y	Y	Y
Bid > 0 Only	N	N	N	N	N	Y	Y	Y	Y
Volume > 0 Only	N	N	N	N	N	Y	Y	Y	Y
Event > 0 Only	N	N	N	N	N	N	Y	Y	Y
Event t-stat > 1.96	N	N	N	N	N	N	N	Y	Y
Event t-stat > 5	N	N	N	N	N	N	N	N	Y
R^2	0.0069	0.0084	0.0084	0.0207	0.1660	0.1685	0.1769	0.1867	0.1924
N	2,178,708	2,178,708	1,928,191	400,633	5,231	4,992	3,635	3,168	2,821

Table 9: Brown Event energy Industry Probability Estimates

	1	2	3	4	5	6	7	8	9
Event Probability	1.0613	1.0290	1.0305	1.0367	0.9541	0.9543	0.9033	0.8644	0.7229
95% CI Lower Bound	1.05	1.02	1.03	1.03	0.90	0.90	0.84	0.80	0.65
95% CI Upper Bound	1.07	1.03	1.04	1.04	1.01	1.02	0.97	0.93	0.81
Spread Ratio Weights	N	Y	Y	Y	Y	Y	Y	Y	Y
Earnings Announcements	Y	Y	N	N	N	N	N	N	N
High Dividend Stocks	Y	Y	N	N	N	N	N	N	N
Nearest Exp. Only	N	N	N	Y	Y	Y	Y	Y	Y
Extreme Strikes Only	N	N	N	N	Y	Y	Y	Y	Y
Bid > 0 Only	N	N	N	N	N	Y	Y	Y	Y
Volume > 0 Only	N	N	N	N	N	Y	Y	Y	Y
Event > 0 Only	N	N	N	N	N	N	Y	Y	Y
Event t-stat > 1.96	N	N	N	N	N	N	N	Y	Y
Event t-stat > 5	N	N	N	N	N	N	N	N	Y
R^2	0.0064	0.0066	0.0045	0.0070	0.0592	0.0570	0.0664	0.0712	0.0942
N	581,517	581,517	525,758	238,631	8,405	7,082	3,668	3,085	1,810

Table 10: SCOTUS Event energy Industry Probability Estimates

	1	2	3	4	5	6	7	8	9	10
Event Probability	0.9835	0.9930	0.9924	0.9933	0.5810	0.5803	0.5515	0.5403	0.5240	0.524
95% CI Lower Bound	0.98	0.99	0.99	0.99	0.53	0.53	0.50	0.49	0.47	0.47
95% CI Upper Bound	0.99	1.00	1.00	1.00	0.64	0.64	0.61	0.60	0.58	0.58
Spread Ratio Weights	N	Y	Y	Y	Y	Y	Y	Y	Y	Y
Event Earnings Announcements	Y	Y	N	N	N	N	N	N	N	N
High Dividend Stocks	Y	Y	N	N	N	N	N	N	N	N
Nearest Expiration Only	N	N	N	Y	Y	Y	Y	Y	Y	Y
Extreme Strikes Only	N	N	N	N	Y	Y	Y	Y	Y	Y
Positive Bids Only	N	N	N	N	N	Y	Y	Y	Y	Y
Positive Volume Only	N	N	N	N	N	Y	Y	Y	Y	Y
Positive Event Only	N	N	N	N	N	N	Y	Y	Y	Y
Event Study t-statistic > 1.96	N	N	N	N	N	N	N	Y	Y	Y
Event Study t-statistic > 5	N	N	N	N	N	N	N	N	Y	Y
Event Study t-statistic> 10	N	N	N	N	N	N	N	N	N	Y
R^2	0.0091	0.0092	0.0092	0.0037	0.0964	0.0977	0.1032	0.1059	0.1110	0.111
N	1,079,327	1,079,327	974,781	310,737	5,052	4,737	4,311	4,144	3,617	3,617

Table 11: McCain Event energy Industry Probability Estimates

	1	2	3	4	5	6	7	8	9	10
Event Probability	1.0669	0.9798	0.9797	0.9718	0.8998	0.8988	0.8959	0.8838	0.8533	0.8533
95% CI Lower Bound	1.05	0.97	0.97	0.96	0.70	0.70	0.66	0.65	0.60	0.60
95% CI Upper Bound	1.08	0.99	0.99	0.98	1.15	1.16	1.22	1.21	1.22	1.22
Spread Ratio Weights	N	Y	Y	Y	Y	Y	Y	Y	Y	Y
Event Earnings Announcements	Y	Y	N	N	N	N	N	N	N	N
High Dividend Stocks	Y	Y	N	N	N	N	N	N	N	N
Nearest Expiration Only	N	N	N	Y	Y	Y	Y	Y	Y	Y
Extreme Strikes Only	N	N	N	N	Y	Y	Y	Y	Y	Y
Positive Bids Only	N	N	N	N	N	Y	Y	Y	Y	Y
Positive Volume Only	N	N	N	N	N	Y	Y	Y	Y	Y
Positive Event Only	N	N	N	N	N	N	Y	Y	Y	Y
Event Study t-statistic $>$ 1.96	N	N	N	N	N	N	N	Y	Y	Y
Event Study t-statistic $>$ 5	N	N	N	N	N	N	N	N	Y	Y
Event Study t-statistic$>$ 10	N	N	N	N	N	N	N	N	N	Y
R^2	0.0133	0.0069	0.0076	0.0131	0.1070	0.1049	0.1085	0.0669	0.0458	0.0458
N	492,986	492,986	387,946	78,398	776	674	520	483	396	396

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

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

for some $\lambda \in [0, 1]$. We can ignore the case with $C_{x,T}(S_x, K) < E_x\left[S_T\right] - K$: our nonparametric bound on $p^H_{\tau-1}$ 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

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

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

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

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

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

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.[40] We seek $p^M_{\tau-1}$.

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

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

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

$$\begin{aligned}
&\Gamma_{\tau-1,T}(S_{\tau-1}, K_1, K_2) \\
\triangleq& C_{\tau-1,T}(S_{\tau-1}, K_1) - C_{\tau-1,T}(S_{\tau-1}, K_2) \\
=&\frac{1}{R_{\tau-1,T}} \int_{K_1}^{K_2} (S_T - K_1) \left[p^L_{\tau-1} f_{\tau-1}(S_T | S^L_{\tau-1}, L) + p^M_{\tau-1} f_{\tau-1}(S_T | S^M_{\tau-1}, M) + p^H_{\tau-1}\, f_{\tau-1}(S_T | S^H_{\tau-1}, H)\right] \, \mathrm{d}S_T \\
&+ \frac{K_2 - K_1}{R_{\tau-1,T}} \int_{K_2}^\infty \left[p^L_{\tau-1} f_{\tau-1}(S_T | S^L_{\tau-1}, L) + p^M_{\tau-1} f_{\tau-1}(S_T | S^M_{\tau-1}, M) + p^H_{\tau-1}\, f_{\tau-1}(S_T | S^H_{\tau-1}, H)\right] \, \mathrm{d}S_T.
\end{aligned}$$

[40] If 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:

$$\begin{aligned}
&\Gamma^M_{\tau-1,T}(S^M_{\tau-1}, K_1, K_2) - \Gamma_{\tau-1,T}(S_{\tau-1}, K_1, K_2)\\
=&(1 - p^M_{\tau-1})\frac{1}{R_{\tau-1,T}}\int_{K_1}^{K_2}(S_T - K_1)\ \left[f_{\tau-1}(S_T|S^M_{\tau-1}, M) - f_{\tau-1}(S_T|S^{\neg M}_{\tau-1}, \neg M)\right]\ \mathrm{d}S_T\\
&+ (1 - p^M_{\tau-1})\frac{K_2 - K_1}{R_{\tau-1,T}}\int_{K_2}^{\infty}\left[f_{\tau-1}(S_T|S^M_{\tau-1}, M) - f_{\tau-1}(S_T|S^{\neg M}_{\tau-1}, \neg M)\right]\ \mathrm{d}S_T\\
=&(1 - p^M_{\tau-1})\,\Gamma^M_{\tau-1,T}(S^H_{\tau-1}, K_1, K_2) - (1 - p^M_{\tau-1})\,\Gamma^{\neg M}_{\tau-1,T}(S^{\neg M}_{\tau-1}, K_1, K_2),
\end{aligned}$$

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

$$p^M_{\tau-1} \leq \frac{\Gamma_{\tau-1,T}(S_{\tau-1}, K_1, K_2)}{\Gamma^M_{\tau-1,T}(S^M_{\tau-1}, K_1, K_2)}.$$

We again have an upper bound on the desired risk-neutral probability.[41] The bound becomes tighter as $\Gamma^{\neg M}_{\tau-1,T}(S^{\neg M}_{\tau-1}, K_1, K_2)$ becomes small, which occurs when $f_{\tau-1}(S_T|S^L_{\tau-1}, L) \to 0$ 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

$$\Gamma^{\neg M}_{\tau-1,T}(S^{\neg M}_{\tau-1}, K_1, K_2) \geq \frac{K_2 - K_1}{R_{\tau-1,T}}\int_{K_2}^{\infty}\left[f_{\tau-1}(S_T|S^L_{\tau-1}, L) + f_{\tau-1}(S_T|S^H_{\tau-1}, H)\right]\ \mathrm{d}S_T.$$

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 $p^M_{\tau-1}$ 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 (8)

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

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

[41] Intuitively, 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 $p^M_{\tau-1}$ 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

$$\begin{aligned} Var_{\tau-1}[S_\tau] =& p^H_{\tau-1} E_{\tau-1}\left[(S^H_\tau)^2\right] + (1-p^H_{\tau-1})E_{\tau-1}\left[(S^L_\tau)^2\right] - [S_{\tau-1}]^2 \\ =& p^H_{\tau-1} Var_{\tau-1}\left[S^H_\tau\right] + (1-p^H_{\tau-1}) Var_{\tau-1}\left[S^L_\tau\right] \\ &+ p^H_{\tau-1}(S^H_{\tau-1})^2 + (1-p^H_{\tau-1})(S^L_{\tau-1})^2 - [S_{\tau-1}]^2 \\ =& \frac{p^H_{\tau-1}}{1-p^H_{\tau-1}}(S^H_{\tau-1} - S_{\tau-1})^2 + p^H_{\tau-1} Var_{\tau-1}\left[S^H_\tau\right] + (1-p^H_{\tau-1}) Var_{\tau-1}\left[S^L_\tau\right], \end{aligned}$$

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

B.2 Proof of Proposition 1

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

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

Therefore:

$$\begin{aligned} V_{\tau-1,T} - \left(\frac{S^H_{\tau-1}}{S_{\tau-1}}\right)^2 V^H_{\tau-1,T} =& E_{\tau-1}\left[\left(\frac{S_\tau - S_{\tau-1}}{S_{\tau-1}}\right)^2 - \left(\frac{S^H_\tau - S^H_{\tau-1}}{S_{\tau-1}}\right)^2\right] \\ &+ (1-p^H_{\tau-1})E_{\tau-1}\left[\sum_{j=1}^{T-\tau}\left\{\left(\frac{S^L_{\tau+j} - S^L_{\tau+j-1}}{\tilde{R}_{\tau-1,\tau+j-1}S_{\tau-1}}\right)^2 - \left(\frac{S^H_{\tau+j} - S^H_{\tau+j-1}}{\tilde{R}_{\tau-1,\tau+j-1}S_{\tau-1}}\right)^2\right\}\right]. \end{aligned} \tag{B-1}$$

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

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

Adding and subtracting $(1-p^H_{\tau-1})E_{\tau-1}[(S^L_\tau - S^L_{\tau-1})^2 - (S^H_\tau - S^H_{\tau-1})^2]$ and simplifying, we then obtain:

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

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

$$\begin{aligned}
V_{\tau-1,T} - \left(\frac{S^H_{\tau-1}}{S_{\tau-1}}\right)^2 V^H_{\tau-1,T} =& \frac{2\tilde{R}_{\tau-1,\tau}-1}{[S_{\tau-1}]^2}\left\{p^H_{\tau-1}[S^H_{\tau-1}]^2 - [S_{\tau-1}]^2 + (1-p^H_{\tau-1})[S^L_{\tau-1}]^2\right\} \\
&+ (1-p^H_{\tau-1})E_{\tau-1}\left[\sum_{j=0}^{T-\tau}\left\{\left(\frac{S^L_{\tau+j} - S^L_{\tau+j-1}}{\tilde{R}_{\tau-1,\tau+j-1}S_{\tau-1}}\right)^2 - \left(\frac{S^H_{\tau+j} - S^H_{\tau+j-1}}{\tilde{R}_{\tau-1,\tau+j-1}S_{\tau-1}}\right)^2\right\}\right],
\end{aligned}$$

which in turn implies:

$$\begin{aligned}
[S_{\tau-1}]^2 V_{\tau-1,T} - [S^H_{\tau-1}]^2 V^H_{\tau-1,T} =& (2\tilde{R}_{\tau-1,\tau}-1)\left(p^H_{\tau-1}[S^H_{\tau-1}]^2 - [S_{\tau-1}]^2 + (1-p^H_{\tau-1})[S^L_{\tau-1}]^2\right) \\
&+ (1-p^H_{\tau-1})\left([S^L_{\tau-1}]^2 V^L_{\tau-1,T} - [S^H_{\tau-1}]^2 V^H_{\tau-1,T}\right).
\end{aligned}$$

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

$$\frac{p^H_{\tau-1}}{1-p^H_{\tau-1}} = \frac{[S_{\tau-1}]^2 V_{\tau-1} - [S^H_{\tau-1}]^2 V^H_{\tau-1}}{(2\tilde{R}_{\tau-1,\tau}-1)\left[S^H_{\tau-1} - S_{\tau-1}\right]^2} + (1-p^H_{\tau-1})\frac{[S^H_{\tau-1}]^2 V^H_{\tau-1,T} - [S^L_{\tau-1}]^2 V^L_{\tau-1,T}}{(2\tilde{R}_{\tau-1,\tau}-1)\left[S^H_{\tau-1} - S_{\tau-1}\right]^2}. \tag{B-2}$$

This is the analogue of equation (9), adapted for the possibility that $\tilde{R}_{\tau-1,T} > 1$ and for the use of simple variance swaps. The first part of the proposition follows from using the implicit function theorem on equation (B-2) to obtain the derivative of $p^H_{\tau-1}$ with respect to $\left([S^H_{\tau-1}]^2 V^H_{\tau-1,T} - [S^L_{\tau-1}]^2 V^L_{\tau-1,T}\right)$ and then taking a first-order Taylor approximation around $\left([S^H_{\tau-1}]^2 V^H_{\tau-1,T} - [S^L_{\tau-1}]^2 V^L_{\tau-1,T}\right) = 0$:

$$p^H_{\tau-1} = \tilde{p} + \frac{[1-p^H_{\tau-1}]\left[(S^H_{\tau-1})^2 V^H_{\tau-1,T} - (S^L_{\tau-1})^2 V^L_{\tau-1,T}\right]}{[2\tilde{R}_{\tau-1,\tau}-1]\left[S^H_{\tau-1} - S^L_{\tau-1}\right]^2} + O\left(\left[(S^H_{\tau-1})^2 V^H_{\tau-1,T} - (S^L_{\tau-1})^2 V^L_{\tau-1,T}\right]^2\right),$$

using $\left[S^H_{\tau-1} - S_{\tau-1}\right]^2 = (1 - p^H_{\tau-1})^2 \left[S^H_{\tau-1} - S^L_{\tau-1}\right]^2$. The second part of the proposition follows from solving for $p^H_{\tau-1}$ in equation (B-2) with assumptions on the relationship between $[S^L_{\tau-1}]^2 V^L_{\tau-1,T}$ and $[S^H_{\tau-1}]^2 V^H_{\tau-1,T}$. Finally, note that if $[S^L_{\tau-1}]^2 V^L_{\tau-1,T} \geq [S^H_{\tau-1}]^2 V^H_{\tau-1,T}$, then $p^H_{\tau-1} \geq 0$ requires $[S_{\tau-1}]^2 V_{\tau-1,T} \geq [S^H_{\tau-1}]^2 V^H_{\tau-1,T}$. Therefore, if $[S_{\tau-1}]^2 V_{\tau-1,T} < [S^H_{\tau-1}]^2 V^H_{\tau-1,T}$, then $[S^L_{\tau-1}]^2 V^L_{\tau-1,T} < [S^H_{\tau-1}]^2 V^H_{\tau-1,T}$. The final part of the proposition follows from applying this inequality in equation (B-2) and solving for $p^H_{\tau-1}$.

C Empirical Appendix

Calculating the Value of Variance Swaps