Reading between the rates: how options price FED uncertainty?

Łukasz Adamski & Robert Ślepaczuk

Quantitative Finance Research Group Seminar

Faculty of Economic Sciences, University of Warsaw, 25 May 2026

Agenda

1. Motivation and objective
2. Contributions
3. Literature review
4. Research questions and hypotheses
5. Data overview
6. Methodology:
    - Interpolation/extrapolation of IV surface
    - Forecasting task
    - Convolutional 2D LSTM
7. Results:
    - Statistical analysis
    - Forecasting
8. Conclusions
9. Limitations
10. Future research

Motivation and objective

Investors anticipate signals given by economic data announcements which translates into elevated uncertainty

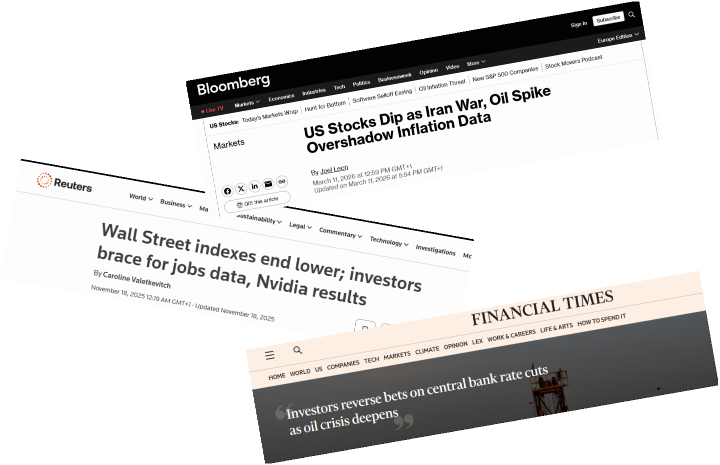

Aims

This uncertainty should be especially relevant for derivatives market which are forward-looking instruments by nature.

As such, the general objective of the paper is to **understand and predict** the pre-announcement effect of elevated volatility in derivatives markets before the economic announcements (the focus is put on announcements of interest rate policy by Fed).

There are two main aspects to this analysis:

- Empirical - statistical description of the pre-announcement effect
- Forecasting - quantitative prediction of the pre-announcement effect

Contributions

Our contribution mainly lies in:

- Quantitative prediction of the pre-announcement effect of elevated implied volatility in derivatives markets
- The inclusion of exogenous information inside Machine learning framework in IV surface forecasting
- A much more comprehensive and explainable diagnostic analysis of the models.
- Application of ML models directly on the IV surface interpolated and exrapolated to a fixed grid

Literature review

Stylised facts regarding pre-announcement effects

- Increased uncertainty raises the undiversifiable risk, which induces a **pre-announcement premium**.
- This premium accounts for as much as 11% excess annualised equity premium in the US market Savor and Wilson [8]. The most tangible effect is for interest rates announcement by Fed as shown in Lucca and Moench [7].
- ATM Implied volatility related to options on equities (S&P 500 index), interest rates and commodities rises before the announcements and drops after the announcement.
- Short-dated options react more heavily (declining term-structure of the pre-announcement effect).

The pre-announcement effect for S&P 500

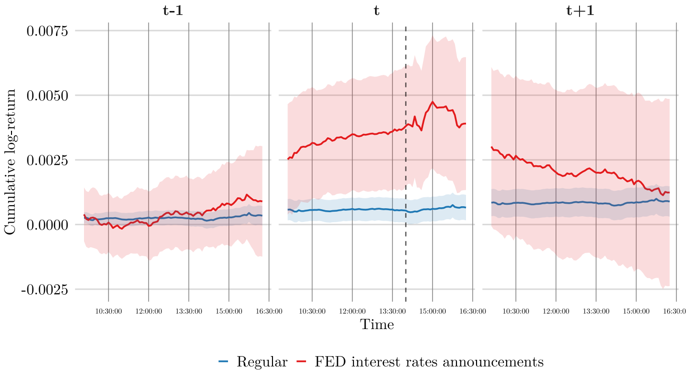

Figure: The cumulative return on S&P 500 index during the 3-day Fed interest rate announcement window compared to the average cumulative return across the period 2004-2025. The ribbon depicts a standard 95% confidence interval. The plot is inspired by Lucca and Moench [7].

Identified gaps

- The literature on derivatives in terms of a (pre-) announcement effect is limited, it focuses on analysing either only ATM implied volatility (the evolution of only one point on the surface) or VIX.
- Models used are rather simple and suited mainly for statistical analysis.
- Lack of trials to quantitatively forecast the pre-announcement premium.
- The frequency of observations is usually low (daily).
- Machine learning framework is underutilised in forecasting implied volatility and the models often lack exogenous information
- Models are often superficially diagnosed and analysed

Research questions and hypotheses

Research questions

1. How does implied volatility surface evolve before and during the announcement of interest rate policy in the US?
2. To what extent machine learning (ML) models can capture the abnormal evolution of implied volatility surface around policy announcement days?
3. Can the inclusion of the information regarding the announcement date meaningfully improve the performance of predictive models?

Hypotheses

1. The implied volatility monotonically increases prior to the announcement.
2. The effect is stronger for short dated options.
3. The effect is stronger in high volatility regimes.
4. The effect is stronger for OTM options.
5. ML framework beats the benchmark random walk in prediction power both globally and during policy announcement days.
6. A better model fit can be achieved for ATM implied volatility and longer maturities.
7. Adding the date of FOMC meetings as a feature improves overall model performance and during policy annnoucement days.

Data overview

Data

The main data for this analysis was extracted from **Chicago Board Options Exchange (CBOE) market data service** Chicago Board Options Exchange [2]. It covers price quotes of options on the S&P 500 index in one minute intervals during the daily trading session between 9:30am to 4:30pm EST enriched by the price of the underlying instrument - here the S&P 500 index. Reliable data with sufficiently granular set of maturities spans 2016-2025. Training set ends with year 2023.

Moreover, a calendar with scheduled data releases was constructed based on the Fed website Board of Governors of the Federal Reserve System [1].

For statistical analysis the data in five minute intervals is used. For forecasting daily data at 2pm EST is used.

No-aribtrage constraints

To filter incorrect data points a set of filters and no-arbitrage conditions have been applied to options data, namely:

- Bid price lower or equal than ask price
- Bid price higher than 0
- Moneyness higher than 0
- Implied volatility higher than 0
- Open Interest higher than 0
- Convexity and monotonicity in strike

Methodology

Implied volatility (IV) surface

The main object of our study is implied volatility which is a volatility such that the Black-Scholes formula produces a price consistent with the market price $C^*(K,T)$ for strike K and maturity T:

$$\sigma^{IV}:(K,T)\rightarrow\sigma^{IV}(K,T)\quad \text{s.t.}\quad C_{\mathsf{BS}}(C,T,\sigma^{\mathsf{IV}}(K,T))=C^*(K,T)$$

Of course we observe only a **discrete** number of strikes and maturities which also change over time, resulting in a sparse grid.

IV surface construction

There are a few approaches to overcome this issue. Following the literature, we construct IV surface using three methods:

- Linear interpolation representing a non-parametric approach
- Stochastic Volatility Inspired (SVI) model representing a parametric approach (SVI - Gatheral and Jacquier [4])
- Ad-hoc-Black-Scholes representing a global, simple parametric approach (Dumas et al. [3])

Notation

Denote the set of maturities $\mathbb{T}_t = \{T_1, T_2, \dots, T_J\}$ and the set of strikes available per fixed maturity $T_{t,j}$ by $\mathbb{K}_t^j = \{K_1, K_2, \dots, K_I\}$ both observed at time t. These describe the real data points.

Then we need the fixed grid which will serve as an input to models. Denote the grid of fixed strikes and maturities as $\mathbb{K}^* \times \mathbb{T}^*$, where $\mathbb{K}^* = \{K_1^*, \dots, K_N^*\}$ lies in the range of $[-0.2, 0.2]$ such that it has a length of 40 with point 0 concatenated and $\mathbb{T}^* = \{T_1^*, \dots, T_M^*\} = \{0, 1, \dots, 20\}$. As such we use $M = 21$ maturities and $N = 41$ strikes.

Also define $k = \ln\left(\frac{S_t e^{rT}}{K}\right)$ as log-forward-moneyness.

Linear interpolation

Serving as a baseline linear interpolation in both log-moneyness and maturity (separately) is the most straightforward approach. For each t we begin with interpolating the surface per fixed maturity $T_{t,j}$. Therefore the fitted IV at $(K_n^*, T_{t,j})$ with $K_{t,i-1} \leq K_n^* \leq K_{t,i}$ is equal to:

$$\sigma_t^{\mathsf{IV}}(K_n^*, T_{t,j}) = \frac{\sigma_t^{\mathsf{IV}}(K_{i-1}, T_{t,j})\,(K_{t,i} - K_{t,i-1}) + \sigma_t^{\mathsf{IV}}(K_i, T_{t,j})\,(K_n^* - K_{t,i-1})}{K_{t,i} - K_{t,i-1}}$$

Then we proceed with interpolating the IV surface in maturity per fitted strikes. Therefore for a fixed $K_n^* \in \mathbb{K}^*$ the fitted IV at (K_n^*, T_m^*) with $T_{t,j-1} \leq T_m^* \leq T_{t,j}$ is equal to:

$$\sigma_t^{\mathsf{IV}}(K_n^*, T_m^*) = \frac{\sigma_t^{\mathsf{IV}}(K_n^*, T_{t,j-1})\,(T_{t,j} - T_{t,j-1}) + \sigma_t^{\mathsf{IV}}(K_n^*, T_{t,j})\,(T_m^* - T_{t,j-1})}{T_{t,j} - T_{t,j-1}}$$

Linear interpolation - wings

If the strikes at any given time were available in the wider range than $[-0.2, 0.2]$ this procedure would produce interpolated IV surface $\hat{\sigma}_t^{\text{IV}}(K_n^*, T_m^*)$ for all pairs $(K_n^*, T_m^*) \in \mathbb{K}^* \times \mathbb{T}^*$.

However, at some dates the data especially in the in-the-money wing is scarce. Because it has been proven in Lee [6] that wings of implied volatility surface are always bounded by a linear function we extrapolate wings linearly using a slope estimate:

$$\Delta\sigma_t^{\text{IV}}(K_I^j, T_{t,j}) = \frac{\sigma_t^{\text{IV}}(K_I^j, T_{t,j}) - \sigma_t^{\text{IV}}(K_{I-2}^j, T_{t,j})}{K_I^j - K_{I-2}^j}$$

$$\Delta\sigma_t^{\text{IV}}(K_1^j, T_{t,j}) = \frac{\sigma_t^{\text{IV}}(K_3^j, T_{t,j}) - \sigma_t^{\text{IV}}(K_1^j, T_{t,j})}{K_1^j - K_3^j}$$

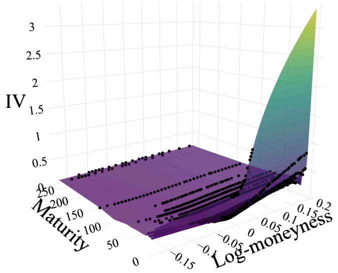

Figure: The implied volatility surface fitted using linear interpolation and linear extrapolation of wings - call options on S&P 500 index as observed at 2pm on 16 May 2024.

Ad-hoc Black Scholes (AHBS)

This approach relies on fitting the polynomial function in strike and moneyness. The surface is described by:

$$\sigma^{\mathsf{IV}}(k,T) = \alpha_0 + \alpha_1 k + \alpha_2 k^2 + \alpha_3 T + \alpha_4 T^2 + \alpha_5 kT + \varepsilon$$

This function can be fitted using least squares at each time t using the whole grid $\bigcup_{j=1}^{J} K_t^j \times \{T_{t,j}\}$ as input points. Then the IV can be evaluated at all points $K^* \times T^*$.

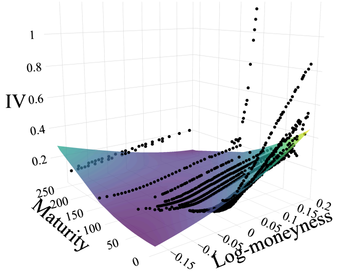

Figure: The implied volatility surface with fitted AHBS model - call options on S&P 500 index as observed at 2pm on 16 May 2024.

Stochastic Volatility Inspired (SVI)

The approach introduced in Gatheral and Jacquier [4] is also based on a parametric approach. Define as $w(k,T) = \sigma^{\mathsf{IV}}(k,T)T^*$ the total variance. The surface is described by a function:

$$w(k,T) = a + b\left(\rho(k-m) + \sqrt{(x-m)^2 + \sigma^2}\right)$$

where $a, m \in \mathbb{R}$, $b \geq 0$, $|\rho| \leq 1$ and $\sigma > 0$.

This function can be fitted using any optimising scheme at each time t using the whole grid $\bigcup_{j=1}^{J} K_t^j \times \{T_{t,j}\}$ as input points. Then the IV can be evaluated at all points $K^* \times T^*$ as $\sigma^{\mathsf{IV}}(k^*,T^*) = \sqrt{\frac{w(k^*,T^*)}{T^*}}$.

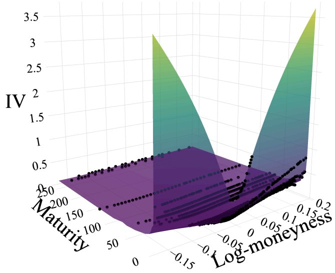

Figure: The implied volatility surface with fitted SVI model - call options on S&P 500 index as observed at 2pm on 16 May 2024.

Forecasting task

The task is to forecast the fixed grid IV surface $\sigma_t^{\mathsf{IV}}(K^*, T^*)$ using its past observations up to time $t - l$ in the first stage. In the second stage also exogeneous dummy variable X_t indicating the day of the FOMC conference which sets forth the monetary policy is introduced.

Therefore, we describe $\sigma_t^{\mathsf{IV}}(K^*, T^*)$ using a (nonlinear) function f:

$$\sigma_t^{\mathsf{IV}}(K^*, T^*) = f\left(\sigma_{t-1}^{\mathsf{IV}}(K^*, T^*), \ldots, \sigma_{t-l}^{\mathsf{IV}}(K^*, T^*), X_t\right) + \varepsilon_t$$

The goal of the machine learning models will be to fit f resulting in trained network $\hat{f}$.

h-step forecasts

To check the long-term performance of the model we forecast the IV for horizons $h = \{1, 2, 5, 10\}$. The forecasts are generated by passing the forecasted values as additional inputs. Therefore, h-step forecast is generated as:

$$\hat{\sigma}^{\mathsf{IV}}_{t+h}(K^*, T^*) = \hat{f}\left(\hat{\sigma}^{\mathsf{IV}}_{t+h-1}(K^*, T^*), \ldots, \hat{\sigma}^{\mathsf{IV}}_{t+h-l}(K^*, T^*), X_t\right)$$

Note that since X is known ahead of time since the FOMC conferences are pre-scheduled, the values of the dummy variable do not have to be forecasted.

Benchmark

The model is benchmarked to a commonly used random walk (RW) prediction, i.e. the prediction of IV at any (K^*, T^*) is a value at this point from the previous day:

$$\hat{\sigma}_t^{\mathsf{IV}}(K^*, T^*) = \sigma_{t-1}^{\mathsf{IV}}(K^*, T^*)$$

For majority of tasks this benchmark is quite easy to beat, however in case of financial time series, which are inherently noisy, susceptible to shocks and if markets are efficient even harder to predict, this benchmark is hypothesised to be a strong competitor.

Evaluation of results

We are using two metrics - root mean squared error (RMSE) and mean percentage error (MPE) to account for directional error.

Root mean squared error for a 1-step prediction is defined as follows:

$$RMSE = \sqrt{\frac{1}{NM}\sum_{i=1}^{N}\sum_{j=1}^{M}\left(\sigma_t^{\mathsf{IV}}(K_i^*, T_j^*) - \hat{\sigma}_t^{\mathsf{IV}}(K_i^*, T_j^*)\right)^2}$$

Mean percentage error for a 1-step prediction is defined as follows:

$$MPE = \frac{1}{NM}\sum_{i=1}^{N}\sum_{j=1}^{M}\frac{\sigma_t^{\mathsf{IV}}(K_i^*, T_j^*) - \hat{\sigma}_t^{\mathsf{IV}}(K_i^*, T_j^*)}{\sigma_t^{\mathsf{IV}}(K_i^*, T_j^*)}$$

Long Short-Term Memory (LSTM)

A type of recurrent neural network introduced in Hochreiter and Schmidhuber [5] designed to capture long-range temporal dependencies. It addresses the vanishing gradient problem in RNNs.

At each time t each cell of the network is updated according to these equations in terms of input x_t. The past data up to lookback parameter l is used:

- Memory cell $c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t$ where $\tilde{c}_t = \tanh(W_c x_t + U_c h_{t-1} + b_c)$
- Hidden state $h_t = o_t \odot \tanh(c_t)$
- Three gates:
  - Input gate $i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i)$
  - Forget gate $f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f)$
  - Output gate $o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o)$

Convolutional 2D LSTM (ConvLSTM)

Standard LSTM ignores spatial dependencies while implied volatility surface is inherently 2-dimensional and can be represented by a small number of latent factors. This issue can be solved through convolution, which was first introduced to LSTM in Shi et al. [9]. All matrix multiplications are replaced by a convolution operation $*$:

- $c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t$ where $\tilde{c}_t = \tanh(W_c * x_t + U_c * h_{t-1} + b_c)$
- Hidden state $h_t = o_t \odot \tanh(c_t)$
- $i_t = \sigma(W_i * x_t + U_i * h_{t-1} + b_i)$
- Forget gate $f_t = \sigma(W_f * x_t + U_f * h_{t-1} + b_f)$
- Output gate $o_t = \sigma(W_o * x_t + U_o * h_{t-1} + b_o)$

ConvLSTM for IV Surface Forecasting

In our case the inputs to 2D ConvLSTM are 3D tensors of the shape $X_t \in \mathbb{R}^{l \times N \times M}$, i.e. lookback and number of fixed grid strikes and maturities.

Convolutions act locally across (K, T) grid allowing the hidden state to capture **spatial dependencies**. Naturally LSTM architecture handles **temporal dependencies**.

IV surface is essentially treated as a moving picture which should account for a better generalisation and smoother, more consistent forecasts.

Model specification

The chosen hyperparameters are:

- Lookback l - 5
- Learning rate - 0.0005
- Kernel dimension k - 4
- Filter size F (dimension of the hidden state space) - 32
- Batch size - 32

Selected optimiser is **Adam** and the loss function is mean squared error (**MSE**).

Architecture - encoding-forecasting structure

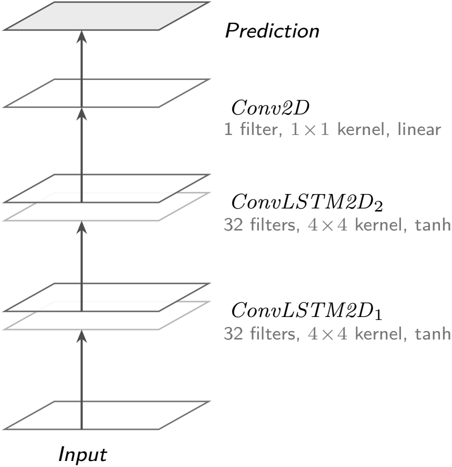

Architecture - including an annoucement dummy

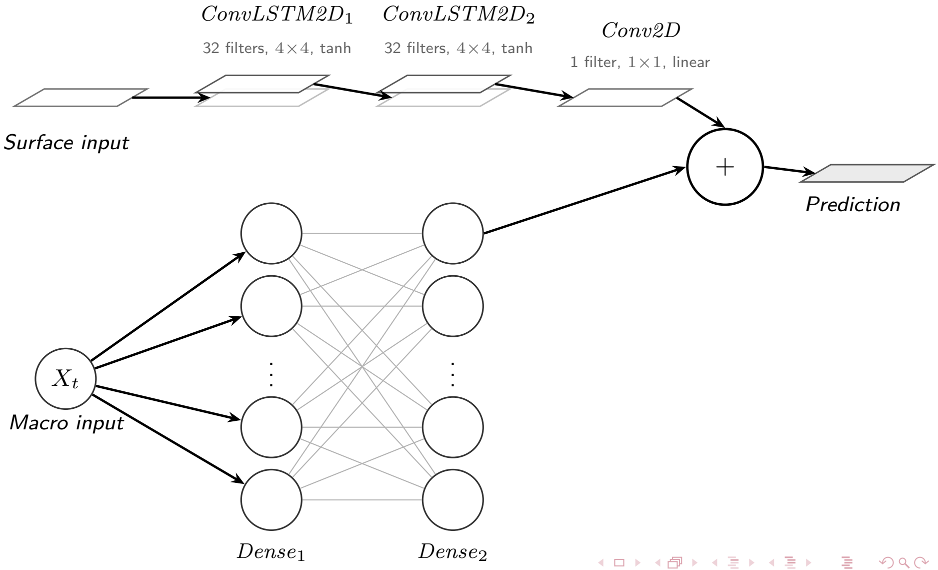

Results

Statistical and graphical analysis of the pre-announcement effect

In the first part we focus on understanding the data and the effect itself. We achieve this by visualising and testing at-the-money (ATM) IV defined as the IV with log-moneyness k closest to 0 (i.e. moneyness closest to 1) per fixed maturity T_j:

$$\sigma_t^{\text{ATM IV}}(T_j) = \sigma_t^{\text{IV}}(K_{t,j}^{\text{ATM}}, T_j)$$

where:

$$K_{t,j}^{\text{ATM}} = \text{arg} \min_{K \in \mathbb{K}_t^j} |k|$$

Statistical and graphical analysis of the pre-announcement effect

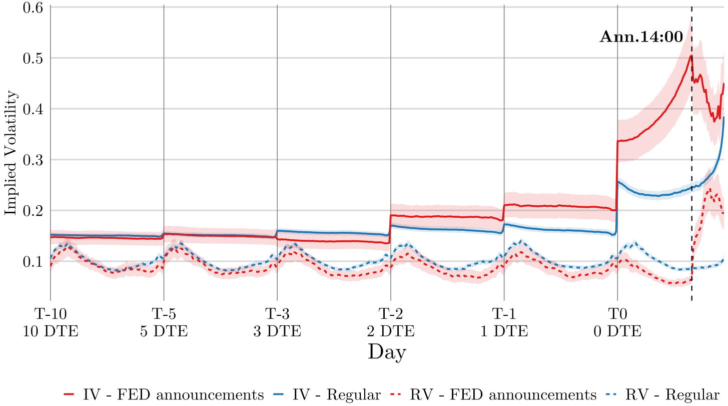

Figure: Average ATM Implied and Realised Volatility for call options expiring on the day of announcement denoted by T_0.

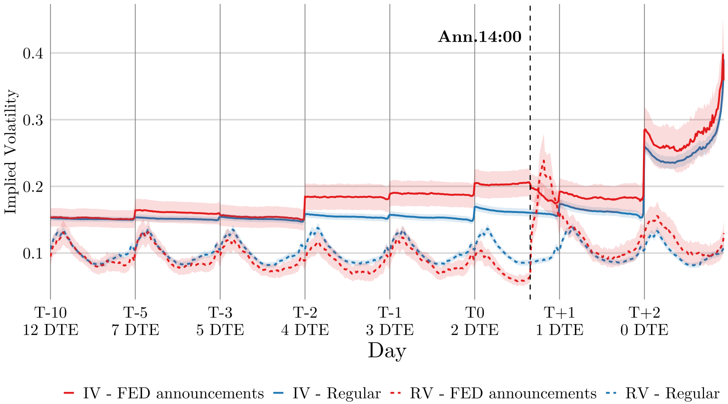

Figure: Average ATM Implied and Realised Volatility for call options expiring 2 trading days after the announcement denoted by T_0.

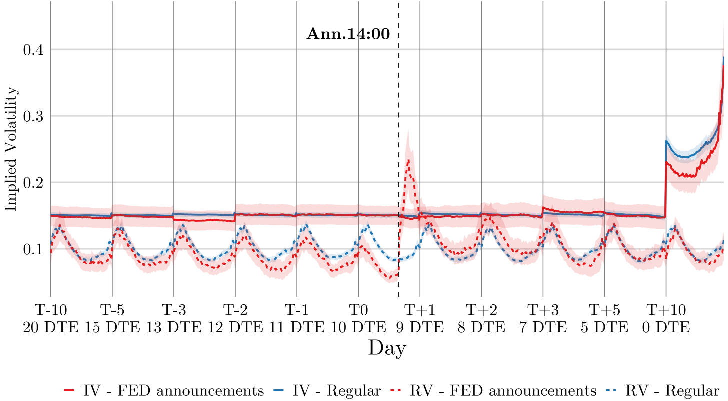

Figure: Average ATM Implied and Realised Volatility for call options expiring 10 trading days after the announcement denoted by T_0.

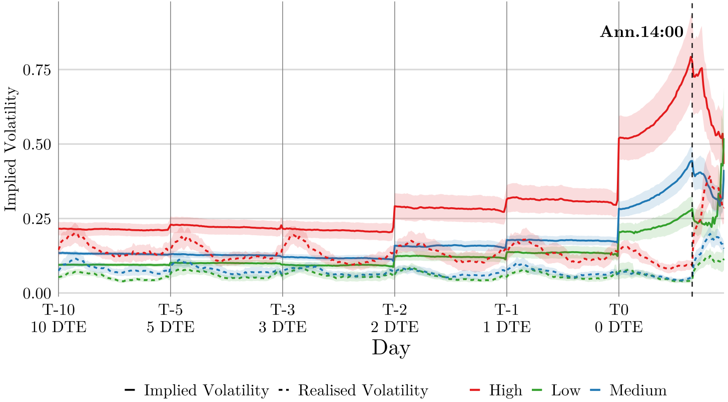

Figure: Average ATM Implied and Realised Volatility in the period 2016-2025 in different volatility regimes for call options expiring on the day of announcement denoted by T_0.

Statistical tests for the change in IV

We test more formally for the change in IV during the days of the Fed announcements of monetary policy using Mann-Whitney U-test (as the assumptions of t-test are violated):

$$\begin{cases} H_0 : \text{The distibution of IV at the date of the announcement is the same} \\ \text{as the distibution of IV during all other periods.} \\ H_1 : \text{The distibution of IV at the date of the announcement differs} \\ \text{from the distibution of IV during all other periods.} \end{cases}$$

Table: The Mann-Whitney U test around FED interest rates announcements across log-moneyness k for call options expiring on the day of the announcement.
*, **, and *** denote significance at 5%, 1%, and 0.1% levels, respectively. Colored cells are significant at at least 5% significance level.

DTE	Time	k = -0.02	k = -0.015	k = -0.01	k = -0.005	k = 0	k = 0.005	k = 0.01	k = 0.015	k = 0.02
3	10:00	0.795	0.883	0.909	0.992	0.138	0.913	0.835	0.820	0.743
	12:00	0.921	0.934	0.953	0.960	0.151	0.893	0.889	0.890	0.875
	14:00	0.802	0.743	0.805	0.797	0.191	0.726	0.652	0.696	0.635
	16:00	0.978	0.930	0.992	0.995	0.141	0.961	0.989	0.924	0.781
2	10:00	0.716	0.577	0.552	0.596	**0.007****	0.463	0.473	0.468	0.434
	12:00	0.648	0.456	0.366	0.339	**0.002****	0.239	0.208	0.187	0.215
	14:00	0.830	0.530	0.390	0.347	**0.001****	0.238	0.241	0.264	0.373
	16:00	0.368	0.216	0.185	0.190	**0.001*****	0.179	0.135	0.209	0.224
1	10:00	0.160	0.059	**0.037***	**0.024***	**0.000*****	**0.017***	**0.025***	**0.018***	0.057
	12:00	0.099	**0.026***	**0.009****	**0.005****	**0.000*****	**0.004****	**0.005****	**0.013***	**0.035***
	14:00	0.062	**0.016***	**0.006****	**0.004****	**0.000*****	**0.005****	**0.006****	**0.023***	0.117
	16:00	**0.022***	**0.004****	**0.001*****	**0.000*****	**0.000*****	**0.001*****	**0.004****	**0.036***	0.155
0	10:00	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****
	12:00	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****
	14:00	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****	**0.000*****

Forecasting results

In the second part we review the performance of machine learning models in IV surface forecasting. We also check their ability to fit the shock related to FOMC conferences.

Models without an annoucnement dummy X - RMSE

Table: Comparison of out-of-sample RMSE for ConvLSTM without an announcement dummy and benchmarks. The smallest RMSE is bolded.

	Type	Call				Put			
	Horizon	1	2	5	10	1	2	5	10
Model type	Model								
ConvLSTM	**AHBS**	0.356	0.357	0.360	0.361	0.178	0.179	0.181	0.184
	SVI	**0.161**	**0.165**	**0.165**	**0.172**	**0.083**	**0.089**	**0.096**	**0.114**
	linear	0.204	0.213	0.211	0.218	0.105	0.112	0.134	0.170
Random Walk	**AHBS**	0.373	0.375	0.378	0.381	0.194	0.194	0.197	0.201
	SVI	0.254	0.385	0.261	0.393	0.158	0.166	0.159	0.166
	linear	0.275	0.281	0.282	0.297	0.168	0.168	0.165	0.174

Models with an annoucnement dummy X - RMSE

Table: Comparison of out-of-sample RMSE for ConvLSTM with an announcement dummy and benchmarks. The smallest RMSE is bolded.

	Type	Call				Put			
	Horizon	1	2	5	10	1	2	5	10
Model type	Model								
ConvLSTM	**AHBS**	0.351	0.351	0.351	0.349	0.181	0.184	0.193	0.207
	SVI	**0.144**	**0.147**	**0.154**	**0.164**	**0.067**	**0.071**	**0.080**	**0.091**
	linear	0.204	0.215	0.218	0.231	0.092	0.094	0.103	0.114
Random Walk	**AHBS**	0.373	0.375	0.378	0.381	0.194	0.194	0.197	0.201
	SVI	0.254	0.385	0.261	0.393	0.158	0.166	0.159	0.166
	linear	0.275	0.281	0.282	0.297	0.168	0.168	0.165	0.174

Models without an annoucnement dummy X - MPE

Table: Comparison of out-of-sample MPE for ConvLSTM without an announcement dummy and benchmarks. The smallest absolute MPE is bolded. Underestimates are marked green, overestimates red.

	Type	Call				Put			
	Horizon	1	2	5	10	1	2	5	10
Model type	Model								
ConvLSTM	**AHBS**	**-0.005**	**-0.010**	**-0.010**	**0.017**	-0.095	-0.104	-0.124	-0.143
	SVI	-0.014	-0.021	-0.046	-0.075	**0.023**	**0.047**	**0.092**	0.148
	linear	-0.009	-0.032	-0.083	-0.147	0.104	0.146	0.237	0.362
Random Walk	**AHBS**	0.034	0.033	0.029	0.027	0.114	0.113	0.110	**0.109**
	SVI	0.109	0.107	0.100	0.093	0.217	0.214	0.212	0.208
	linear	0.113	0.112	0.103	0.097	0.223	0.220	0.218	0.214

Models with an annoucnement dummy X - MPE

Table: Comparison of out-of-sample MPE for ConvLSTM with an announcement dummy and benchmarks. The smallest absolute MPE is bolded. Underestimates are marked green, overestimates red.

	Type	Call				Put			
	Horizon	1	2	5	10	1	2	5	10
Model type	Model								
ConvLSTM	**AHBS**	-0.026	-0.038	-0.064	-0.106	0.045	0.079	0.167	0.284
	SVI	**0.010**	**0.013**	**0.021**	0.043	**0.004**	**0.010**	**0.026**	**0.055**
	linear	0.041	0.054	0.077	0.117	-0.007	-0.018	-0.042	-0.083
Random Walk	**AHBS**	0.034	0.033	0.029	**0.027**	0.114	0.113	0.110	0.109
	SVI	0.109	0.107	0.100	0.093	0.217	0.214	0.212	0.208
	linear	0.113	0.112	0.103	0.097	0.223	0.220	0.218	0.214

Statistical significance of predictions - ConvLSTM vs. benchmark

Table: Comparison of out-of-sample Diebold-Mariano p-value for ConvLSTM without announcement dummy X and benchmarks. Statistically significant results are colored red.

	Type	Call				Put			
	Horizon	1	2	5	10	1	2	5	10
Model type	Model								
ConvLSTM	**AHBS**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
	SVI	**0.000**	0.135	**0.000**	0.125	**0.000**	**0.000**	**0.000**	**0.000**
	linear	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**

Statistical significance of predictions - models with and without announcement dummy X

Table: Comparison of out-of-sample Diebold-Mariano p-value for ConvLSTM with and without announcement dummy X.

	Type	Call				Put			
	Horizon	1	2	5	10	1	2	5	10
Model type	Model								
ConvLSTM	**AHBS**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
	SVI	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
	linear	0.881	0.086	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**

Statistical significance of predictions - models with and without announcement dummy X only during announcement days

Table: Comparison of out-of-sample Diebold-Mariano p-value for ConvLSTM with and without announcement dummy X only during annoucement days.

	Type	Call				Put			
	Horizon	1	2	5	10	1	2	5	10
Model type	Model								
ConvLSTM	**AHBS**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
	SVI	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
	linear	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**

Analysis and diagnostics

In what follows we will analyse and diagnose the model which learns on the surface interpolated/extrapolated by SVI because it performs the best.

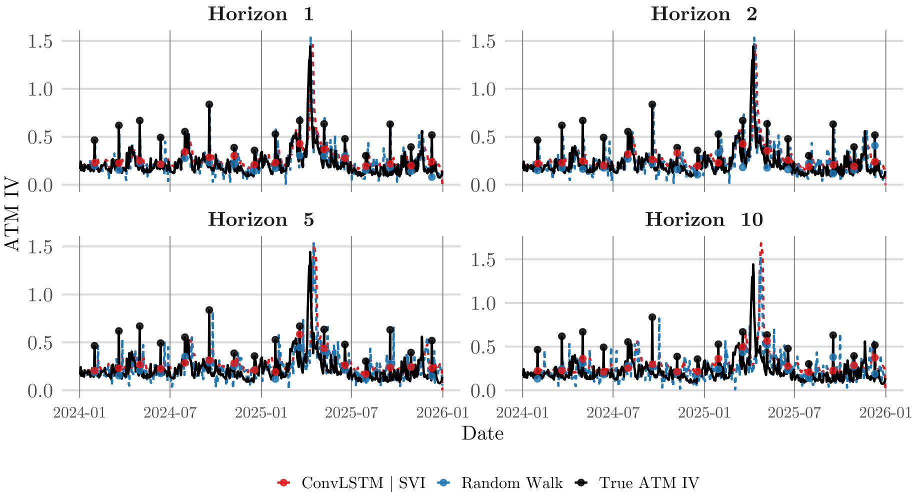

Figure: Comparison between naive Random Walk and ConvLSTM model out-of-sample predictions of ATM IV predictions for SVI model, call options. Dots represent an annoucement date.

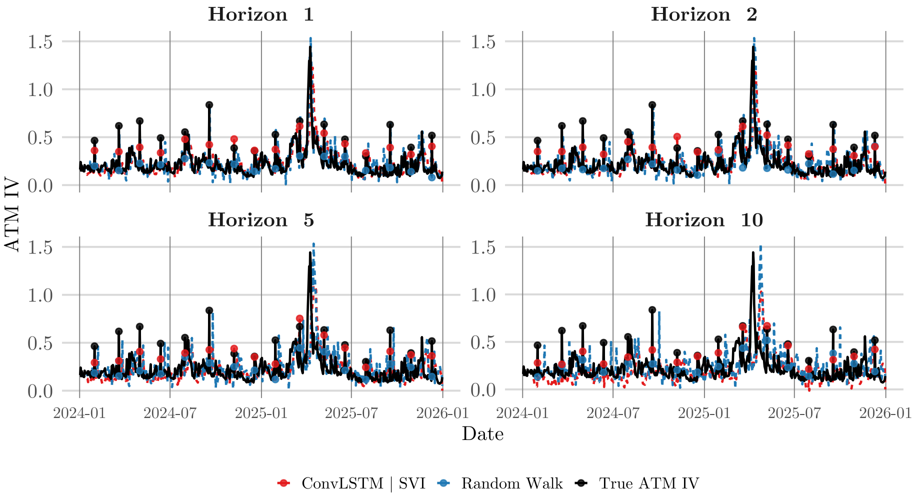

Figure: Comparison between naive Random Walk and ConvLSTM model with announcement dummy out-of-sample predictions of ATM IV predictions for SVI model, call options. Dots represent an annoucement date.

Announcement shock - call options

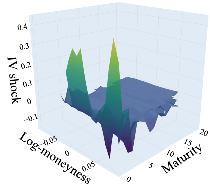

(a) Average difference between the implied volatility of call options at the time of the FOMC conference and the day before.

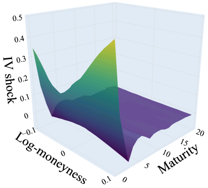

(b) Average difference between the fitted (with the ConvLSTM model with an announcement dummy and interpolated by SVI) implied volatility $\hat{\sigma}^{\mathrm{IV}}$ of call options at the time of the FOMC conference and the day before.

Announcement shock - put options

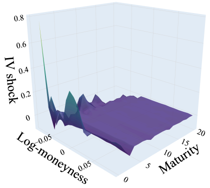

(a) Average difference between the implied volatility of put options at the time of the FOMC conference and the day before.

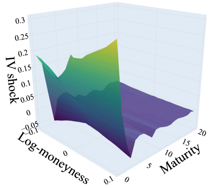

(b) Average difference between the fitted (with the ConvLSTM model with an announcement dummy and interpolated by SVI) implied volatility $\hat{\sigma}^{\mathrm{IV}}$ of put options at the time of the FOMC conference and the day before.

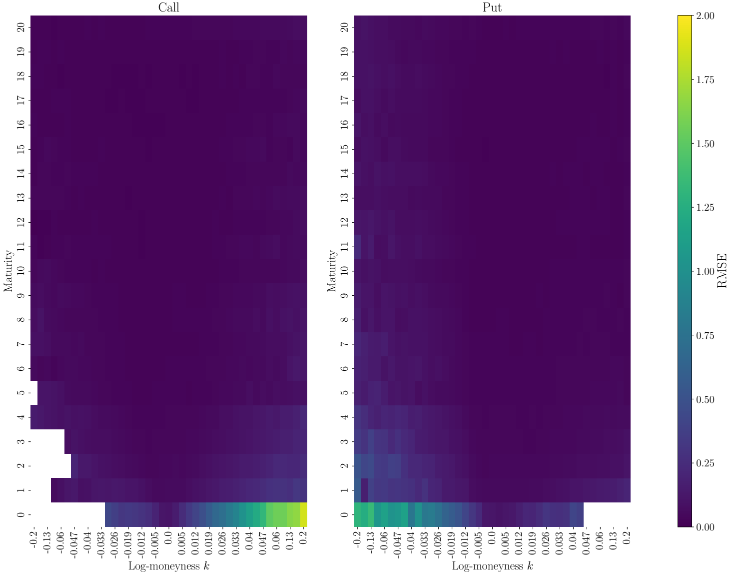

Figure: Out-of-sample RMSE for predictions from ConvLSTM with announcement dummy - surface interpolated by SVI model, prediction horizon 1.

Impulse response function

Training the models allow to simulate the model implied reaction of the implied volatility to a shock related to the annoucnement of interest rates. Therefore, assuming the shock happens at time T we follow the procedure:

- Compute the average pre-announcement surfaces $(\bar{\sigma}(K^*, T^*) = \bar{\hat{\sigma}}^{\mathsf{IV}}_{T-1}(K^*, T^*), \ldots, \bar{\hat{\sigma}}^{\mathsf{IV}}_{T-l}(K^*, T^*))$
- Predict the surface $\mathbb{E}\left[\hat{\sigma}^{\mathsf{IV}}_{T}(K^*, T^*)\right]$ setting the dummy to 1 or 0
- Predict subsequent surfaces using two versions and calculate the difference

We thus predict the announcement and non-annoucement impulse response for horizon h for all pairs (K_i^*, T_j^*):

$$IRF_h(K_i^*, T_j^*) = \mathbb{E}\left[\sigma^{\mathsf{IV}}_{T+h}(K_i^*, T_j^*)|X_T = 1, \bar{\sigma}(K^*, T^*)\right] - \mathbb{E}\left[\hat{\sigma}^{\mathsf{IV}}_{T+h}(K_i^*, T_j^*)|X_T = 0, \bar{\sigma}(K^*, T^*)\right]$$

Figure: The difference between the announcement and non-announcement IRF for ATM IV 0 DTE of call options for the ConvLSTM model trained on the surface interpolated by SVI.

Conclusions

Hypotheses review

Hypothesis	Conclusion
The implied volatility monotonically increases prior to the announcement.	✓
The effect is stronger for short dated options.	✓
The effect is stronger in high volatility regimes.	✓
The effect is stronger for OTM options.	×
ML framework beats the benchmark random walk in prediction power both globally and during policy announcement days.	✓
A better model fit can be achieved for ATM implied volatility and longer maturities.	✓
Adding the date of FOMC meetings as a feature improves overall model performance and during policy announcement days.	✓

Limitations

- The model is rather small in ML framework terms and only one type of architecture was considered
- There was only one exogeneous variable included
- The data quality in the wings is of a poorer quality
- The results are not tested in terms of trading, only statistical significance was considered

Future research

- Inclusion of more exogeneous variables, like the return and volatility of the underlying, VIX, as well as economic variables
- Consideration of other architectures - attention mechanism, more layers.
- Economic significance of results - application to trading around policy announcement days

General takeaways

- Pre-announcement effect of elevated uncertainty translates prominently into options market, where investors price the policy and systemic risk
- ML models can effectively forecast the IV surface also during abnormal days
  - The convolutional LSTM paired with SVI interpolation/extrapolation generally achieves the best performance
- The edge of the ML framework can be limited due to the noisy characteristics of the IV surface
- Inclusion of scheduled policy announcements improves the quality of predictions

Thank you for your attention!

Questions?

Łukasz Adamski

l.adamski@uw.edu.pl

References I

[1] Board of Governors of the Federal Reserve System. *Federal Open Market Committee - Historical Materials by Year*. URL: https://www.federalreserve.gov/monetarypolicy/fomc_historical_year.htm (visited on 12/30/2025).

[2] Chicago Board Options Exchange. *FRED Data*. URL: https://datashop.cboe.com/option-quote-intervals (visited on 03/11/2026).

[3] Bernard Dumas, Jeff Fleming, and Robert E. Whaley. "Implied Volatility Functions: Empirical Tests". In: *The Journal of Finance* 53.6 (1998), pp. 2059–2106. (Visited on 04/13/2026).

References II

[4] Jim Gatheral and Antoine Jacquier. "Arbitrage-free SVI volatility surfaces". In: *Quantitative Finance* 14.1 (2014), pp. 59–71.

[5] Sepp Hochreiter and Jürgen Schmidhuber. "Long Short-Term Memory". In: *Neural Computation* 9.8 (1997), pp. 1735–1780.

[6] Roger W. Lee. "The Moment Formula for Implied Volatility at Extreme Strikes". In: *Mathematical Finance* 14.3 (2004), pp. 469–480.

[7] David O. Lucca and Emanuel Moench. "The Pre-FOMC Announcement Drift". In: *The Journal of Finance* 70.1 (2015), pp. 329–371.

References III

[8] Pavel G. Savor and Mungo I. Wilson. "How Much Do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements". In: *Journal of Financial and Quantitative Analysis* 48.2 (2013), pp. 343–372.

[9] Xingjian Shi et al. "Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2015.