Appendix

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

A Supporting Analysis for the 2016 Election

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

A.1 Robustness Tests for $\tilde{p}$

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

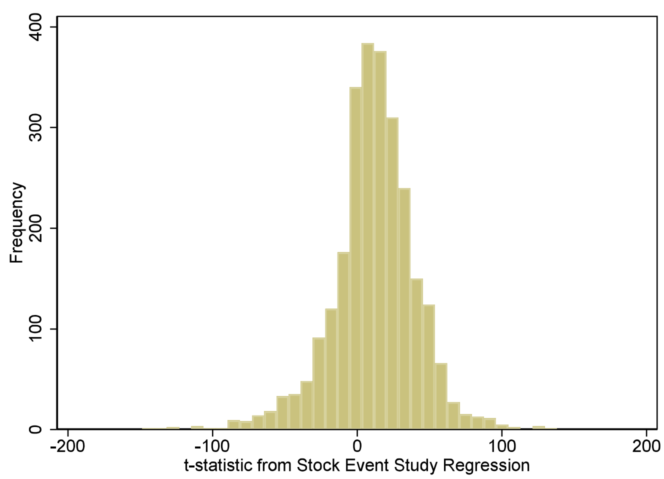

Figure A-1: Distribution of statistical significance of firms' responses to the 2016 U.S. election.

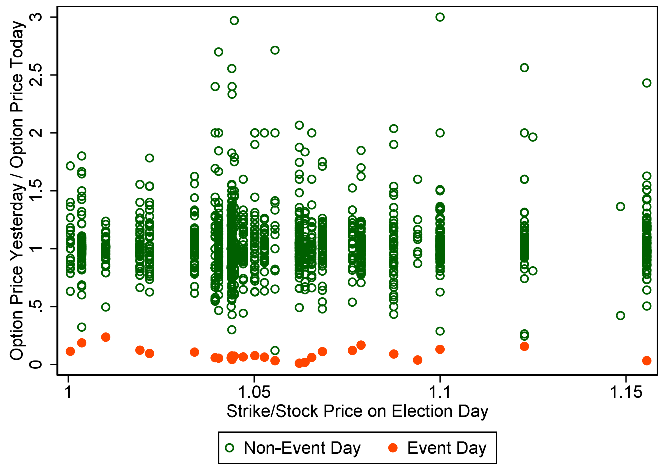

Figure A-2: Distribution of option prices on day $t-1$ to option prices on day t. For non-event and event days for firms with election stock event-study t-statistics of 75 or higher.

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

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

A.2 Robustness Tests for $\tilde{p}$

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

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

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

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

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

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

1. We convert observed option prices into implied volatilities. We obtain equity options' implied volatilities directly from OptionMetrics (following Carr and Wu, 2009). This approach strips away early exercise premia, yielding the European-style option prices required by the variance swap formula (9). These premia tend to be small in practice.

2. We calculate the forward price by adapting the methodology used to construct the familiar VIX index.[2]

3. We drop options that are likely to be illiquid, as in the main text.

4. We construct a fine grid of option prices by linearly interpolating implied volatilities at 10,000 different moneyness levels.[3] For moneyness levels below (above) the smallest (largest) observed moneyness, we use the implied volatility from the smallest (largest) observed liquid strike.

5. We convert back to option prices using the Black-Scholes formula for equities (following Carr and Wu, 2009).

6. We calculate the variance swap rate, using $R_{\tau-1,T}$ from a daily version of the 3-month LIBOR rate (as in the main text).

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

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

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

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

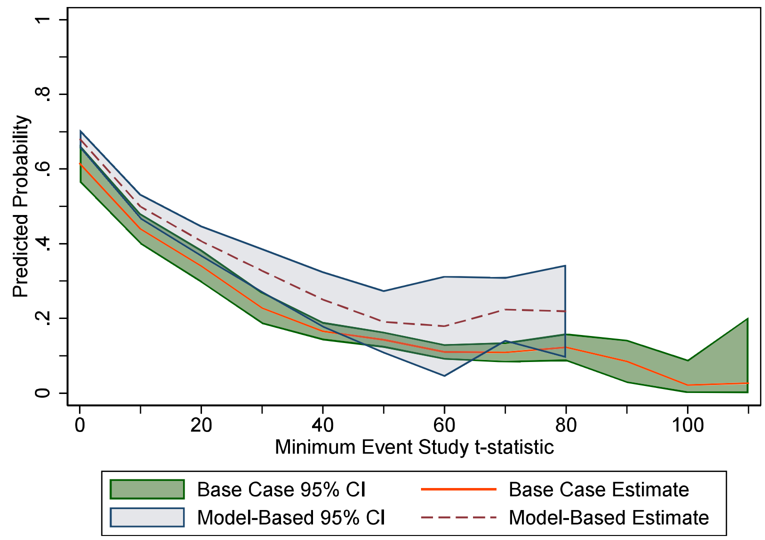

Figure A-3: Sensitivity of $\tilde{p}$ to discretization methods for 2016 election

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

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

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

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

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

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

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

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

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

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

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

Standard errors in parentheses, clustered by date.

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

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

Standard errors in parentheses, clustered by date.

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

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

Standard errors in parentheses, clustered by date.

B Supporting Analysis for OPEC Meetings

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

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

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

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

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

Table B-6: Estimated Probabilities for OPEC Meetings

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

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

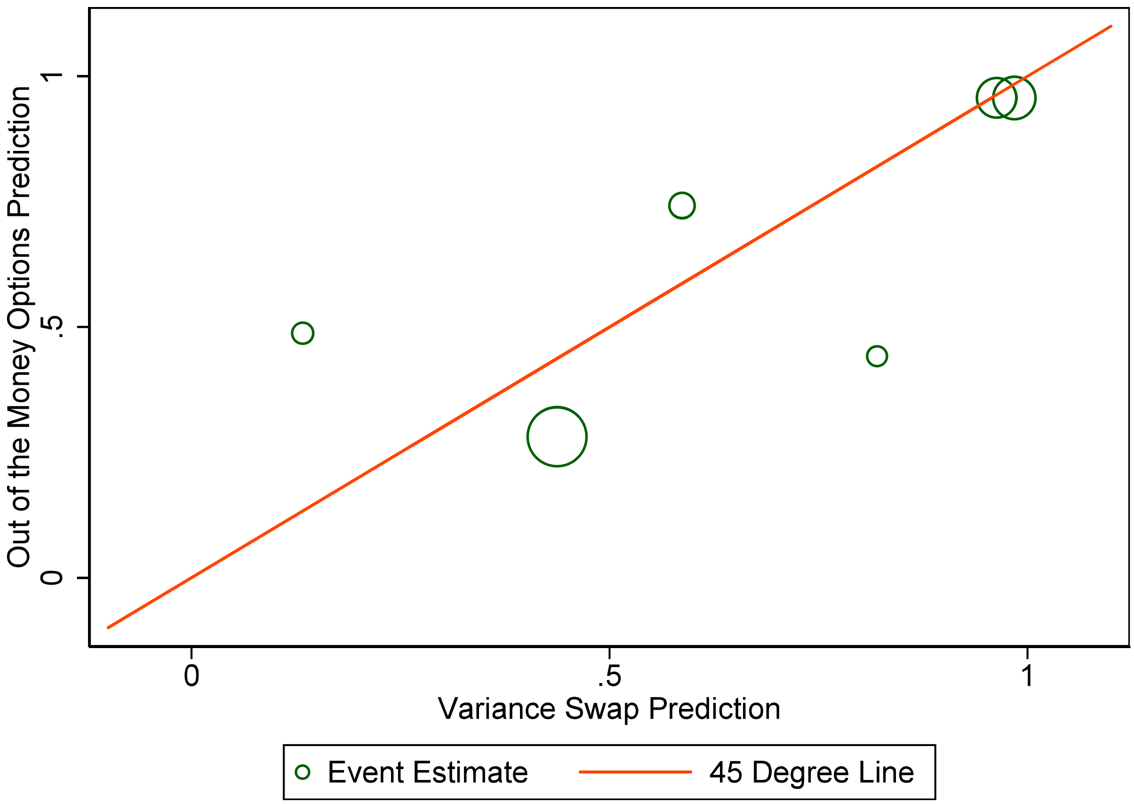

Figure B-4: Comparison Between $\bar{p}$ and $\tilde{p}$ Estimates

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

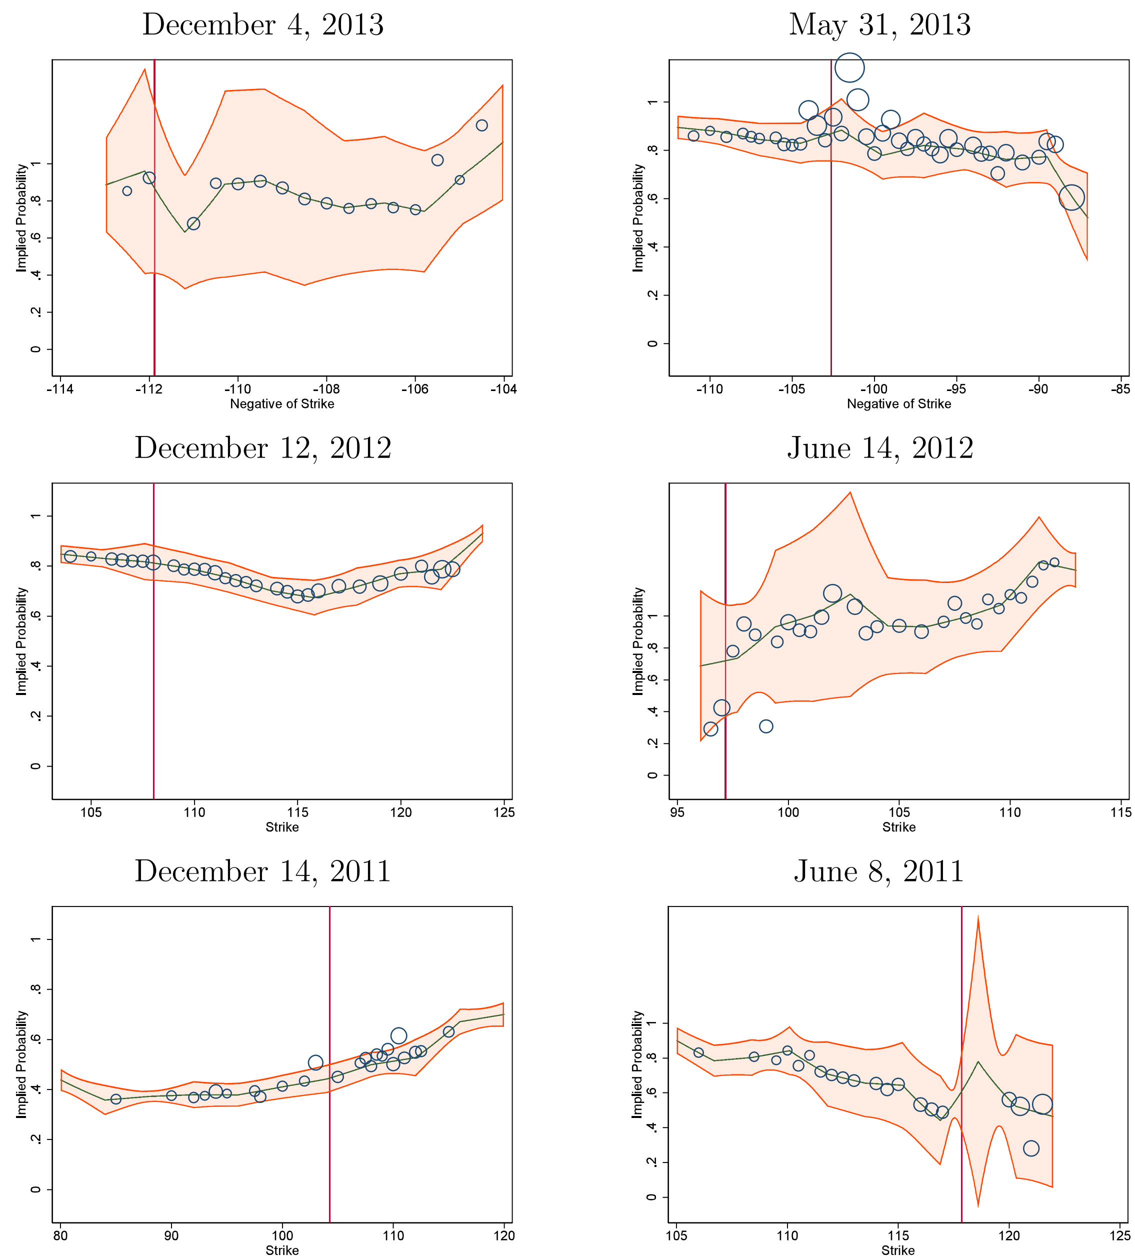

Figure B-5: Spline fits for additional OPEC meeting dates in 2011-2013.

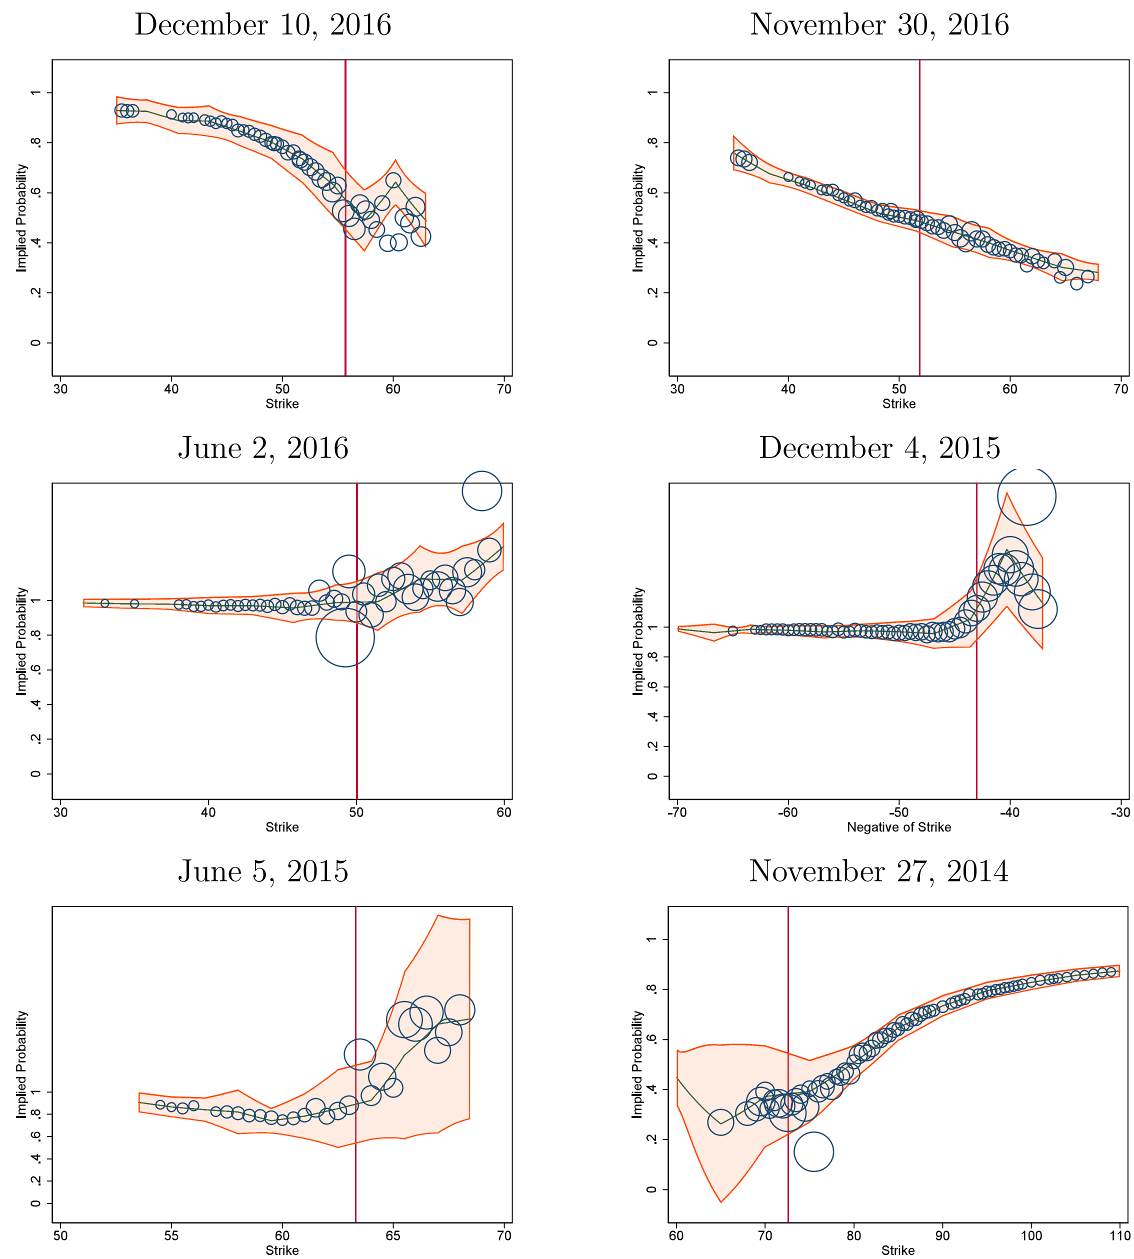

Figure B-6: Spline fits for additional OPEC meeting dates in 2014-2016.

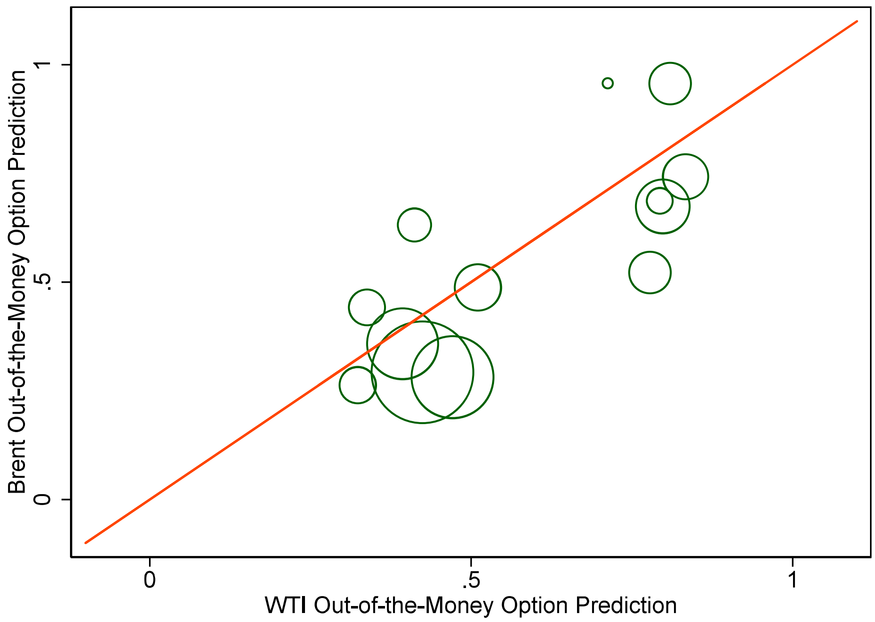

Figure B-7: Comparison Between Brent and WTI Estimates of $\bar{p}$

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

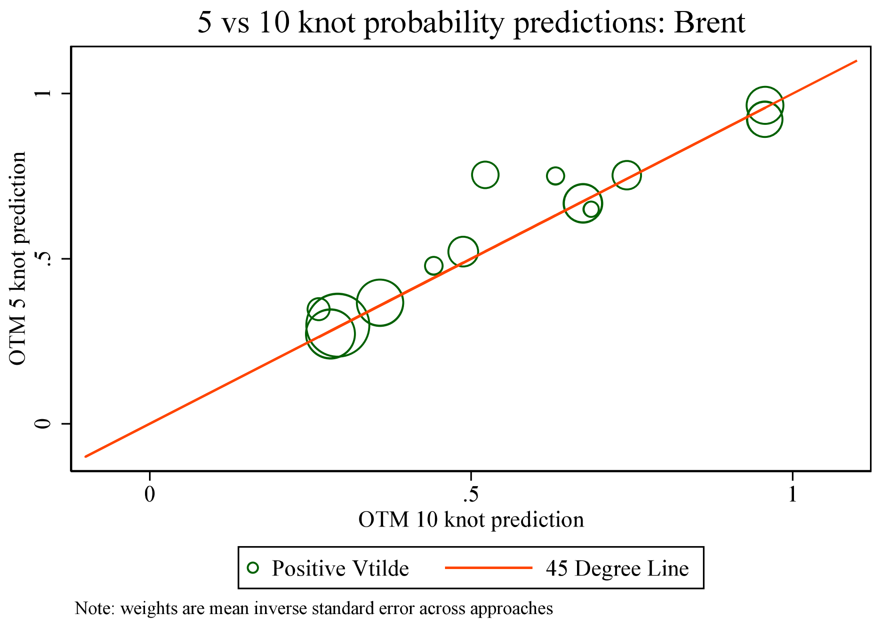

Figure B-8: Comparison Between 5-knot and 10-knot Estimates of $\bar{p}$

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

C Extensions to Section 3.1

C.1 American Options

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

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

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

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 e. 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.

C.2 When the Realized Event Was Not Extreme

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

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

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

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

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

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

where $\neg M$ means that $e\in\{L,H\}$. 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.[7] 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.

D Derivations for the Variance Swap Analysis

D.1 Equation (7)

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

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

[7] 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.

D.2 Proof of Proposition 1

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

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

Therefore:

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

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

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

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

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

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