Recovering the FOMC Risk Premium*

Hong Liu
Washington University in St. Louis

Xiaoxiao Tang
University of Texas at Dallas

and

Guofu Zhou
Washington University in St. Louis†

First Draft: November, 2019
Current Version: May, 2020

*We are grateful to Xiwei Tang for his help on maching learning methodology. We also thank seminar participants at Washington University in St. Louis for helpful comments and suggestions.

†Hong Liu, Olin Business School, Washington University in St. Louis, liuh@wustl.edu. Xiaoxiao Tang, Naveen Jindal School of Management, University of Texas at Dallas, xiaoxiao.tang@utdallas.edu. Guofu Zhou, Olin Business School, Washington University in St. Louis, zhou@wustl.edu.

Recovering the FOMC Risk Premium*

Abstract

The Federal Open Market Committee (FOMC) meetings have significant impact on market returns. We propose a methodology to recover the risk premium associated with FOMC meetings from option prices. We also estimate the sizes of upward/downward market price jumps after an imminent FOMC meeting. In our empirical analysis, with observed price data for 67 meetings and with data backed out via machine learning for the remaining 109 meetings from 1996 to 2017, we find that the risk premium varies from 15 to 88 basis points (bps), with an average of 38 bps which is consistent with the average realized returns documented in the literature. The average upward jump size is 103 bps, and the average downward jump size is 137 bps.

JEL Classification: G11, G14

Keywords: Options, FOMC Meeting, Jumps, Recovery

1 Introduction

The Federal Open Market Committee (FOMC) meetings announce key decisions about interest rates and the growth of the United States money supply. It is generally recognized that an FOMC announcement has a very significant influence on the market (see, e.g., Bernanke and Kuttner (2005)). In particular, Savor and Wilson (2013) and Lucca and Moench (2015) find that investors require a much higher risk compensation during a time interval of 24 to 48 hours before a meeting. Studies so far have only computed the ex post average realized excess return after FOMC meetings, which can be a good estimate of the FOMC unconditional risk premium. However, if like many other types of risk premia, the FOMC risk premium also changes with economic conditions, then the ex post average realized excess return would not be a good estimate of the risk premia right before the meetings. The economic question of interest is whether the FOMC risk premia are really time-varying, and if so, how we can estimate them.

In this paper, we propose a methodology to recover the FOMC risk premium around the FOMC meetings using options written on the S&P 500 index and a simple equilibrium model. The main idea is as follows. If there are options that expire right after the announcement of an FOMC meeting, then in an efficient option market, the option prices right before the announcement should correctly capture *exclusively* the risk from the imminent FOMC decision. The probability that the release of information about another risk factor coincides exactly with the FOMC announcement is virtually zero. As a result, the implied risk premium from these option prices should be equal to the FOMC risk premium without any contamination from other potential risk factors.

The S&P 500 options are the most liquid options written on the broad market index and are the most widely-studied derivatives for analyzing the U.S. financial market. We assume that immediately after an FOMC announcement, the S&P 500 index can take only two possible values, either jumps up (corresponding to a "good" surprise) or jumps down (corresponding to a "bad" surprise).[1] Accordingly, we consider a simple one-period binomial tree model for the S&P 500 index and use the just-before-announcement prices of options that expire right after the announcement to estimate the jump sizes. Then, in a representative agent framework with a constant relative risk aversion (RRA) utility, we recover the corresponding physical probabilities of the potential upward and downward jumps, which, combined with the estimated jump sizes, yield the FOMC risk premium for a given FOMC meeting.

Empirically, consider first the estimation of jump sizes due to FOMC meetings. For 67 FOMC meetings over the recent 12 years from 2006 to 2017, the option prices with maturity covering the 3 days around the meetings are available.[2] Our estimates indicate an average upward jump size of 116 basis points (bps) and an average downward jump size of 153 bps for the S&P 500 index ex post the meetings. These estimates show a remarkable performance in a "pseudo" prediction in which we use up or down jump returns to analyze returns conditional on up or down. The "pseudo" predicted return has a correlation of 89% with the realized return, and the "pseudo" out-of-sample R-squared is as high as 76.5%. This suggests that our estimation of jump sizes is surprisingly precise. For the entire sample of the 176 FOMC meetings between 1996 and 2017, we fill in the missing data via machine learning to estimate the implied volatility surface, which allows in general to extrapolate option prices for any desired strike prices and time-to-maturity from existing ones. The average upward and downward jump sizes are 103 and 137 bps, respectively. These values are lower than those for the 67 meetings, but their values match closely with the realized returns during the entire sample.

The FOMC risk premium is found to be time-varying with strong empirical evidence. Its value varies from 15 to 88 bps over the period from 1996 to 2017, assuming a typical value of 5 for the RRA. The volatility, however, is 10 bps, much smaller than the range.

[1] When the time horizon under consideration is almost zero, this assumption is without loss of generality.

[2] Ideally, we need the just-before-announcement prices of options that mature right after the announcement to capture only the effects of the meetings rather than other new risks arrived before or after the meetings.

Its long-term average is 38 bps, matching well those documented in the literature based on realized returns.

Our paper is the first to recover the time-varying FOMC risk premium from option market. Our empirical results show that the investors' demand of risk compensation around the FOMC meeting is well reflected by the option market. This, in turn, allows us to extract forward-looking information, and to estimate the time-dependent conditional FOMC risk premium from option prices regarding each meeting. The large time-variation in the estimated FOMC risk premium suggests that it is important to understand the risks of each meeting for its asset pricing implications on investments and corporate planning.

Our paper is closely related to the literature that shows the existence of the FOMC risk premium. Savor and Wilson (2013) and Lucca and Moench (2015) present empirical evidence of the premium. Ai and Bansal (2018) explain the premium with a theoretical model. Savor and Wilson (2014) and Ai et al. (2019b) show that the cross-section of stocks also behaves differently around macroeconomic announcements. We add to the literature by finding time-varying risk premium evidence from the option market.

Our paper also contributes to the recovery literature. Ross (2015) proposes a theory to recover the entire physical distribution of market returns from options written on the S&P 500 index. Subsequent papers focus on recovering the expected return of assets from option prices under normal market conditions and over a relatively long time interval (e.g., Martin (2017), Kadan and Manela (2019), Tang (2018), Jensen et al. (2019), Kadan and Tang (2020), and Martin and Wagner (2019)). Our paper introduces a methodology to recover expected returns with jumps within a short time horizon on news or events.

Our paper also adds to the literature that explores the relation between the FOMC meetings and the option market behavior. Vähämaa and Äijö (2011) and Amengual and Xiu (2018) find that the FOMC meeting significantly influences the behavior of implied volatility of market options through an uncertainty channel. Neuhierl and Weber (2016) construct a slope factor based on implied volatilities of different maturities to predict future market returns during the FOMC meetings. Ai et al. (2019a) show that S&P 500 index option prices around the FOMC announcements identify investors' preference for the timing of resolution of uncertainty. Our paper shows that option prices can be directly used to measure and predict stock jumps around the FOMC meetings.

The rest of the paper proceeds as follows. Section 2 introduces our theoretical framework and estimation methodology. In Section 3, we conduct estimation using option prices of the S&P 500 index. Section 4 presents further analysis and Section 5 concludes.

2 Theoretical Framework

In this section, we consider an asset whose price will likely experience either an upward or a downward jump after an imminent event (for example, earnings announcement, important economic news release). We begin with a one-period binomial tree model to estimate the jump sizes from observed option prices. Then, we employ a representative agent equilibrium model to recover physical probabilities of the upward jump and the downward jump, and the implied risk premium for the asset associated with the event. This method can be applied to any asset that will likely experience a jump in its price after an imminent event and has liquidly traded options on it that mature shortly after the event.

2.1 Jump Sizes and State Prices

Consider a one-period binomial tree model for an asset where the time starts at $t = 0$ and ends at $t = 1$. Let S_0 denote the price of the asset at $t = 0$. Assume an event occurs at $t = 1_-$, immediately after which the price of this asset either jumps up by $u\, S_0$, or jumps down by $d\, S_0$, with $u, d > 0$. Assume that there are two call options and two put options written on this asset, all maturing at $t = 1$. These options have prices C_1, C_2, P_1, and P_2 with respective strike prices of K_1^C, K_2^C, K_1^P, and K_2^P such that $(1-d)\,S_0 < K_1^C, K_2^C < (1+u)\,S_0$ and $(1-d)\,S_0 < K_1^P, K_2^P < (1+u)\,S_0$. Thus, at time $t=1$, the call options are exercised only when the realized state is u, with payoff $(1+u)\,S_0 - K_i^C$, and the put options are exercised only when the realized state is d, with payoff $(1-d)\,S_0 - K_i^P$, for $i = 1, 2$. Otherwise, the options are not exercised and the buyer gets zero payoff.

Let π_u denote the state price of state u (i.e., the price of an Arrow-Debreu security that pays \$1 in the up state and 0 otherwise). Then the call options are priced as

$$C_i = \pi_u \left((1+u)\,S_0 - K_i^C\right), \qquad i = 1, 2.$$

Thus we have the state price for the up state as

$$\pi_u = \frac{(1+u)\,S_0 - K_1^C}{C_1} = \frac{(1+u)\,S_0 - K_2^C}{C_2}. \tag{1}$$

The solution to upward jump size u is then

$$u = \frac{K_1^C C_2 - K_2^C C_1}{S_0\,(C_2 - C_1)} - 1. \tag{2}$$

Plugging (2) into (1), we obtain our estimate of the state price of the upward jump state,

$$\pi_u = \frac{C_1 - C_2}{K_2^C - K_1^C}. \tag{3}$$

Therefore, with the information contained in a pair of call options, we are able to estimate the future upward jump size of an asset, as well as the corresponding state price.

Similarly, we can estimate the size of downward jumps and the state price π_d for the down state by a pair of put options. Following the same steps as above, we have that the downward jump size d can be estimated as

$$d = 1 - \frac{K_1^P P_2 - K_2^P P_1}{S_0 \left(P_2 - P_1\right)}, \tag{4}$$

and the state price of the down state d can be given by

$$\pi_d = \frac{P_1 - P_2}{K_1^P - K_2^P}. \tag{5}$$

Hence, using a pair of calls and a pair of puts, we can estimate both the jump sizes and states prices, which will be used below to estimate the risk premium.

For the choice of these options, options close to the money (CTM) are preferable. This is because CTM options are more liquid and thus their prices are less distorted by illiquidity. We use call option prices to estimate the upward jump size and the up state price because call options are likely to better reflect information about an upward jump.[3] Similarly, we use put option prices to estimate the downward jump size and the down state price.[4]

2.2 Jump Risk Premium

Now, to estimate the risk premium, we consider a one-period representative agent equilibrium model where a representative firm produces one consumption good at time 0 and time 1. The consumption good produced at time 0 is δ_0 and at time 1 is δ_u with probability p_u and δ_d ($< \delta_u$) with probability $p_d = 1 - p_u$. The representative agent owns the firm which issues one share of stock. The firm pays the stock holder the consumption good produced as dividend at each point in time. The agent chooses her consumptions to maximize her expected utility

[3]For example, as shown in An et al. (2014), informed traders with positive news mostly trade by buying call options.

[4]Since the asset itself can be viewed as an option with zero strike price, it can also be used to help estimate the state prices and jump sizes. In an efficient financial market, this would result in the same estimates. In our empirical analysis later, we find that the difference is smaller than 0.1%, implying a high degree of market efficiency.

subject to a budget constraint:

$$\max_{c_0,c_u,c_d} \quad v(c_0) + \rho\left(p_u v(c_u) + (1-p_u)\, v(c_d)\right),$$
$$s.t. \quad c_0 + \pi_u c_u + \pi_d c_d = w_0,$$

where $v(.)$ is the agent's utility function, c_0 denotes the investor's consumption at $t = 0$, c_u and c_d are her consumptions in the corresponding states at $t = 1$, ρ is the time discount factor, p_u is the physical probability of state u, and w_0 is the investor's initial wealth.

The first order conditions with respect to c_0, c_u, and c_d yield

$$\rho p_u \frac{v'(c_u)}{v'(c_0)} = \pi_u,$$
$$\rho(1-p_u)\frac{v'(c_d)}{v'(c_0)} = \pi_d,$$

which leads to

$$\frac{p_u}{1-p_u}\frac{v'(c_u)}{v'(c_d)} = \frac{\pi_u}{\pi_d}.$$

Solving for p_u, we have

$$p_u = \frac{\pi_u}{\pi_u + \frac{v'(c_u)}{v'(c_d)}\pi_d}. \tag{6}$$

Accordingly, the physical probability of a downward jump is given by

$$\begin{aligned} p_d &= 1 - p_u \\ &= 1 - \frac{\pi_u}{\pi_u + \frac{v'(c_u)}{v'(c_d)}\pi_d}. \end{aligned} \tag{7}$$

The equation makes intuitive sense. Everything else equal, the greater the upside (downside) state price is, the greater the upside (downside) probability is.

By the market clearing condition, we have

$$c_0 = \delta_0, c_u = \delta_u, c_d = \delta_d.$$

Using the same notation as before, let S_0 be the initial stock price at $t = 0$, $(1 + u)\, S_0$ be the up state stock price at $t = 1$, and $(1 - d)\, S_0$ be the down state stock price at $t = 1$. The (cum-dividend) stock price at time 1 in equilibrium must be equal to the dividend payment in each state. Therefore, we have

$$(u + 1)\, S_0 = \delta_u, (1 - d)\, S_0 = \delta_d.$$

As a result, we must have

$$c_u = (1 + u)S_0, \tag{8}$$

$$c_d = (1 - d)S_0. \tag{9}$$

Thus, for any given form of the representative agent's utility function, we are able to recover the physical probabilities with the estimates presented in Section 2.1. In particular, based on the estimates of physical probabilities, the estimation of the jump risk premium is given by:

$$E\,(\tilde{r}) - r_f = p_u u - p_d d - r_f, \tag{10}$$

where $\tilde{r}$ is the net return of the stock and r_f is the observed risk-free rate between time $t = 0$ and $t = 1$. In our empirical analysis, because the time period under consideration is short (only up to a couple of days), the level of the risk free rate is not important and accordingly, we assume that $r_f = 0$.

2.3 CRRA Utility

To completely determine the physical probabilities from (6) and (7), we need to specify a form of utility function. Here we focus on the constant relative risk aversion utility:

$$v(c) = \frac{c^{1-\gamma}}{1-\gamma},$$

where γ is the relative risk aversion (RRA) coefficient. Then, the marginal utility is given by $v'(c) = c^{-\gamma}$, and so the estimate of physical probability of an upward jump is

$$\begin{aligned} p_u &= \frac{\pi_u}{\pi_u + \frac{v'(c_u)}{v'(c_d)}\pi_d} \\ &= \frac{\pi_u}{\pi_u + \frac{c_u^{-\gamma}}{c_d^{-\gamma}}\pi_d} \\ &= \frac{\pi_u}{\pi_u + \frac{(1-d)^\gamma}{(u+1)^\gamma}\pi_d}, \end{aligned}$$

where the last equality holds because of (8) and (9). In our estimation below, we consider three levels of RRA: 5, 8, and 10, which are found to be reasonable levels of RRA in the literature.[5] These different levels of RRA imply different values of p_u. In the following analysis, however, we focus on discussing the results with $\gamma = 5$.

Recall that in Section 3.3, we estimate the state price π_u (π_d) of the up (down) state using two cal (put) options under the assumption that call (put) options provide more precise information about the up (down) state. Consistent with this assumption, we next estimate physical probabilities in two steps. In the first step, we estimate the physical probability p_u of state u, based on the up state price π_u estimated using two call options and the implied

[5] See, for example, Mehra and Prescott (1985), Vissing-Jørgensen and Attanasio (2003), and Bliss and Panigirtzoglou (2004)

down state price $1-\pi_u$ (because the risk free rate r_f is assumed to be 0 as explained before),

$$p_u = \frac{\pi_u}{\pi_u + \frac{(1-d)^\gamma}{(u+1)^\gamma}\left(1-\pi_u\right)}, \tag{11}$$

and that of state d based on the down state price π_d estimated using two put options and the implied up state price $1-\pi_d$:

$$p_d = 1 - \frac{1-\pi_d}{1-\pi_d + \frac{(1-d)^\gamma}{(u+1)^\gamma}\pi_d}. \tag{12}$$

As a result, we obtain two pairs of the physical probability estimates: $(p_u, 1-p_u)$ and $(1-p_d, p_d)$. Then in the second step, we calculate a weighted average of the two sets of probabilities. In particular, we set the weights to be proportional to the jump sizes:

$$\begin{aligned} \hat{p}_u &= \frac{u}{u+d} p_u + \frac{d}{u+d}\left(1-p_d\right), \\ \hat{p}_d &= 1-\hat{p}_u, \end{aligned} \tag{13}$$

to obtain our final probability estimates. Clearly, $\hat{p}_u + \hat{p}_d = 1$ with both $\hat{p}_u$ and $\hat{p}_d$ being positive. Intuitively, when the upward (downward) jump size is relatively large, the up (down) state information is more important and thus we should rely more on the up (down) state estimates from call (put) options.

2.4 Discussions

The proposed methodology is general enough for estimating the jump sizes and jump risk premium for any assets with liquid option trading. In this paper, we consider the Federal Open Market Committee (FOMC) meeting day as a candidate time interval during which a jump in the market price is likely to happen. Indeed, during the 24- to 48-hour time window right before the FOMC meetings, significant amount of news about the monetary policy is expected to release, bringing potential shocks to the market. Savor and Wilson (2013) and

Lucca and Moench (2015) find that market excess returns over this time interval account for 60% to 80% of the market annual returns. As the dates of regular FOMC meetings are pre-specified, the jumps are expected by the investors. Therefore, it is reasonable to assume that a jump in the market index arrives during this time window, and the investors can correctly predict it.

Under our framework, if we consider $t = 0$ as 24 or 48 hours before the FOMC meeting, and $t = 1$ as the meeting day, then the above methodology can be used to estimate the upward and downward jump sizes for each meeting whenever we have option prices available. Such estimations are time conditional and forward-looking. In addition, for any given utility function of the representative agent, we are able to recover the risk premium of the market right before and on the meeting day, which is the FOMC risk premium.

Notably, the FOMC risk premium is estimated in a forward-looking manner, incorporating time-conditional information. Most of the current literature that examines the FOMC risk premium relies on the realized market returns, which only contain information about the realized states of the jumps. In other words, there is no way to obtain information about the downward jump state when the realization is an upward jump, and vice versa. Consequently, the premium estimated based on the realized returns is the unconditional one, thus cannot be utilized to examine the time-series variation of the FOMC risk premium. In contrast, we show that, with our methodology, we are able to recover and observe the dynamic changes of the FOMC risk premium.

3 Empirical Results

3.1 Data

The aim of our empirical analysis is to recover the FOMC risk premium from option prices. We obtain S&P 500 option pricing data from OptionMetrics with a sample period spanning

January 1996 to December 2017. For each year in our sample, there are 8 regular FOMC meetings, with a total of 176 meetings. We choose options with life span covering the 24-hour time interval right before the closing time on the FOMC announcement days. We use the most recent prices of these options before the FOMC meeting to estimate jumps of S&P 500 index on the meeting days. To cleanly identify the FOMC risk premium, we only consider options that mature within three days.[6] For each meeting day in our sample, we retain two calls and two puts that have strike prices closest to the current S&P 500 index level and with the shortest available maturity horizon. In addition, we obtain S&P 500 index levels from CRSP.

However, there are available data for only 67 out of 96 FOMC meetings over the recent 12 years from 2006 to 2017 that satisfy the above three-day expiration window. Prior to 2006, there are substantially fewer options available. In particular, there are no options that have maturity shorter than three days but still covers the 24-hour time window right before any FOMC meetings. The data availability is, however, not an issue going forward, as in the current market, there are options written on the S&P 500 index maturing on every Monday, Wednesday, and Friday (except for days that the market is closed). This promises that we almost always have available options to estimate jumps and other related variables for any FOMC meeting in the future.

The available options from 2006 to 2017 on the 67 meetings are all liquidly traded. Indeed, the average trading volume is 2325.37, with the smallest trading volume being 52 contracts. Thus, we believe that the options prices are reliable. We use the midpoint of their highest closing bid and lowest closing ask prices for the estimation. For each pair of call and put options, we follow the methodology presented in Section 2 to recover information about the upward and downward jump states, respectively.

To have more samples and to make our study comparable to those with the full 176 meetings from 1996 to 2017, we also back out the rest of the data from the implied volatilities

[6] As a robustness check, we relax this condition to four and five days. See Section 4.1 for details.

of options that are available 48 hours before the FOMC meetings and that expire after the meeting announcement date. This requires us to extrapolate the volatility relation between relatively short- and relatively long-term options. We use a recently developed machine learning method, the smoothing-embedded matrix completion, to fit the implied volatility surface. Details are provided in the Appendix.

3.2 Jump Estimates

Consider first the case with the actual option data for the recent 67 meetings. Panel A of Table 1 reports summary statistics of the variables used in the estimation. The mean prices of C_1, C_2, P_1, and P_2 are \$8.7776, \$11.1220, \$9.5299, and \$8.9813, respectively. These options are all very close to ATM, with moneyness ranging from 0.9569 to 1.0232. Thus their strike prices are quite likely to fall between the downward and upward jumps. There are 22.37% of the options mature in one day, 7.46% mature in two days and the rest of them mature in three days.

Panel B reports the summary statistics of the estimated jump sizes. The upward jump sizes range from 30 basis points (bps) to 3.47%, with a mean of 116 bps. The downward jump sizes are from 26 bps to 7.56%. The mean downward jump size is 153 bps. While mean seems quite reasonable, the lowest and highest jumps deserve some analysis and explanation.

The top panel of of Figure 1 plots the time series of the jump estimates along with the S&P 500 returns. The jump sizes are associated with volatility. The extreme low level of 7.56% downside jump occurs during the financial crisis period. Other than this one outlier, the estimated downside jumps are in fact quite close to the realized values when the market indeed goes down. In contrast, the extreme upward jump of 3.47% does not seem high as compared with the actual returns when the market indeed goes up. Overall, the magnitudes of the estimates are consistent with the realized jumps.

To better illustrate the precision of the jump estimates, we consider a pseudo prediction.

Let

$$\hat{r} = \begin{cases} u & \text{if } \tilde{r} > 0, \\ -d & \text{if } \tilde{r} < 0, \end{cases} \tag{14}$$

where $\tilde{r}$ is the realized net return of the S&P 500 index. Then $\hat{r}$ is the predicted jump by the option market conditional on either upside or downside jumps. This is a pseudo prediction because we combine the estimated jump sizes, u and d, with the directions of future realized jumps, and so we do not observe $\hat{r}$ before the jump realizations. Also it should be noted that, for each meeting, only one of u and d is realized.

Nevertheless, $\hat{r}$ contains information on how well the market anticipates the size of the jump. The bottom panel of Figure 1 plots $\hat{r}$ and $\tilde{r}$ over time. It is clear that there is a strong link between pseudo predictions and realizations, indicating that our estimation of the jump size is quite precise. Indeed, the correlation between the pseudo prediction and the realized return is 89%. Predicting $\tilde{r}$ with $\hat{r}$ yields a pseudo out-of-sample R-squared of 76.5%.

Consider now the full sample results. Figure 2 includes plots parallel to Figure 1 for the extended sample period with the extrapolated data. The top panel plots the estimated upward and downward jumps together with the realized returns. We can see that for the first few years (before 2003), the jump estimates based on the extrapolated data do not capture some of the high levels of realized jumps. This is reasonable, as the option trading before year 2003 is rather thin, leading to the loss of certain information. The machine learning model is trained only based on a few observations, resulting in less accuracy in our estimation. Similarly, during the financial crisis, the estimates based on extrapolation are not able to capture the spike in October 2008, where the one-day return of the S&P 500 index is 9.6%. Perhaps this is also because, when the volatility is too high, the estimation becomes less accurate, and hence the gap widens between what is estimated and what is realized in the presence of extreme events.

The bottom panel in Figure 2 presents the time-series of the realized returns with the pseudo prediction estimated from (14). The comparison shows that except for the early years and the extreme cases during financial crisis, the estimates match the realized return pattern closely. Even for the two periods that are lack of information, the estimates are still able to match the trend of the jumps. This demonstrates that the matrix completion methodology applies to the option data quite efficiently.

Panel C of Table 1 reports summary statistics of the jump estimates with the extrapolated data. The average estimated upward and downward jump sizes are 103 and 137 bps, respectively. Both are slightly below the average estimated jump sizes with real data. The reason might be due to financial crisis. During the crisis, extreme events are more likely to happen, and so there are quite a few greater jump sizes that make the average slightly higher. Now, on the pseudo prediction, the out-of-sample R-squared for the entire sample is 67.6%, which is still impressive as it is usually difficult to get predictors with such a high R^2. However, this value is about 10% lower than 76.5% for the subsample. It is likely caused by using the augmented data rather than the true but unavailable ones.

3.3 State Price Estimates

We next estimate the state prices following (3) and (5). As the time interval of interest is as short as only a few days, we assume that $r_f = 0$ as mentioned before. To preserve information from both call and put options, we estimate two pairs of the state prices. In particular, we estimate $(\pi_u, 1 - \pi_u)$ following (3) with two call options, and $(1 - \pi_d, \pi_d)$ following (5) with two put options.

For the sub-sample of 67 meetings, Panel A of Table 2 reports the summary statistics. There are two interesting facts. First, the two pairs of estimates are quite close to each other, which means that $\pi_u + \pi_d = 1$ roughly holds. Indeed, the average of $\pi_u + \pi_d$ is 0.9975, with a standard deviation of 0.0202. This is interesting because we use two different pairs of options, and there is no guarantee that $\pi_u + \pi_d = 1$ empirically. Furthermore, across all the quantiles, the two pairs of estimates are very close too. Figure 3 plots their values over time and their time-series patterns are very close as well. Therefore, in what follows, we continue to use both pairs of estimates to recover physical probabilities and combine the physical ones to recover the FOMC risk premium.

The second interesting fact is that, in general, the upward jump state has a higher state price than the downward state. This compensates the larger average downward jump size. The time-series plot in Figure 3 shows that the upward (downward) state prices are countercyclical (procyclical), which implies that good news is more desired during financial crisis than normal times.

Panel B of Table 2 presents summary statistics of the estimated state prices for the full sample. Again, we can see that the estimates based on call and put prices are quite close. The average of $\pi_u + \pi_d$ is 0.9921, with a standard deviation of 0.0505. Time series plot of the state prices for the full sample in Figure 4 further confirms this.

3.4 Recovery of Physical Probabilities

Table 3 and Figure 5 present results of estimated physical probabilities based on (11) and (12) for RRA (γ) levels of 5, 8, and 10. The results are for all the 176 meetings (the subperiod results are omitted for brevity). As expected, the physical probabilities of an upward (downward) jump are always higher (lower) than the risk neutral ones, with the spread being larger as γ increases from 5 to 10. For example, the average physical probability of an upward jump increases from 62% to 65% as γ increases from 5 to 10, compared to an average of 60% for the corresponding risk-neutral probability. This is due to the risk adjustment made by the representative agent. From the plots we can easily see that the spread of the estimates based on different levels of RRA is relatively small compared to the time-series volatility of the probabilities. Thus, our estimates provide reasonable and informative bounds on physical probabilities.

Table 4 and Figure 6 report estimations of weighted physical probabilities following (13).

By comparing them with results in Table 3 and Figure 5, we can see that the estimates are quite close. This is expected as the results in Section 3.3 show that the estimates from calls and puts are quite similar.

3.5 Recovery of the FOMC Risk Premium

With the estimates from previous steps, now we are able to estimate the FOMC risk premium as

$$E(\tilde{r}) = \hat{p}_u u - \hat{p}_d d, \tag{15}$$

for different levels of relative risk aversion, γ. Note that this equation is the same as (10) under the simplifying assumption that $r_f = 0$ in the one-day horizon.

Panel A of Table 5 presents summary statistics of the estimated FOMC risk premium for RRA levels of 5, 8, and 10 for the entire sample. The average risk premium lies between 38 bps and 50 bps depending on the level of risk aversion of the representative agent. The time-series volatility of the risk premium estimates ranges from 10 to 16 bps. This shows substantial fluctuations of the FOMC risk premium over time. For example, when $\gamma = 8$, the FOMC risk premium ranges from 16 bps to 120 bps over the 176 FOMC meetings considered in our sample.[7]

Figure 7 presents time-series plots of the estimated FOMC risk premiums for each RRA level. The estimates fluctuate significantly over time, consistent with large time-series volatility reported in Table 5. The estimated risk premium is always positive. While during the crisis, the FOMC risk premium is substantially higher, which indicates that investors demand higher compensation for monetary policy change during the market distress time. Under normal market conditions, the risk premium level is moderate, and is sometimes very

[7]For comparison, we have also estimated the risk premiums using two other approaches, Breeden and Litzenberger (1978) with CRRA utility and Martin (2017), respectively, and found that the estimates are too small (no greater than 3.5 bps), due to noisy out-of-money option prices when the maturity is short. In contrast, our methodology relies only on at-the-money or nearby options whose prices are much more reliable.

close to zero.

We also report summary statistics of the realized S&P 500 net returns. Notice that the realizations only reflect information of the realized jump states, thus it is expected to see that the realized returns are much more volatile than the estimated risk premium. Indeed, the time-series volatility of the realized returns is 160 bps and they range from -309 to 537 bps. The average realized return is 47 bps, which can be considered as an ex-post estimate of the unconditional FOMC risk premium for the 176 FOMC meetings of interest. This estimate falls within the interval provided by our methodology. In fact, it lies between the option-implied FOMC risk premium under cases when $\gamma = 8$ and when $\gamma = 10$.

For completeness, we report the summary statistics of the FOMC risk premium as well as the corresponding realized returns for the 67 subsample in Panel B of Table 5. The average FOMC risk premium ranges from 12 to 23 bps depending on the level of relative risk aversion. The average realized returns over life span of the corresponding options (which covers the 24-hour before the FOMC meeting) is 14 bps, falling between the estimated ranges based on our model. Note that the average risk premium for the sub sample, which contains the crisis, is lower than that of the entire sample. This appears counter-intuitive, but can be explained with two reasons. First, the sub-sample (the 67 meetings) do not cover every meeting in the recent 12 years. In particular, there are 4 missing observations for the year 2008 and 2 missing observations for the year 2009. It happens that the sub-sample does not cover the FOMC associated with the huge upward jump in 2008, so it does not fully capture the volatility during the crisis, and hence the associated risk premium. These missing data all play a role in the entire sample. Second, excluding the financial crisis (year 2008 and 2009), we find the realized return volatility before 2006 is 1.54%, much greater than 1.22%, the return volatility after year 2006. This indicates that we can expect a greater risk premium for the early sample than the later sample. Empirically, the data reveal the relatively sizable difference. In contrast to the early jump size estimation whose average seems to be driven more by the large values during the crisis, the average of the risk premium is less affected because the numerical values of the risk premium are smaller than those of the jumps.

4 Further Analysis

In this section, we present some further analysis to complement our main results. First, we investigate the sensitivity of our estimates to the maximum maturity horizon allowed. Then we present results excluding years 2008 and 2009 to learn the behavior of the FOMC risk premium during normal times. Finally, we analyze the relation between the FOMC risk premium with market uncertainty.

4.1 Maximum Maturity Horizon

We first check the robustness of maximum maturity allowed. In the results provided in Sections 3.2 and 3.3, we require the maturity of options to be less than or equal to three days, which leads to only 67 observations with real data. In this robustness check, we relax the maximum maturity to four or five days. Notice that the relaxation does not change the estimates for the 67 FOMC meetings, as the shortest option horizons on these days are already shorter than four or five days. This only allows us to include more observations with real data.

Table 6 reports the summary statistics of the estimated jump sizes under these specifications. When relaxing the maximum maturity horizon to four days, the number of days in our sample increases to 89. It increases further to 114 when the maturity horizons up to five days are allowed. The summary statistics are comparable to those in the main analysis.

Figure 8 plots the time-series of the pseudo prediction (estimated by (14)) along with the realized returns. From the plots we can see that allowing more options in the calculation extends our sample back to year 1996. When the maximum maturity is relaxed to four or five days, the pseudo predictions still agree with the realizations quite well. The corresponding out-of-sample R-squared are 72.9% and 72.5%, respectively, just slightly lower compared to the results when we restrict the option maturity to be less or equal to three days. These informative predictions indicate that the jump in the market index level right before the FOMC meetings dominates other fluctuations within the neighborhood of the meeting. However, to cleanly identify the jump sizes during the FOMC meeting, we need to focus on a short-time window. A bias in jump size estimation could potentially happen if more days are allowed after the announcement. Thus, all the conclusions in this paper only apply to options with short-term maturity horizons.

4.2 Excluding Financial Crisis Period

Results in Section 3 show that the estimates, as well as the market itself, behave differently during the financial crisis compared to normal times. In this subsection, we examine the results by excluding years 2008 and 2009 that cover the crisis. Note that removing the two years drops the number of FOMC meetings to 160.

Panel A of Table 7 reports the summary statistics of the estimated jump sizes. The average upward jump size is 103 bps and the average downward jump size is 100 bps. The average upward jump sizes are comparable to the results in Table 1, while average downward jump sizes are 30 bps smaller. Also note that the standard deviations of the jump estimates drop by 8 to 25 bps, which is consistent with the notion that market is more volatile during financial crisis (See Roll (1988), Schwert (1990), and Hong et al. (2007)).

Panel B of Table 7 reports summary statistics of the FOMC risk premium for these 160 meetings. After excluding years 2008 and 2009, the average FOMC risk premium drops to 33 (when $\gamma = 5$) to 44 (when $\gamma = 10$) bps, with a decrement of 5 to 6 bps compared to the whole sample. This is consistent with the observation in Section 3.5 that investors require higher risk compensation for news announcement during financial crisis.

4.3 Relation to Volatility

We next check the relation between our estimated FOMC risk premium and economic variables. Here we consider the CBOE Volatility Index (VIX), which is often considered as a popular measure of the option market's expectation of uncertainty. Similar to our FOMC risk premium, the VIX is also estimated from options written on the S&P 500 index. Therefore, the VIX also contains time-conditional information and can be estimated with high frequency (daily or even intra-daily) and in a forward-looking manner. It is thus interesting to investigate whether our recovered FOMC risk premium is related to the VIX.

In particular, we run the following time-series regression:

$$E_t(R) = \beta_0 + \beta_1 VIX_t + \epsilon_t, \tag{16}$$

where $E_t(R)$ is the FOMC risk premium evaluated at time t, and VIX_t is the closing level of the VIX at time t. We work with the estimated FOMC risk premiums with $\gamma = 5$, 8, and 10, and consider both the full sample and the sample based on real data (67 meetings during 2006 to 2017).

Table 8 reports the regression results. Columns (1) to (3) include the coefficient estimates for the full sample. We can see that for all levels of γ, the VIX is significantly related to the FOMC risk premium at the 1% level. For example, a one-percentage-point increase in the VIX relates to a 50-basis-point of 1% increase in the FOMC risk premium given $\gamma = 8$. The results with real data (reported in Columns (4) to (6)) exhibit similar patterns. Therefore, we conclude that the FOMC risk premium has a positive association with the VIX, which shows that the investors require a higher risk compensation during periods that the future uncertainty is expected to be high.

5 Conclusion

The Federal Open Market Committee (FOMC) meetings are major events that significantly impact the stock market. In this paper, we propose a simple model to estimate the risk premium of the FOMC meetings based on option prices. Our results indicate that the risk premium varies from 15 to 88 basis points (bps) depending on investors' risk aversion, with an average of 38 bps consistent with the related findings in the existing literature. Modeling the price move as a two-state jump process, we find that the average upward jump size is 103 bps, and the average downward jump size is 137 bps.

Our methodology applies not only to FOMC meetings, but also to any events after which asset prices are likely to experience significant jumps and there are options traded on these assets. In particular, for future research, it will be of interest to apply our approach to study the risk premium associated with earnings announcements for individual stocks in order to provide new insights about the findings of the vast literature on the effect of earnings announcement.

Appendices

I. Data Augmentation

In this appendix, we present a new machine learning method, the smoothing-embedded matrix completion (Dai et al., 2019), to extrapolate the implied volatility surface, and hence to fill in the option data with expirations right after the FOMC meeting announcements.

The method uses low-rank matrix factorization to complete a sparse matrix based on its underlying structure. Specifically for our problem, the implied volatility surface has a well known matrix structure. Each row in the matrix represents a moneyness level, and each column corresponds to a specific maturity. The observed data fill certain matrix entries, while leaving others blank. The matrix completion method is capable of filling the empty entries based on a penalized low-rank matrix decomposition.[8] We then use these extrapolated implied volatilities to obtain synthetic option prices.

Specifically, we consider the implied volatility data with moneyness levels from 0.5 to 1.5 and maturity values (day) from 1 to 100. Since the moneyness level is essentially a continuous variable, we divide the considered interval into 200 segments evenly, each with a range of 0.005. Consequently, this results in a 200-by-100 matrix in which the implied volatility values fill the entries with corresponding moneyness levels and maturity values. Given the selected segments, most of the observations have unique row and column indexes. If there are multiple observations sharing the same indexes (no more than two in our data sample), we just fill the matrix entry with the average value. The matrices for call and put options are constructed separately.

The low-rank matrix factorization model is a prevalent tool for matrix completion (Hastie et al., 2015), which enables us to complete the sparse matrix based on a small number of latent factors. Let $\mathbf{V}$ denote the target matrix of implied volatilities, and let $\mathbf{W}$ denote the observation indicator matrix with the entries to be one if the corresponding implied volatility in $\mathbf{V}$ is observed and to be zero otherwise. Analogous to the singular value decomposition of an arbitrary matrix, the low-rank matrix factorization model considers approximating the target matrix with a few latent factors, $\mathbf{a}_r$'s and $\mathbf{b}_r$'s, through optimizing the following penalized loss function:

$$\min_{\mathbf{a}_r's,\mathbf{b}_r's} \left\| \mathbf{W} \circ \left(\mathbf{V} - \sum_{r=1}^{K} \mathbf{a}_r \otimes \mathbf{b}_r \right) \right\|_F^2 + \lambda \sum_{r=1}^{K} (\|\mathbf{a}_r\|_2^2 + \|\mathbf{b}_r\|_2^2), \tag{17}$$

where "$\circ$" denotes the point-wise Hadamard product, "$\otimes$" denotes the outer product, and $\|\cdot\|_F$ denotes the Frobenius norm and $\|\cdot\|_2$ denotes the L_2 norm. Note that the vectors

[8] See Jain et al. (2013) and Hastie et al. (2015) for details.

$\mathbf{a}_r = (a_{r1}, \ldots, a_{r200})'$ and $\mathbf{b}_r = (b_{r1}, \ldots, b_{r100})'$ are essentially the latent row- and column-factors corresponding to moneyness level and maturity, respectively. Any missing entry of $\mathbf{V}$ can be imputed as $\widehat{V}_{ij} = \sum_{r=1}^{K} \hat{a}_{ri}\hat{b}_{rj}$, for the ith row, the jth column. The rank K and the tuning parameter λ can be pre-selected based on the cross-validation method. Dai et al. (2019) discuss extensively on the implementation. Following their analysis, we choose $K = 1$ in our context.

We assume that the implied volatility surface is smooth, and further embed smoothing into the matrix completion by posing a spline structure onto the latent factors. In particular, we consider the cubic spline model as

$$a_{ri} \approx \sum_{m=1}^{M} \beta_{rm} h_m(x_i), \quad b_{rj} \approx \sum_{m=1}^{M} \bar{\beta}_{rm} h_m(t_j), \tag{18}$$

where $h_m(\cdot)$'s are spline basis functions, x_i and t_j are corresponding moneyness level and maturity value, respectively. By combing (17) and (18), we carry out a smoothing-embedded matrix completion which yields an extrapolation of implied volatility surface to certain moneyness level and maturity value. We then use these extrapolated implied volatilities to obtain synthetic option prices, and estimate jumps as well as state prices. With the above procedure, we are able to recover information for all of the 176 meetings during 1996 to 2017.

References

Ai, H. and Bansal, R. (2018), "Risk preferences and the macroeconomic announcement premium," *Econometrica*, 86, 1383–1430.

Ai, H., Bansal, R., Guo, H., and Yaron, A. (2019a), "Identifying preference for early resolution from asset prices," .

Ai, H., Han, L. J., Pan, X. N., and Xu, L. (2019b), "The Cross-section of monetary policy announcement premium," *Lai, The Cross-Section of Monetary Policy Announcement Premium (February 13, 2019)*.

Amengual, D. and Xiu, D. (2018), "Resolution of policy uncertainty and sudden declines in volatility," *Journal of Econometrics*, 203, 297–315.

An, B.-J., Ang, A., Bali, T. G., and Cakici, N. (2014), "The joint cross section of stocks and options," *The Journal of Finance*, 69, 2279–2337.

Bernanke, B. S. and Kuttner, K. N. (2005), "What explains the stock market's reaction to Federal Reserve policy?" *The Journal of finance*, 60, 1221–1257.

Bliss, R. R. and Panigirtzoglou, N. (2004), "Option-implied risk aversion estimates," *The journal of finance*, 59, 407–446.

Breeden, D. T. and Litzenberger, R. H. (1978), "Prices of state-contingent claims implicit in option prices," *Journal of business*, 621–651.

Dai, B., Wang, J., Shen, X., and Qu, A. (2019), "Smooth neighborhood recommender systems," *The Journal of Machine Learning Research*, 20, 589–612.

Hastie, T., Mazumder, R., Lee, J. D., and Zadeh, R. (2015), "Matrix completion and low-rank SVD via fast alternating least squares," *The Journal of Machine Learning Research*, 16, 3367–3402.

Hong, Y., Tu, J., and Zhou, G. (2007), "Asymmetries in stock returns: Statistical tests and economic evaluation," *The Review of Financial Studies*, 20, 1547–1581.

Jain, P., Netrapalli, P., and Sanghavi, S. (2013), "Low-rank matrix completion using alternating minimization," in *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pp. 665–674.

Jensen, C. S., Lando, D., and Pedersen, L. H. (2019), "Generalized recovery," *Journal of Financial Economics*, 133, 154–174.

Kadan, O. and Manela, A. (2019), "Estimating the value of information," *The Review of Financial Studies*, 32, 951–991.

Kadan, O. and Tang, X. (2020), "A bound on expected stock returns," *The Review of Financial Studies*, 33, 1565–1617.

Lucca, D. O. and Moench, E. (2015), "The pre-FOMC announcement drift," *The Journal of Finance*, 70, 329–371.

Martin, I. (2017), "What is the Expected Return on the Market?" *The Quarterly Journal of Economics*, 132, 367–433.

Martin, I. W. and Wagner, C. (2019), "What is the Expected Return on a Stock?" *The Journal of Finance*, 74, 1887–1929.

Mehra, R. and Prescott, E. C. (1985), "The equity premium: A puzzle," *Journal of monetary Economics*, 15, 145–161.

Neuhierl, A. and Weber, M. (2016), "Monetary policy and the stock market: Time-series evidence," Tech. rep., National Bureau of Economic Research.

Roll, R. (1988), "The international crash of October 1987," *Financial analysts journal*, 44, 19–35.

Ross, S. (2015), "The recovery theorem," *The Journal of Finance*, 70, 615–648.

Savor, P. and Wilson, M. (2013), "How much do investors care about macroeconomic risk? Evidence from scheduled economic announcements," *Journal of Financial and Quantitative Analysis*, 48, 343–375.

— (2014), "Asset pricing: A tale of two days," *Journal of Financial Economics*, 113, 171–201.

Schwert, G. W. (1990), "Stock volatility and the crash of' 87," *The review of financial studies*, 3, 77–102.

Tang, X. (2018), "Variance asymmetry managed portfolios," *Working Paper. Available at SSRN 3108007.*

Vähämaa, S. and Äijö, J. (2011), "The Fed's policy decisions and implied volatility," *Journal of Futures Markets*, 31, 995–1010.

Vissing-Jørgensen, A. and Attanasio, O. P. (2003), "Stock-market participation, intertemporal substitution, and risk-aversion," *American Economic Review*, 93, 383–391.

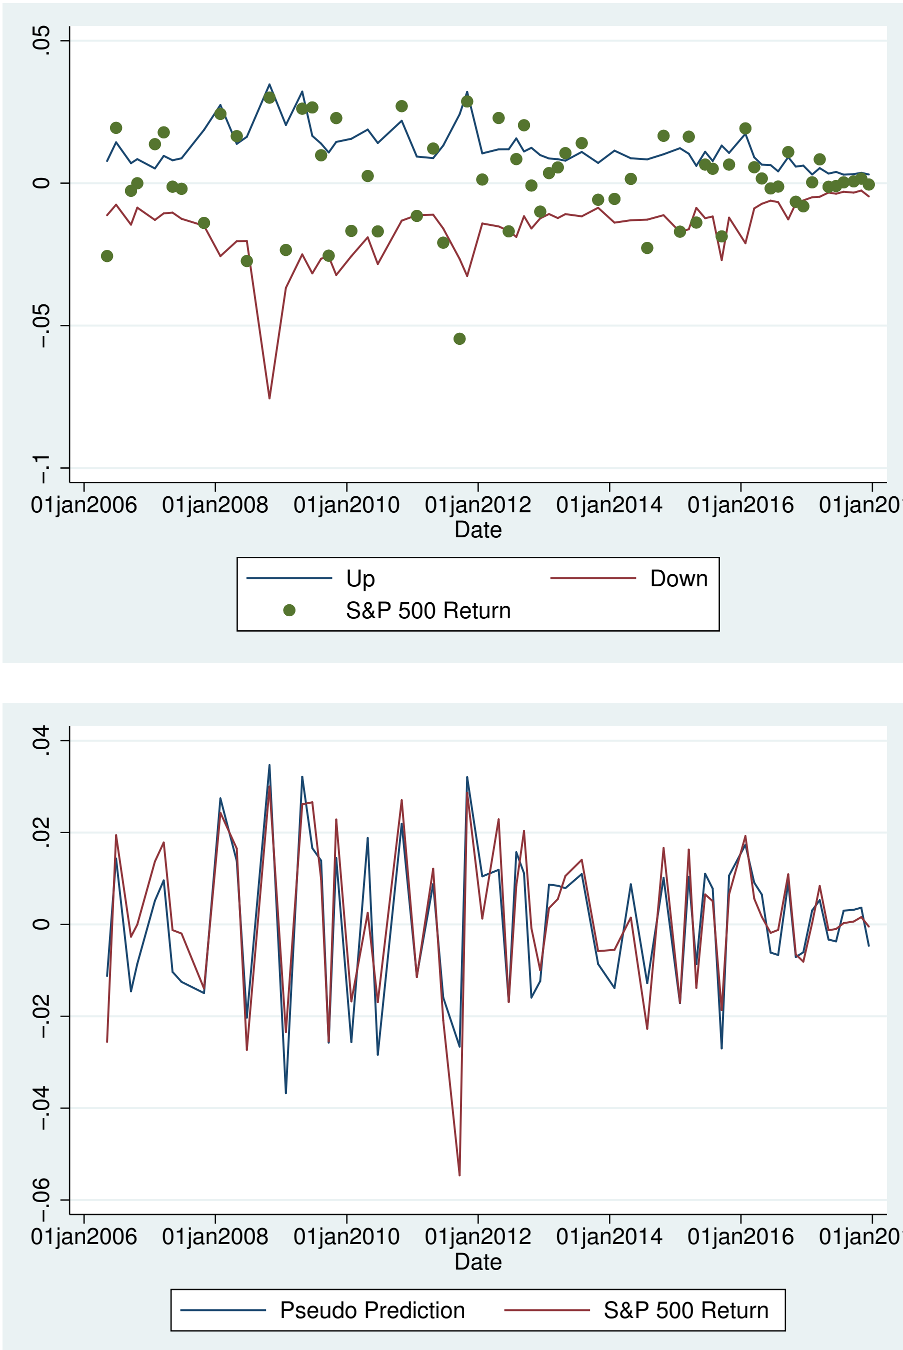

Figure 1: Jump Size Estimates with Real Data

Notes: The top panel presents the time-series of the jump size estimations as well as the realized S&P 500 returns. The bottom panel presents the time-series of pseudo prediction of the S&P 500 returns by (14) and the realized S&P 500 returns. The sample includes options written on the S&P 500 index with a life span shorter than three days and those cover the 24-hour time window before the FOMC meeting between January 1996 and December 2017. Realized returns are over the life span of the corresponding options.

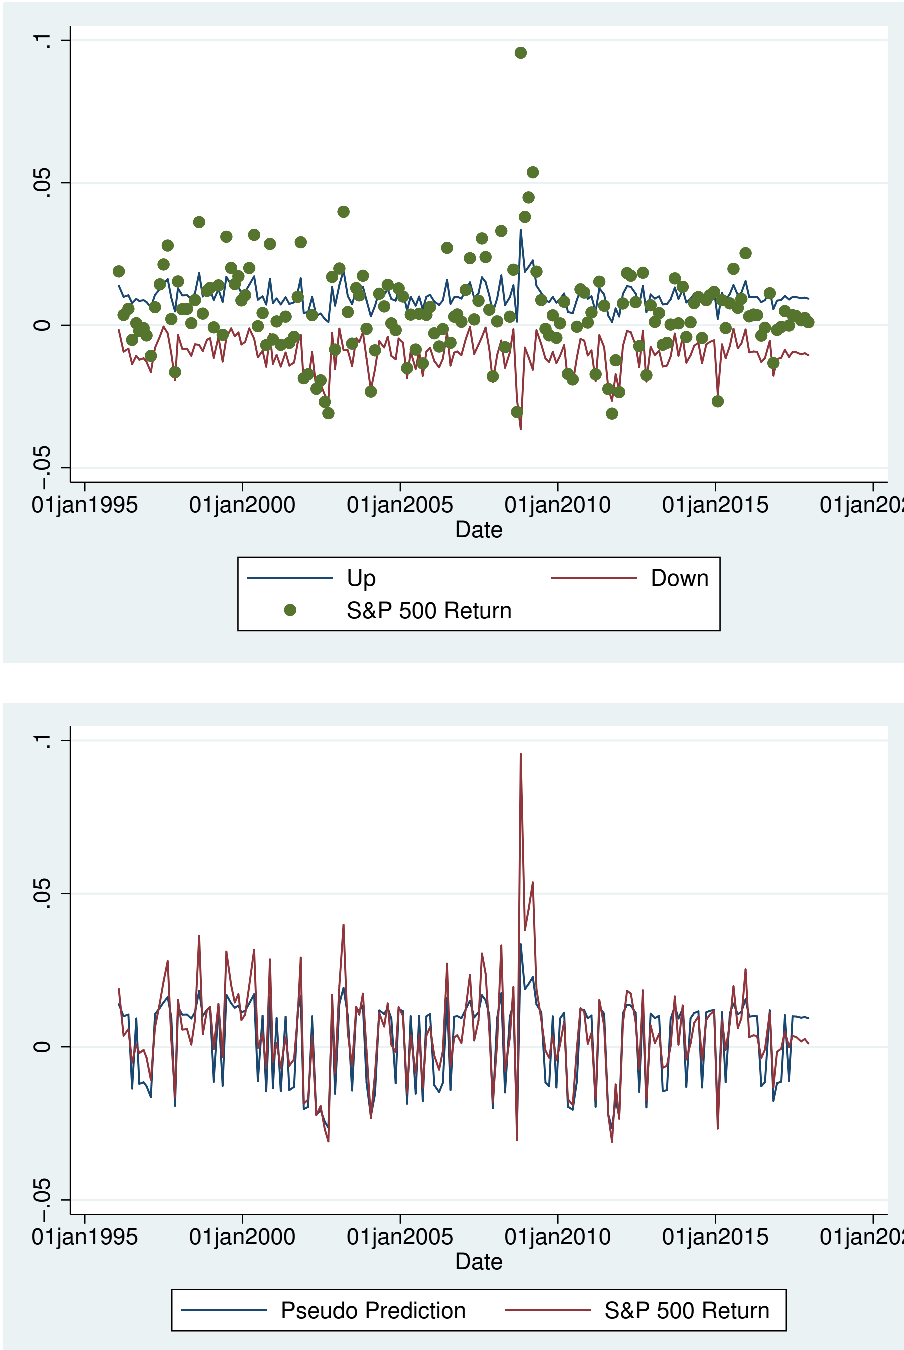

Figure 2: Jump Sizes Over Extended Sample Period

Notes: The top panel presents the time-series of the jump size estimations as well as the realized S&P 500 returns. The bottom panel presents the time-series of pseudo prediction of the S&P 500 returns by (14) and the realized S&P 500 returns. The sample includes all FOMC meetings between January 1996 and December 2017. Realized returns are returns of the &P 500 index over the 48-hour time window before the FOMC meeting.

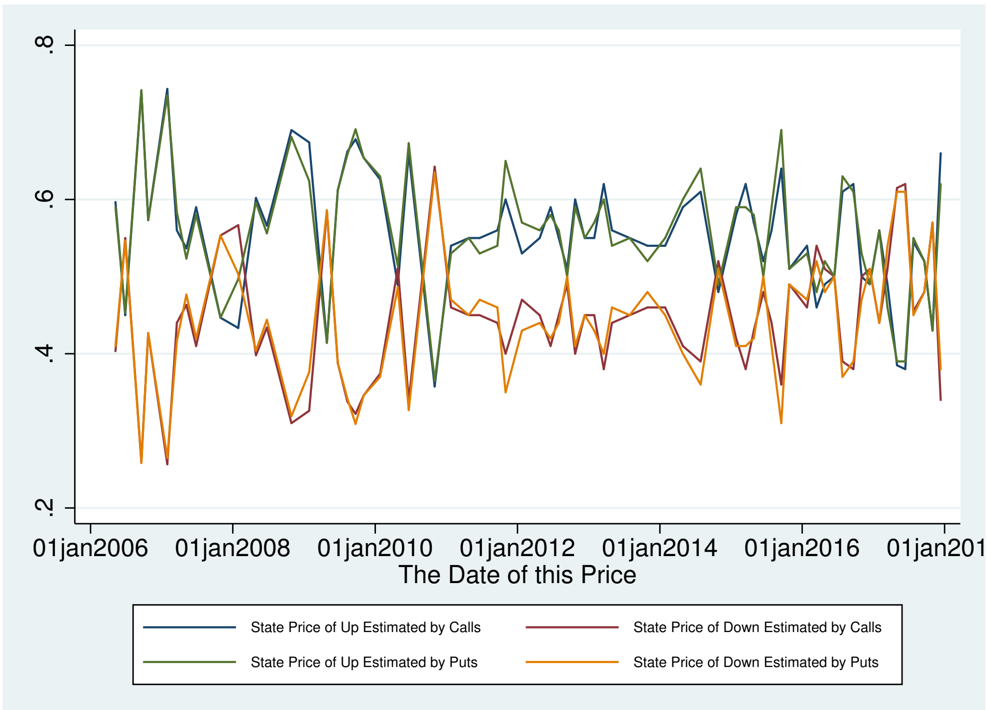

Figure 3: State Price Estimates with Real Data

Notes: This figure presents the time-series of the state price estimations for upward and downward jumps. The sample includes options written on the S&P 500 index with a life span shorter than three days and that covers the 24-hour time window before the FOMC meeting between January 1996 and December 2017. There are two sets of estimations based on a pair of calls and a pair of puts, respectively.

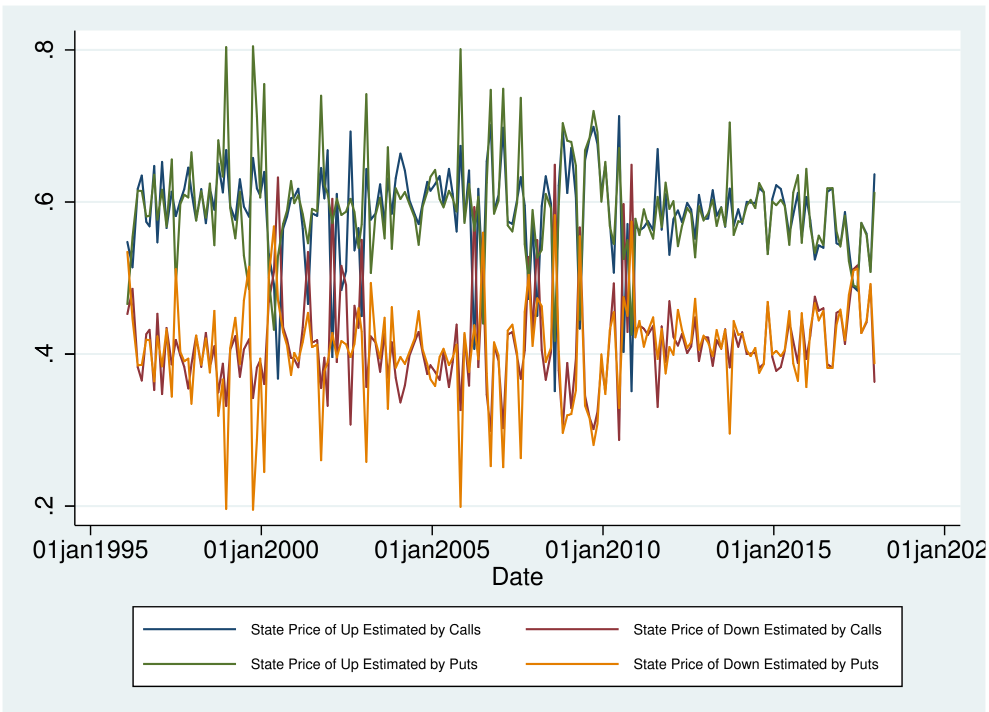

Figure 4: State Prices Over Extended Sample Period

Notes: This figure presents the time-series of the state price estimations for upward and downward jumps. The sample includes all FOMC meetings between January 1996 and December 2017. There are two sets of estimations by a pair of calls and a pair of puts, respectively.

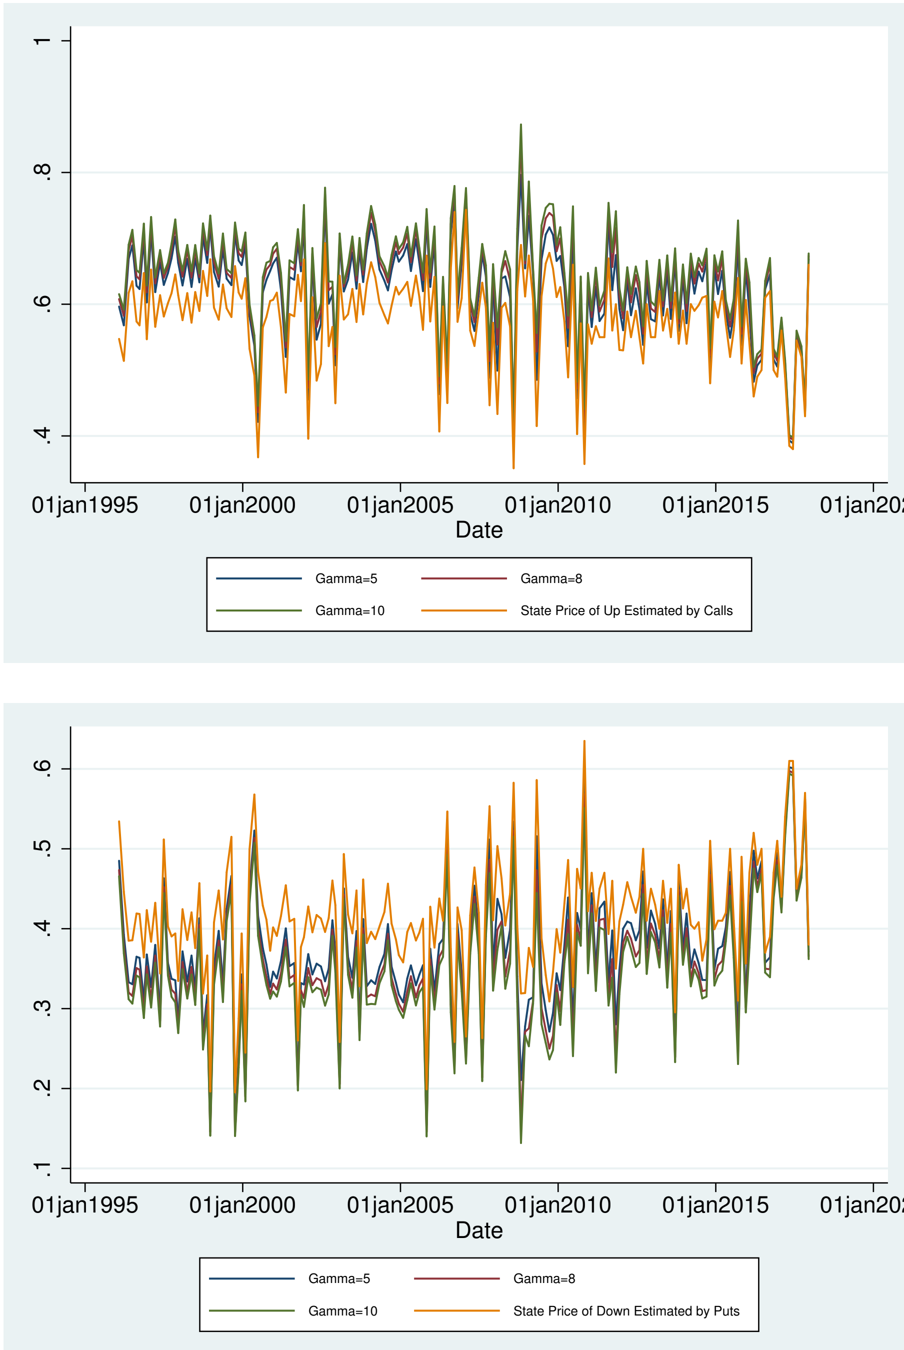

Figure 5: Physical Probabilities of Jumps

Notes: This figure presents the time-series of the estimation of physical probabilities for upward jumps following (11) in the top panel and downward jumps following (12) in the bottom panel, as well as the corresponding state prices. We choose the relative risk aversion levels from 5 to 10. The sample includes all FOMC meetings between January 1996 and December 2017.

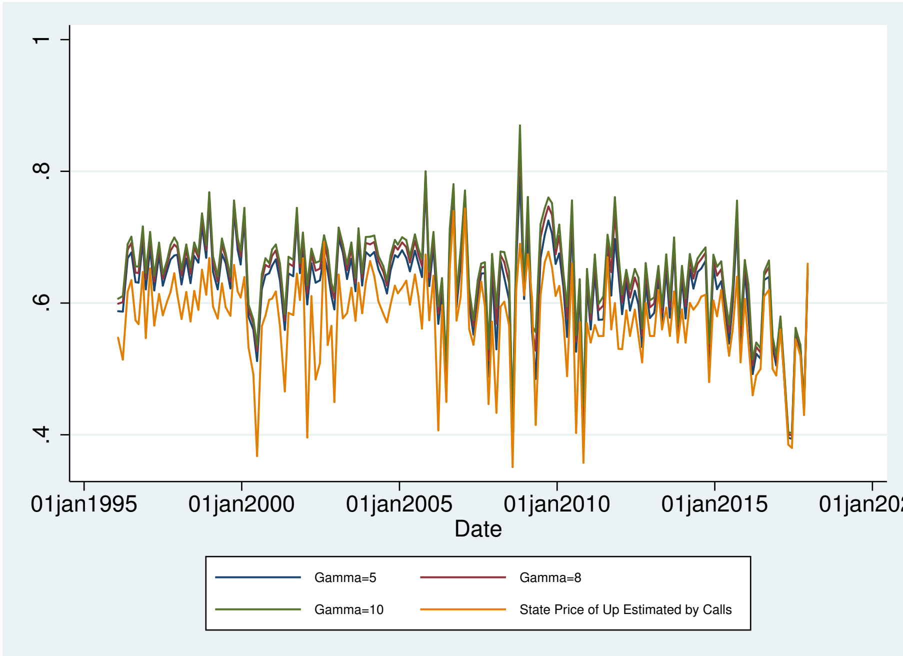

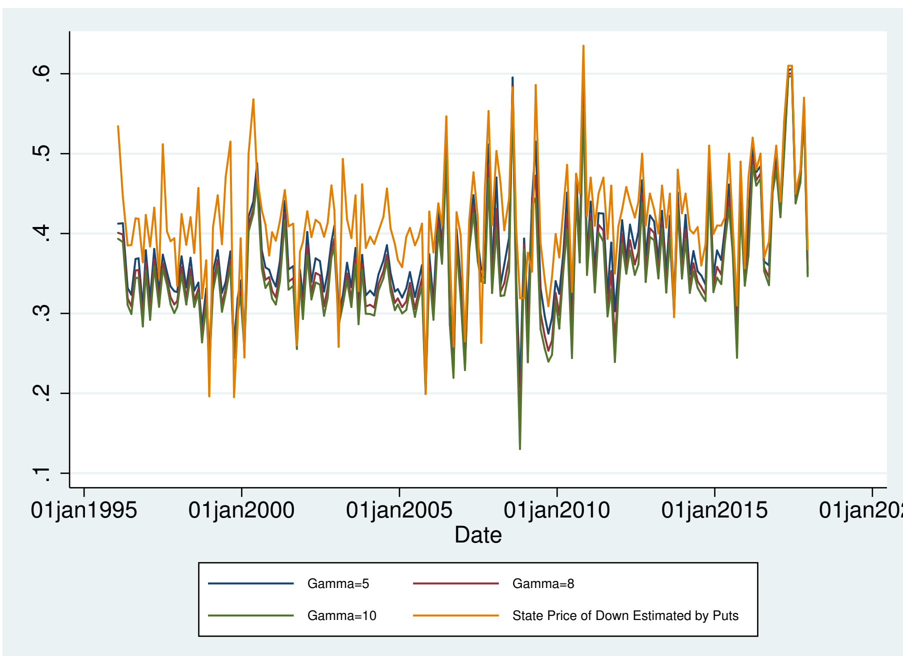

Figure 6: Weighted Physical Probabilities of Jumps

Notes: This figure presents the time-series of the estimation of physical probabilities for upward and downward jumps following (13) as well as their corresponding state prices. We choose the relative risk aversion levels from 5 to 10. The sample includes all FOMC meetings between January 1996 and December 2017.

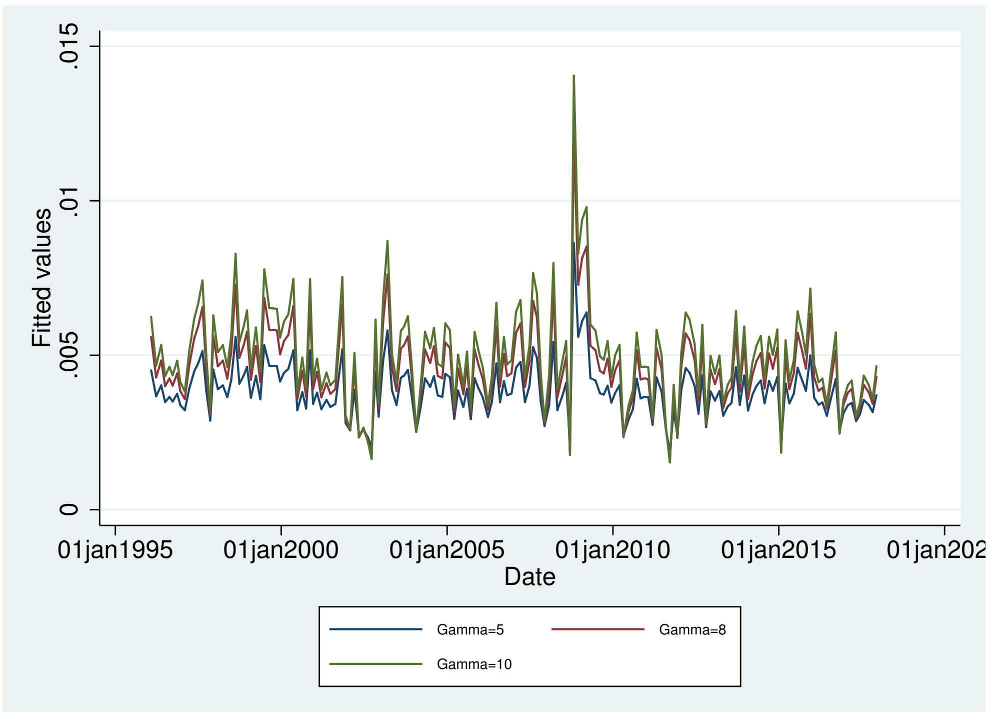

Figure 7: The FOMC Risk Premium

Notes: This figure presents the time-series of estimation of the FOMC risk premium following (15). We choose the relative risk aversion levels from 5 to 10. The sample includes all FOMC meetings between January 1996 and December 2017.

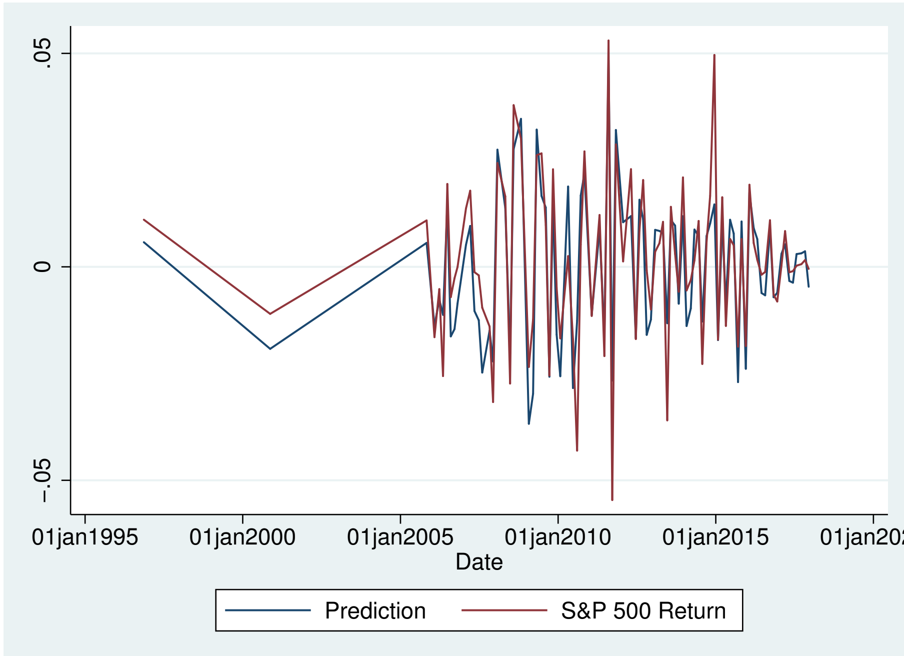

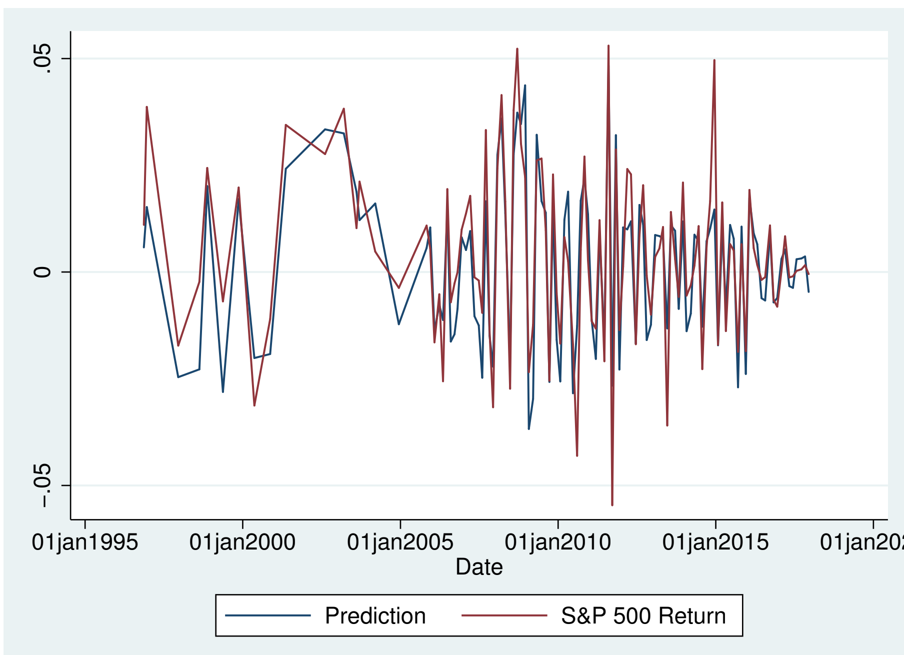

Figure 8: Robustness Check – Maximum Maturity

Notes: This figure presents the time-series of pseudo prediction of the S&P 500 returns by (14) and the realized S&P 500 returns. The sample includes options written on the S&P 500 index with a life span shorter than four (top panel) or five (bottom panel) days and that covers the 24-hour time window before the FOMC meeting between January 1996 and December 2017. Realized returns are over the life span of the corresponding options.

Table 1: Summary Statistics of Jump Estimation

This table reports summary statistics of the variables associated with jump estimation (Panel A) and the estimated jump sizes (Panel B and Panel C). The sample in Panels A and B includes options written on the S&P 500 index with a life span shorter than three days and that covers the 24-hour time window before the Federal Open Market Committee meeting between January 1996 and December 2017. The sample in Panel C includes all Federal Open Market Committee meetings from January 1996 to December 2017. On each day in our sample, we use or estimate the synthetic prices of two call and two put options that have strike prices closest to the S&P 500 index level, and follow (2) and (4) to estimate the index jump sizes.

Panel A: Summary Statistics of Related Variables

Variable	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
C_1	67	8.7776	4.0905	2.6	3.55	4.6	5.75	7.7	10.9	15.15	17.1	22.4
C_2	67	11.1220	8.2631	1.625	2	2.475	5.05	9.7	15.1	21.95	23.6	50.45
P_1	67	9.5299	4.3272	1.9	3.65	4.6	6.2	9.1	12.1	14.5	16.6	25.7
P_2	67	8.9813	5.2673	1.225	1.95	3.2	5.25	8.45	11.8	17.15	20.4	24.15
$\frac{K_1^C}{S_0}$	67	1.0008	0.0030	0.9926	0.9958	0.9983	0.9992	1.0001	1.0015	1.0054	1.0065	1.0101
$\frac{K_2^C}{S_0}$	67	0.9978	0.0099	0.9569	0.9798	0.9815	0.9969	0.9984	1.0022	1.0064	1.0112	1.0232
$\frac{K_1^P}{S_0}$	67	1.0008	0.0030	0.9926	0.9958	0.9983	0.9992	1.0001	1.0015	1.0054	1.0065	1.0101
$\frac{K_1^P}{S_0}$	67	0.9978	0.0099	0.9569	0.9798	0.9815	0.9969	0.9984	1.0022	1.0064	1.0112	1.0232
Maturity	67	2.4776	0.8413	1	1	1	2	3	3	3	3	3

Panel B: Summary Statistics of u and d

Variable	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
u	67	0.0116	0.0069	0.0030	0.0032	0.0040	0.0071	0.0102	0.0141	0.0204	0.0274	0.0347
d	67	0.0153	0.0110	0.0026	0.0033	0.0048	0.0087	0.0125	0.0190	0.0270	0.0322	0.0756

Panel C: Summary Statistics of u and d for the full sample

Variable	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
u	176	0.0103	0.0041	0.0012	0.0034	0.0048	0.0081	0.0100	0.0122	0.0145	0.0170	0.0228
d	176	0.0137	0.0080	0.0015	0.0032	0.0055	0.0090	0.0133	0.0148	0.0193	0.0222	0.0264

Table 2: Summary Statistics of State Price Estimation

This table reports summary statistics of estimated state prices. The sample in Panel A includes options written on the S&P 500 index with a life span shorter than three days and that covers the 24-hour time window before the Federal Open Market Committee meeting between January 1996 and December 2017. The sample in Panel B includes all Federal Open Market Committee meetings from January 1996 to December 2017. On each day in our sample, we use or stimate the synthetic prices of two call and two put options that have strike prices closest to the S&P 500 index level, and follow (3) and (5) to estimate the state prices of index jumps. The table reports summary statistics of two pairs of estimates: $(\pi_u, 1-\pi_u)$ based on call options, and $(1-\pi_d, \pi_d)$ based on put options.

Panel A: Sub-Sample with Real Data

Variable	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
π_u	67	0.5561	0.0803	0.3575	0.415	0.4467	0.51	0.55	0.61	0.66	0.678	0.7433
$1-\pi_u$	67	0.4439	0.0803	0.2567	0.322	0.34	0.39	0.45	0.49	0.5533	0.585	0.6425
$1-\pi_d$	67	0.5585	0.0785	0.365	0.414	0.4533	0.514	0.56	0.6	0.658	0.69	0.7417
π_d	67	0.4415	0.0785	0.2583	0.31	0.342	0.4	0.44	0.486	0.5467	0.586	0.635

Panel B: Full Sample with Extrapolated Data

Variable	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
π_u	176	0.5951	0.0618	0.3510	0.4658	0.5306	0.5801	0.6024	0.6224	0.6578	0.6737	0.7196
$1-\pi_u$	176	0.4049	0.0618	0.2804	0.3263	0.3422	0.3776	0.3976	0.4199	0.4694	0.5342	0.6490
$1-\pi_d$	176	0.6031	0.0631	0.4271	0.5007	0.5430	0.5786	0.5967	0.6160	0.6812	0.7398	0.8037
π_d	176	0.3969	0.0631	0.1963	0.2602	0.3188	0.3840	0.4033	0.4214	0.4570	0.4993	0.5729

Table 3: Summary Statistics of Physical Probability Estimation

This table reports summary statistics of estimated physical probabilities. The sample includes all Federal Open Market Committee meetings from January 1996 to December 2017. On each day in our sample, we use the observed prices (if available) or estimate the synthetic prices of two call and two put options that have strike prices closest to the S&P 500 index level, and follow (11) and (12) to estimate the physical probabilities of index jumps based on results reported in Tables 1 and 2. We set levels of relative risk aversion (γ) to 5, 8 and 10.

Panel A: Physical Probability of Upward Jumps

γ	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
5	176	0.6205	0.0759	0.3929	0.4636	0.5072	0.5831	0.6343	0.6687	0.7026	0.7221	0.7603
8	176	0.6356	0.0767	0.3977	0.4799	0.5249	0.5973	0.6488	0.6817	0.7170	0.7385	0.7720
10	176	0.6455	0.0773	0.4009	0.4908	0.5334	0.6057	0.6567	0.6901	0.7288	0.7506	0.7863

Panel B: Physical Probability of Downward Jumps

γ	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
5	176	0.3717	0.0784	0.1533	0.2381	0.2804	0.3324	0.3645	0.4125	0.4708	0.5157	0.6009
8	176	0.3567	0.0789	0.1456	0.2265	0.2662	0.3176	0.3504	0.3988	0.4547	0.4863	0.5953
10	176	0.3468	0.0794	0.1401	0.2190	0.2488	0.3071	0.3401	0.3901	0.4447	0.4794	0.5916

Table 4: Summary Statistics of Weighted Physical Probability Estimation

This table reports summary statistics of weighted physical probability estimations. The sample includes all Federal Open Market Committee meetings from January 1996 to December 2017. On each day in our sample, we follow (13) to estimate the physical probabilities of index jumps based on results reported in Table 3. We set levels of relative risk aversion (γ) to 5, 8 and 10.

Panel A: Physical Probability of Upward Jumps

γ	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
5	176	0.6246	0.0698	0.3954	0.4885	0.5298	0.5882	0.6365	0.6684	0.6943	0.7253	0.7830
8	176	0.6397	0.0708	0.4002	0.5137	0.5423	0.6054	0.6507	0.6816	0.7075	0.7397	0.7933
10	176	0.6496	0.0716	0.4034	0.5206	0.5559	0.6136	0.6607	0.6906	0.7191	0.7559	0.8001

Panel B: Physical Probability of Downward Jumps

γ	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
5	176	0.3754	0.0698	0.2170	0.2747	0.3057	0.3316	0.3635	0.4118	0.4702	0.5115	0.6046
8	176	0.3603	0.0708	0.2067	0.2603	0.2925	0.3184	0.3493	0.3946	0.4577	0.4863	0.5998
10	176	0.3504	0.0716	0.1999	0.2441	0.2809	0.3094	0.3393	0.3864	0.4441	0.4794	0.5966

Table 5: Summary Statistics of the FOMC Risk Premium Estimation

This table reports summary statistics of the recovered FOMC risk premium. The sample in Panel A includes all Federal Open Market Committee meetings from January 1996 to December 2017. The sample in Panel B includes options written on the S&P 500 index with a life span shorter than three days and that covers the 24-hour time window before the Federal Open Market Committee meeting between January 1996 and December 2017. On each day in our sample, we use the observed prices (if available) or estimate the synthetic prices of two call and two put options that have strike prices closest to the S&P 500 index level, and follow (15) to estimate the risk premium of index jumps. We set levels of relative risk aversion (γ) to 5, 8 and 10. We also report summary statistics of realized returns of the S&P 500 index in the last row of each panel.

Panel A: Full Sample with Extrapolated Data

γ	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
5	176	0.0038	0.0010	0.0020	0.0025	0.0029	0.0034	0.0038	0.0043	0.0047	0.0053	0.0064
8	176	0.0046	0.0013	0.0018	0.0025	0.0030	0.0038	0.0044	0.0053	0.0060	0.0068	0.0085
10	176	0.0050	0.0016	0.0016	0.0025	0.0031	0.0041	0.0049	0.0059	0.0067	0.0077	0.0098
Realized Returns	176	0.0047	0.0160	-0.0309	-0.0223	-0.0165	-0.0036	0.0036	0.0122	0.0213	0.0311	0.0537

Panel B: Sub-Sample with Real Data

γ	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
5	67	0.0012	0.0012	0.0001	0.0002	0.0003	0.0005	0.0007	0.0015	0.0025	0.0036	0.0125
8	67	0.0019	0.0019	0.0001	0.0002	0.0004	0.0005	0.0011	0.0023	0.0033	0.0055	0.0172
10	67	0.0023	0.0023	0.0001	0.0003	0.0005	0.0007	0.0014	0.0029	0.0040	0.0068	0.0203
Realized Returns	67	0.0014	0.0164	-0.0546	-0.0255	-0.0209	-0.0081	0.0013	0.0137	0.0229	0.0266	0.0300

Table 6: Robustness Test – Maximum Maturity

This table reports summary statistics of estimated jumps.The sample includes options written on the S&P 500 index with a life span shorter than four (Panel A) or five (Panel B) days and that covers the 24-hour time window before the FOMC meeting between January 1996 and December 2017. On each day in our sample, we select two call and two put options that have strike prices closest to the S&P 500 index level, and follow (2) and (4) to estimate the index jump sizes.

Panel A: Maximum Maturity: 4 Days

Variable	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
u	89	0.0125	0.0079	0.0030	0.0034	0.0052	0.0078	0.0106	0.0146	0.0224	0.0310	0.0484
d	89	0.0162	0.0115	0.0026	0.0037	0.0061	0.0106	0.0132	0.0192	0.0270	0.0322	0.0756

Panel B: Maximum Maturity: 5 Days

Variable	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
u	114	0.0142	0.0089	0.0030	0.0036	0.0058	0.0084	0.0119	0.0171	0.0274	0.0334	0.0437
d	114	0.0177	0.0119	0.0030	0.0047	0.0071	0.0111	0.0146	0.0222	0.0297	0.0350	0.0703

Table 7: Robustness Test – Excluding Financial Crisis

This table reports summary statistics of estimated jumps and the FOMC risk premium excluding years 2008 and 2009. The sample includes all Federal Open Market Committee meetings from January 1996 to December 2017. On each day in our sample, we estimate the synthetic prices of two call and two put options that have strike prices closest to the S&P 500 index level, and follow (2) and (4) to estimate the index jump sizes (reported in Panel A). We follow (15) to estimate the FOMC risk premium (reported in Panel B). We set levels of relative risk aversion (γ) to 5, 8 and 10. We also report summary statistics of realized returns of the S&P 500 index in the last row.

Panel A: Jump Sizes

Variable	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
u	160	0.0103	0.0033	0.0022	0.0042	0.0061	0.0085	0.0101	0.0122	0.0142	0.0164	0.0192
d	160	0.0100	0.0055	0.0011	0.0018	0.0028	0.0061	0.0094	0.0130	0.0194	0.0210	0.0230

Panel B: FOMC Risk Premium

γ	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
5	160	0.0033	0.0007	0.0015	0.0020	0.0023	0.0029	0.0033	0.0038	0.0042	0.0046	0.0051
8	160	0.0039	0.0011	0.0013	0.0020	0.0025	0.0033	0.0039	0.0047	0.0053	0.0060	0.0068
10	160	0.0044	0.0014	0.0011	0.0020	0.0025	0.0035	0.0043	0.0053	0.0060	0.0068	0.0078
Realized Returns	160	0.0035	0.0133	-0.0309	-0.0208	-0.0168	-0.0036	0.0036	0.0113	0.0199	0.0276	0.0362

Table 8: Relation to VIX

This table reports the coefficient estimates of regression in (16). The FOMC risk premiums are based on relative risk aversion of 5, 8, and 10. We consider both the whole sample (Columns (1) to (3)) during 1996 to 2017 and the 67 meetings estimated with real option data (Columns (4) to (6)).

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	The FOMC Risk Premium Based On Relative Risk Aversion					
	5	8	10	5	8	10
VIX	0.0031	0.0050	0.0061	0.0050	0.0081	0.0100
	(0.0007)***	(0.0011)***	(0.0014)***	(0.0010)***	(0.0016)***	(0.0019)***
$Intercept$	0.0032	0.0035	0.0037	0.0028	0.0029	0.0029
	(0.0002)***	(0.0003)***	(0.0003)***	(0.0002)***	(0.0003)***	(0.0004)***
Adj R-Squared	0.0896	0.0953	0.0992	0.2513	0.2726	0.2865
#Obs	176	176	176	67	67	67