



# Unlocking the black box: Non-parametric option pricing before and during COVID-19

Nikola Gradojevic<sup>1,2</sup> · Dragan Kukolj<sup>2</sup>

Accepted: 18 January 2022 / Published online: 25 February 2022

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022

## Abstract

This paper addresses the interpretability problem of non-parametric option pricing models by using the explainable artificial intelligence (XAI) approach. We study call options written on the S&P 500 stock market index across three market regimes: pre-COVID-19, COVID-19 market crash, and post-COVID-19 recovery. Our comparative option pricing exercise demonstrates the superiority of the random forest and extreme gradient boosting models for each market regime. We also show that the model's pricing accuracy has worsened from the pre-COVID-19 to the recovery period. Moneyness was the most important price determinants across the market regimes, while the implied volatility and time-to-maturity inputs contributed intermittently to a lesser extent. During the COVID-19 crash, open interest gained more economic importance due to the increased behavioral tendencies of traders consistent with market distress.

**Keywords** Option pricing · COVID-19 · Random forest · Extreme gradient boosting · Explainable artificial intelligence · Interpretability

**JEL Classification** G01 · C45 · C55 · C58

## 1 Introduction

The systemic risk related to the financial contagion effects of the COVID-19 pandemic, on February 20, 2020, turned a steadily increasing S&P 500 index market into panic and distress. Between February 20 and March 23, the market plummeted at a record pace and lost more

---

✉ Nikola Gradojevic  
ngradoje@uoguelph.ca

Dragan Kukolj  
dragan.kukolj@rt-rk.com

<sup>1</sup> Department of Economics and Finance, University of Guelph, Lang School of Business and Economics, 50 Stone Road East, Guelph, ON N1G 2W1, Canada

<sup>2</sup> Faculty of Technical Sciences, University of Novi Sad, Novi Sad, Serbia

than 30%. During this time, the volatility index (VIX) peaked at 82% on March 16, 2020.<sup>1</sup> This was similar to the subprime mortgage crisis when the VIX closed at 79% on October 24, 2008. However, the bear market that ensued did not last for too long and the recovery started on March 24, 2020. By June 8, 2020, the S&P 500 index market recuperated about 85% of the losses from its record closing high on February 19, 2020.

The option market followed the underlying S&P 500 index market closely. Specifically, the average premium of call options written on the S&P 500 index (across all strike prices and maturities) declined from \$139.34 on February 19, 2020 to \$101.31 on March 23, 2020. More generally, considering that the S&P 500 index option prices represent discounted expectations of future outcomes of the S&P 500 index, if the market expects a crash, the prices of out-of-the-money put options would be relatively large compared to the out-of-the-money call options. In fact, on February 19, 2020, the average price of all the out-of-the-money put (call) options was \$25.40 (\$39.64). Hence, it appears that the COVID-19 crash that followed had not been anticipated by the market participants. Then, over the period until March 23, 2020, the average daily price of all the out-of-the-money put (call) options increased (decreased) to \$105.83 (\$34.99). This indicates the emergence of extreme pessimism of option traders at the trough of the crisis. In the same vein, by analyzing the prices of options on commodity futures in March, 2020, Vercammen (2020) also reported traders' growing confusion and concern about the future of commodity prices and disruptions to supply chains due to COVID-19.

Broadly speaking, the COVID-19 contagion has had a devastating global impact on various businesses and markets over both short- and long-run (Sharif et al. 2020). Certain industries such as tourism, transportation, hospitality and automotive manufacturing were halted abruptly. Mazur et al. (2021) investigated the performance of the stock constituents of the S&P 1500 index during the COVID-19 crisis. They found that health care, food, natural gas, and software sectors performed abnormally well, while crude petroleum, real estate, entertainment and hospitality sectors underperformed and produced high negative returns. Similarly, for the Asian stock markets, Liu et al. (2020) documented large cumulative abnormal returns for the pharmaceutical, software and IT services, while transportation, lodging and catering incurred heavy losses. Such a discouraging outlook for profitability was further confirmed in Harjoto et al. (2021) in the sense that COVID-19 generated substantial negative shocks to emerging markets and for small firms. As explained in Bellalah et al. (2020), the COVID-19 crisis was marked by a significant worldwide reduction in allocated investments and their model was able to assist in portfolio decisions in presence of regime-switching.

In the context of financial contagion, Ramelli and Wagner (2020) illustrated how the effects of the COVID-19 pandemic were amplified through financial channels. Morelli and Petrella (2021) suggested that the contagion was spread between option and stock markets during the COVID-19 outbreak and that it affected both European and American options in a similar manner. Furthermore, Akhtaruzzaman et al. (2021c) demonstrated that the strength of the relationship between Chinese and G7 financial and non-financial companies' stock returns increased substantially during the COVID-19 crash. In regards to oil prices, Akhtaruzzaman et al. (2021a) found that the COVID-19 pandemic appeared to have moderated the oil price risk exposure of both financial and non-financial firms. As for the U.S. stock market, by studying the frequency domain causality, Lento and Gradojevic (2021) documented that the S&P 500 index returns caused oil returns prior to the pandemic, but that the causality direction was reversed during the market crash and recovery. They also reported mostly bi-directional causalities between the S&P 500 index returns and various asset classes when

<sup>1</sup> The VIX is an index providing 30-day market expectations based on the S&P 500 Index. Higher values of VIX indicate the risk that the market will make a large swing.

the market crashed during the February 20, 2020 to March 23, 2020 period. With respect to gold prices, Akhtaruzzaman et al. (2021b) showed that gold lost its ‘safe haven’ property during the market recovery (from March 17, 2020 to April 24, 2020). Their evidence basically corroborates that from Lento and Gradojevic (2021) that observed price spillover effects from gold to the stock market at this time.

This paper studies the behavior of the S&P 500 index option market in response to the COVID-19 pandemic through the lens of explainable artificial intelligence (XAI). First, our objective is to track the option pricing accuracy of an array of competing models across three time periods: from January 1, 2020 to February 19, 2020 (pre-COVID-19), from February 20, 2020 to March 23, 2020 (COVID-19 market crash), and from March 24, 2020 to June 15, 2020 (post-COVID-19 recovery). Since the seminal works of Bates (1991, 2000), there has been a few papers that analyzed the performance of option pricing models during a financial crisis (Bates 2012; Fulop et al. 2014; Calvet et al. 2015; Driouchi et al. 2018; Luo et al. 2018; Kukolj et al. 2012). The consensus in the literature is that a crisis represents a distinct market state characterized by psychological biases, regime-switching, tail behavior and price jumps, and that it should be approached with more adaptable option pricing models that relax the assumptions of the classic Black-Scholes framework (Black and Scholes 1973). For such a reason, we employ four sophisticated non-parametric models (support vector regression, feedforward artificial neural network, extreme gradient boosting and random forest) and we test their out-of-sample pricing performance. The results suggest that the extreme gradient boosting and random forest models deliver the most accurate pricing performance on all time periods. In particular, the model’s pricing error increased over time and it appears that it was slightly more affected by the uncertainty of the recovery period than by the actual COVID-19 crash. We also find that, irrespective of the time period tested, it is generally more difficult to price long-term options with lower volume that are not “near-the-money”. Overall, we conclude that market volatility adversely impacted the pricing ability of all competing models.

Our second objective is to provide an economic interpretation of the pricing accuracy on the sub-periods of data by (i) clustering of options based on the ranges of input variables, and (ii) estimating the relative importance of each input. The former approach is an extension of Gradojevic et al. (2011) where the options will be clustered by observing their moneyness,<sup>2</sup> time-to-maturity, implied volatility, open interest and volume. The latter relies on the XAI concept (Ding 2018) in which we interpret the contribution of inputs of a non-parametric option pricing model that resulted in its optimal pricing accuracy. To the authors’ best knowledge, this is the first research work on option pricing models that employs sophisticated non-parametric methods (including the XAI) and studies a wide variety of potential inputs during a market regime shift.<sup>3</sup>

The results from the clustering exercise reveal that the optimal number of clusters was reduced during the COVID-19 crisis. Essentially, the market became more homogeneous with regard to the variety of option types that were traded. This can also be understood as the reduction in the heterogeneity of market participants (and their beliefs) that were less dispersed due to the COVID-19 fears. Furthermore, in the pre-COVID-19 market, option prices were greatly influenced by their moneyness and volatility, and to a lesser extent by other inputs. However, when the volatility increased substantially and the market switched to

<sup>2</sup> Moneyness is the ratio between the price of the underlying security (S&P 500 index:  $S$ ) and its strike (or exercise) price ( $K$ ). Call options are “in (out-of)-the-money” when  $S > K$  ( $S < K$ ).

<sup>3</sup> The objective of this paper is not to use the proposed models in risk management, portfolio optimization or hedging. Rather, the current paper aims at understanding the option market microstructure effects by gauging the sentiment of traders exposed to regime shifts.

crisis mode, option prices were less sensitive to the volatility input and more sensitive to open interest. The importance of such a “non-standard” input could be interpreted as the strength of confidence of market participants that the current (bearish) trend would continue and, also, that traders panicked, whereas market psychology was at play at this time. This effect was diminished during the post-COVID-19 recovery period while moneyness and volatility regained economic importance. Surprisingly, time-to-maturity was not identified as a major driving force in option pricing, but it is worthwhile to note that it was among the three most important inputs during the COVID-19 crash.

Our paper fits within the broader literature on non-parametric option pricing and hedging models that rely on machine learning. One of the pioneering papers in this area was Hutchinson et al. (1994), followed by Qi and Maddala (1996), Garcia and Gençay (2000), Gençay and Qi (2001), Gençay and Altay-Salih (2003), and Gradojevic et al. (2009). More recent work can be seen in, for example, Gradojevic (2016) and Jang and Lee (2019). The goal of such scholarly efforts is to harness the learning ability and flexibility of machine learning models to achieve better prediction accuracy than the classical (parametric) financial option models.

In relation to the existing option pricing literature, the novelty of this paper lies in the following: (1) making use of the advanced AI methods such as extreme gradient boosting and random forest in option pricing; (2) introducing two additional inputs (open interest and volume) to a traditional non-parametric model and measuring their significance; (3) proposing a new clustering approach based on a re-organizing neural network to facilitate option pricing and model interpretation; (4) utilizing XAI in interpreting the non-parametric models; and (5) testing option pricing models across various market regimes surrounding the COVID-19 market meltdown.

In short, the main results of this paper can be summarized as follows: (1) non-parametric models (extreme gradient boosting and random forest) produce the most accurate forecast performance across all market regimes; (2) the Black-Scholes model’s pricing accuracy is comparable to that of the non-parametric models during the market crash; (3) the relative importance of traditional option pricing model’s inputs (moneyness, time-to-maturity and volatility) is robust during regime shifts; (4) open interest becomes a more important input at times of market volatility and distress; and (5) the proposed data clustering method is beneficial for providing deeper insights into the behavior of market participants and data generating mechanisms of option prices.

The paper is laid out as follows: The next section reviews the option pricing literature with the special emphasis on non-parametric models. Section 3 describes the S&P 500 index options data. Section 4 presents the methodological steps involved in our research design as follows: (i) empirical models (support vector regression, artificial neural network, and random forest), (ii) optimal clustering (with re-organizing neural networks), (iii) sub-period analysis, and (iv) feature importance estimation. Section 5 provides the empirical findings, while Sect. 6 concludes.

## 2 Background literature: option pricing

To address the well-known pricing biases of the Black-Scholes model (Black and Scholes 1973) such as the “volatility smile” and “volatility smirk” phenomena, research efforts have expanded into developing both parametric and non-parametric alternatives.<sup>4</sup> The research on

<sup>4</sup> For an exhaustive review of the option pricing literature please see Renault (2010).

parametric models has mainly focused on introducing the stochastic volatility (SV), stochastic volatility random jump (SVJ) and stochastic interest rate (SI) to the Black-Scholes setting. These parametric extensions have been shown to be superior to the Black-Scholes model in out-of-sample pricing and hedging exercises (Bakshi et al. 1997). Specifically, the SV model has been shown to have first-order importance over the Black-Scholes model (Gençay and Gibson 2009). The SVJ model further enhances the SV model for pricing short-term options, while the SI model extends the SVJ model in regards to the pricing of long-term options. Additional notable parametric research contributions include Bates (1991, 2000), Heston (1993), Christoffersen et al. (2009), Calvet et al. (2015), He and Zhu (2016), Wong and Lo (2009), Cai and Kou (2011) and Gaß et al. (2018).

Although parametric models in general improve the pricing accuracy of the Black-Scholes model, these models exhibit certain pricing biases and are often inferior to non-parametric approaches. Non-parametric models are not bound by the normality of return distributions (or other unjustifiable parametric assumptions) and can benefit from their adaptive learning abilities (Jang and Lee 2019; Gradojevic et al. 2009). For example, in Gençay and Gibson (2009), the out-of-sample performance of an artificial neural network model was compared to the SVJ, SI and SV parametric approaches for the S&P 500 stock market index. The paper showed that all three parametric models were dominated by the neural network pricing model with the GARCH (1, 1) volatility. Similarly, a semi-parametric approach from Andreou et al. (2008) was able to improve upon the SV and SVJ models.<sup>5</sup>

The non-parametric approaches to option pricing have also been used by Hutchinson et al. (1994), Garcia and Gençay (2000), Qi and Maddala (1996), Gençay and Qi (2001), and Gençay and Altay-Salih (2003). More recent examples can be found in Barone-Adesi et al. (2008), Fan and Mancini (2009), von Spreckelsen et al. (2014) and Guidolin and Hansen (2016). An innovative strand of literature concerns option pricing with wavelets (Liu et al. 2019). This paper used a set of daily options written on the DAX-30 index over the 2009–2012 period. The wavelet-based option pricing model outperformed the SVJ parametric benchmark on out-of-sample data, but it was still inferior to the neural network non-parametric model.

In addition, the predictive accuracy and hedging prowess of neural network-based models has been highlighted in a very recent paper by Cao et al. (2021). They develop a novel hybrid gated neural network option pricing model, where they used a similar method to predict the S&P 500 index implied volatility. This paper basically reinforces the work of Culkin and Das (2017) that showed how neural networks could be trained to mimic option pricing traders. Buehler et al. (2019) is another related contribution that demonstrates the advantages of hedging with deep learning neural networks on the S&P500 index option data (2013–2018) in comparison to the Black-Scholes model.<sup>6</sup>

Some other important option pricing approaches are the mixture of distributions model by Melick and Thomas (1997) and Bhat and Kumar (2012), and the semi-parametric estimator by Ait-Sahalia and Lo (1998). These models have shown sizable improvements in option pricing accuracy compared to the Black-Scholes model; however, their out-of-sample pricing is inferior to non-parametric modular neural network models (Gradojevic et al. 2009). It is also worthwhile to mention non-parametric approaches that are based on the affine jump-diffusion models (Carr and Wu 2004) and the normal inverse Gaussian models (Eriksson et al. 2009; Barndorff-Nielsen and Shephard 2001). Furthermore, fuzzy logic has been proven very

<sup>5</sup> As the literature shows that models based on neural networks are able to price options more accurately than parametric alternatives such as the SV and SVJ approaches, we focus our attention on non-parametric methods (Gençay and Gibson 2009; Andreou et al. 2008; Gradojevic 2016; Jang and Lee 2019).

<sup>6</sup> Ruf and Wang (2020) offer an excellent up-to-date review of the relevant literature on the applications of neural networks for option pricing and hedging.

useful for option pricing in an uncertain market environment (Agliardi and Agliardi 2009). Nevertheless, fuzzy logic models have never been compared to non-parametric competitors in an out-of-sample pricing exercise with real market data.

### 3 Data

The data are provided by DeltaNeutral and represent the daily closing S&P 500 index European call option prices (the average of the bid-ask quotes), taken from the Chicago Board Options Exchange. Call options across different strike prices and maturities are considered for the period from January 1, 2020 to June 15, 2020. Since it is one of the deepest and the most liquid option markets in the United States, the S&P 500 index option market is sufficiently close to the theoretical setting of the Black-Scholes model. The implied volatility used in the estimations is a proprietary mean estimate provided by DeltaNeutral. The risk-free rate is approximated by the monthly yield of the U.S. Treasury bills. To reduce the size of the data set, options with zero volume and open interest on a given day were eliminated.

The data were divided into three non-overlapping sub-samples as follows:

- January 1, 2020 to February 19, 2020 (pre-COVID-19): 12989 observations;
- February 20, 2020 to March 23, 2020 (COVID-19 market crash): 19547 observations;
- March 24, 2020 to June 15, 2020 (post-COVID-19 recovery): 38909 observations.

The statistical properties of the data set are presented in Table 1. The data across three sub-periods are divided into several categories in terms of moneyness and time-to-maturity ( $\tau$ ). A call option is defined to be out-of-the money (OTM) if  $(S_t/K) < 0.95$ , near-the-money (NTM) if  $0.95 \leq (S_t/K) \leq 1.05$  and in-the-money (ITM) if  $(S_t/K) > 1.05$ . An option is classified as short-term if  $\tau < 60$  days, medium-term if  $60 \leq \tau \leq 180$  days and long-term if  $\tau > 180$  days. The reported numbers are the average quoted bid-ask midpoint price, the standard deviation of the bid-ask midpoint prices (shown in parentheses) and the total number of observations for each moneyness-maturity category. OTM, NTM, and ITM options take approximately 50%, 33%, and 17% of the total sample, respectively. The average prices of call options range from \$0.91 for the OTM, short-term options in the pre-COVID-19 period to \$807.46 for the ITM, long-term options in the post-COVID-19 period. Standard deviations become extremely large for the ITM options on all sub-samples and they generally increase with maturity.

It is useful to note that the COVID-19 and post-COVID-19 sub-periods markedly affected all options, regardless of their moneyness or maturity. First, the prices of both OTM and NTM medium- and short-term options on average increased during the COVID-19 crash and, then, slightly decreased during the recovery period. This is not the case, however, with long-term options whose prices were the highest in the post-COVID-19 period for all maturities. Clearly, long-term uncertainty was priced at its peak after the crash. Standard deviations followed the similar inverted-U pattern across the sub-periods for all OTM and NTM options, while the standard deviations for ITM options did not show any discernible temporal patterns.

**Table 1** SAMPLE PROPERTIES OF THE S&P 500 INDEX EUROPEAN CALL OPTION PRICES

		$(S_t/K) < 0.95$	$0.95 \leq (S_t/K) \leq 1.05$	$(S_t/K) > 1.05$
pre-COVID-19	$\tau < 60$	0.91	49.31	427.96
		(1.12)	(48.38)	(413.09)
		693	3698	1059
	$60 \leq \tau \leq 180$	6.76	70.50	581.38
		(6.70)	(47.35)	(515.43)
		990	2201	503
	$\tau > 180$	42.87	187.28	670.35
		(46.46)	(71.40)	(556.46)
		1718	1501	626
COVID-19 crash	$\tau < 60$	17.07	108.26	491.21
		(30.44)	(67.11)	(413.16)
		4897	2288	801
	$60 \leq \tau \leq 180$	37.83	172.79	586.69
		(46.08)	(74.51)	(365.96)
		3327	1438	549
	$\tau > 180$	57.31	241.24	765.62
		(64.28)	(79.27)	(468.53)
		4317	1161	769
post-COVID-19	$\tau < 60$	10.86	87.88	452.45
		(16.82)	(54.92)	(374.77)
		6670	5631	2961
	$60 \leq \tau \leq 180$	32.64	165.90	656.40
		(33.79)	(54.77)	(443.10)
		5442	3250	2111
	$\tau > 180$	61.82	268.04	807.46
		(62.16)	(63.24)	(476.31)
		7992	2548	2304

The reported numbers are the average quoted bid-ask midpoint prices of call options, the standard deviation of the bid-ask midpoint prices (shown in parentheses), and the total number of observations for each sub-period (pre-COVID-19, COVID-19 crash and post-COVID-19) and moneyness-maturity category.  $S$  denotes the daily spot S&P 500 index level,  $K$  is the exercise price and  $\tau$  is time-to-maturity. The options are categorized based on maturity (short-term:  $\tau < 60$ , medium-term:  $60 \leq \tau \leq 180$ , long-term:  $\tau > 180$ ), moneyness (out-of-the-money:  $(S_t/K) < 0.95$ , near-the-money:  $0.95 \leq (S_t/K) \leq 1.05$ , in-the-money:  $(S_t/K) > 1.05$ ). The sample period extends from January 1, 2020 through June 15, 2020, for a total of 71445 calls

## 4 Methodology

### 4.1 The model

The option pricing formula is defined in the spirit of Hutchinson et al. (1994), Garcia and Gençay (2000) and Gradojevic et al. (2009):

$$C_t = \phi(S_t, K, \tau, \sigma_{IV}, OPI, VOL) \quad (1)$$



where  $C_t$  is the call option price,  $S_t$  is the price of the underlying asset,  $K$  is the strike price,  $\tau$  is the time-to-maturity,  $OPI$  denotes open interest (number of open call contracts),  $VOL$  is volume (number of contracts traded) and  $\sigma_{IV}$  is the implied volatility. Assuming the homogeneity of degree one of the pricing function  $\phi$  with respect to  $S_t$  and  $K$ , one can write the option pricing function as follows:

$$\underbrace{\frac{C_t}{K}}_{c_t} = \phi \left( \underbrace{\frac{S_t}{K}}_{x_1}, 1, \underbrace{\tau}_{x_2}, \underbrace{\sigma_{IV}}_{x_3}, \underbrace{OPI}_{x_4}, \underbrace{VOL}_{x_5} \right) = \phi(x_1, x_2, x_3, x_4, x_5). \quad (2)$$

In general, options are often referred to as plain vanilla derivatives because their payoff (or price) is determined by the so-called underlying, which is in our case the S&P 500 stock market index. Call options are more profitable for the buyer when, *ceteris paribus*, the price of the underlying ( $S_t$ ) increases or the strike price ( $K$ ) decreases. Therefore, intuitively, these two variables must be integral parts of the option pricing formula. Further, when time-to-maturity ( $\tau$ ) increases, call options in general become more valuable. This is explained by the fact that it is more likely that the option will be in the money ( $S_t - K > 0$ ) and, thus, worthwhile exercising at maturity. The preceding explanatory variables are extended with the implied volatility that is a standard input to an option pricing model. The two non-standard inputs that are used in our model are open interest and volume. Apart from the possibility that such inputs may gain importance in a distressed market (e.g., COVID-19 crash), the rationale for using these two predictors is hinted at in Gârleanu et al. (2009) who found that demand pressures could influence option prices. Volume and open interest reflect the activity of option traders (i.e., market sentiment) that is related to demand for options.

## 4.2 Clustering of options

First, we determine the optimal number of clusters based on the partitioning of the input space for all sub-periods. For such a purpose we utilize the Davies-Bouldin ( $DB$ ) index (Davies and Bouldin 1979):

$$DB = \frac{1}{M} \sum_{i=1}^M R_i, \quad R_i = \max(R_{ij}), \quad j = 1, \dots, M; j \neq i \quad (3)$$

where  $R_{ij} = (S_i + S_j)/D_{ij}$  is a similarity measure between observations in data partitions (clusters)  $i$  and  $j$ ,  $S_i$  is a dispersion measure of the observations of  $i^{th}$  cluster calculated as the average Euclidean distance of the data points in cluster  $i$  to its center, and  $D_{ij}$  is a cluster's dissimilarity measure (distance between the centers of cluster pairs). The partition that produces the minimum  $DB$  is considered to be optimal.

The actual clustering is performed by using a competitive learning algorithm called Re-Organizing Neural Network (RONN) (Kukolj and Levi 2004; Kukolj et al. 2006). The algorithm relies on training data that contain  $N$  input observations  $x_k, k = 1, \dots, N$ , where  $N$  is the sample size. If we denote the dimension of the input vector by  $n$ , then the clusters we find will represent areas where data values are highly concentrated.

The RONN generates a one-layered network with one input layer where  $x_k$ 's are fed. Nodes in the output layer represent cluster centers. The total number of clusters is an input parameter of the algorithm. In tandem with the  $DB$  index measure, for each pre-selected number of clusters, the RONN algorithm finds the cluster centers, which is followed by the selection of optimal clustering structure, based on the minimal  $DB$  index. The RONN



model is an iterative learning algorithm and its main loop contains two smaller loops within itself. The first loop performs iterative adjustments of the node coordinates using the  $k$ -means algorithm until the nodes stabilize in their geometric positions. Coordinates of the nodes are calculated as arithmetic means of observations in each cluster. A cluster's mean-squared error ( $MSE$ ) is a measure of deviation of observations in the cluster from its center.

The second loop concerns 'dead nodes', i.e., cluster centers which have ended up without any observations in their vicinity. If there are no 'dead nodes', the cluster with the smallest number of observations is identified. The center of that cluster is then considered a 'dead node', while its observations are re-distributed to the closest clusters. During the iterative adjustments each 'dead node' observation is allocated to the node with the largest  $MSE$  and the new coordinates of this node are then given by:

$$v_i^{new} = v_{\max}^q + \delta, \quad i = 1, \dots, q \quad (4)$$

where  $v_{\max}^q$  is the location of the selected node among the  $q$  nodes with the largest  $MSE$  values,  $v_i^{new}$  is the position of the new node with respect to the center of a possible new cluster, and  $\delta = [\delta_1, \delta_2, \dots, \delta_n]^T$  are small random numbers. The procedure depicted by Eq. 4 is repeated for all 'dead nodes' until they are allocated to clusters. A detailed description of the RONN algorithm can be found in Kukolj et al. (2006).

#### 4.3 Feedforward artificial neural network (FF-ANN)

To explain the concept of an FF-ANN model, we will refer to Eq. 2 and our goal is to estimate the parameters of non-linear function ( $\phi$ ). In the context of the FF-ANN model, the parameters are called connection weights ( $\alpha_{ij}$  and  $\beta_j$ ) and node biases ( $\alpha_{j0}$  and  $\beta_0$ ). The model consists of three building blocks that are the input layer (where predictors are fed into the model), hidden layers (where the functional approximation or learning takes place) and the output layer (where the option price prediction is generated).

If we initially assume the model architecture involves only one hidden layer with  $q$  computational elements (nodes or neurons), the FF-D-ANN model is estimated as:

$$c_t = \phi \left( \beta_0 + \sum_{j=1}^q \beta_j \psi \left( \alpha_{j0} + \sum_{i=1}^k \alpha_{ij} x_{it} \right) \right) \quad (5)$$

In this example,  $q$  is the number of hidden nodes, where the single hidden and the output layers are characterized by two flexible classes of non-linearities:  $\psi$  and  $\phi$ , respectively.  $\alpha_{ij}$  and  $\beta_j$  denote appropriate connection weights between the adjacent layers. Subscripts 0 for  $\alpha$  and  $\beta$  stand for the biases.

The training algorithm is Adam, a first-order gradient-based optimizer based on adaptive estimates of lower-order moments. It is computationally efficient and suitable for large datasets with a lot of parameters to estimate. The default parameters follow those provided in Kingma and Ba (2015). The activation function types used in the hidden layer are sigmoid, with the linear function in the output layer. Validation data is used to select the optimal model architecture. The model will set apart 10% of the training data, will not train on it, and will evaluate the loss (i.e., mean-squared error) and any model metrics on this data at the end of each epoch. If, during the last 5 epochs the monitored quantity (validation loss) has no improvement, the training will be stopped.

#### 4.4 Support vector machine (SVM) model

SVM is a supervised learning model for prediction and classification, introduced by Vapnik (1999). The basic idea of this algorithm can be described by two steps: (a) transform training samples into a space of higher dimension using a non-linear mapping function; and (b) perform linear regression in the space of higher dimension in order to separate the data samples. Transforming data from original to the new (higher order) space is performed using a predetermined kernel function. The kernel function is defined as dot product of mapped data input vectors ( $x_{it}$ ,  $i = 1, \dots, k$ ) from the original space that enables computations of points in the new feature space without explicitly calculating the unknown non-linear mapping ( $\phi(x_{1t}, x_{2t}, \dots, x_{kt})$ ). In this implementation, we employ a second order polynomial for the initial kernel function ( $K(x_{it}, x)$ ). This function is redefined by a normalization, as given by the following expression:

$$\tilde{K}(x_{it}, x) = \frac{K(x_{it}, x)}{\sqrt{K(x_{it}, x_{it})K(x, x)}} \quad (6)$$

The normalization of the kernel function can be viewed as a simultaneous rescaling of data rows and columns to obtain a matrix with all diagonal entries set to one (Graf and Borer 2001). In the second step of the SVM algorithm (frequently called SV Regression—SVR), we perform data separation by constructing linear regression in the higher dimension feature space:

$$f(x) = \omega^T \phi(x) + b, \quad (7)$$

where  $\omega$  represents the vector of weights and  $b$  is a bias term. The optimization goal is to determine the trade-off between the flatness of  $f(x)$  while making sure that it has at most an  $\epsilon$  deviation between the obtained targets and training data outputs. In our work, we use the SVR implementation given in Shevade et al. (2000). The  $\epsilon$  parameter of the  $\epsilon$ -insensitive loss function is set to 0.001.

#### 4.5 Random forest (RF) model

A random forest (RF) is an ensemble machine learning technique introduced by Breiman (2001). It consists of a collection of regression trees whose averaged outputs determine the final prediction of the ensemble. The RF learning is based on a bootstrap aggregation (bagging) and random features subspace selection. Through a bagging procedure, each tree in the ensemble has randomly selected portions of training samples (with replacement) from the original dataset. In order to avoid possible correlations between constituent random trees and enhance the estimation performance of the RF model, the idea of a random features subspace is applied. As a result, each tree is grown on the basis of a randomly chosen input subset. For each node, the splitting algorithm searches over a subset of selected features to determine the best split point. The RF growth during the learning process is determined by two parameters: the size of the features subset used in each regression tree and the number of trees that form the forest. We tuned these two parameters to find the best regression RF model, as presented in Probst et al. (2019). Because of the small number of input features, a grid search was adopted as the simplest tuning strategy. Varying the number of trees on the [50, 150] interval and the features subset size from 2 to 6 showed that the RF with 80 trees and the features subset size of four has the lowest *RMSE* across the three data sub-periods considered.

In the process of constructing each decision tree, the RF learning algorithm uses the classification and regression trees (CART) learning algorithm (Breiman et al. 1984) which adopts the Gini index as the impurity-based measure. More precisely, the Gini index is used every time a split of a node is made on a certain variable. The learning algorithm is able to cope with numeric variables, characteristic for the problem considered in this study, by discretizing the continuous scale into two intervals. The optimal cut-off point is determined on the basis of evaluation of a threshold pool consisting of adjacent distinct values. The Gini index describes a decrease in the node impurity weighted by the probability of reaching that node:

$$G = \sum_{i=1}^C p_i (1 - p_i), \quad (8)$$

where  $p_i$  is the node probability calculated as the number of observations that reach the node divided by the total number of observations and  $C$  is the number of classes in the target variable. The variance reduction of a given tree node describes a decrease of the variance of the target variable due to the split at this node. Adding up and averaging the variance reductions for every node over all trees in the forest gives the value of relative importance of an input variable. Feature importance in the form of Gini index is commonly used to identify the contribution of individual predictors toward explaining the output (option price). Such an approach that relies on a decision tree classifier is considered intrinsically XAI in nature and it will assist us in obtaining the economic interpretation of our results.

## 5 Results

### 5.1 Comparative performance on sub-periods

We run our predictive models and price call options on the last 10% of data (testing set) across the three sub-periods: (1) pre-COVID-19, (2) COVID-19 crash, and (3) post-COVID-19 recovery. Our goal is to compare the forecast performance of the competing models (FF-ANN, SVM, RF, XGB, and the Black-Scholes benchmark) on each subsample and to also identify the most important predictors of option price fluctuations.<sup>7</sup> Table 2 shows the pricing accuracy reflected in the root mean-squared prediction error ( $RMSE$ ) of the models that were tested on out-of-sample data.

First, we find that the non-parametric option pricing models strongly outperformed the Black-Scholes model during the pre-COVID-19 period. Overall, the most accurate pricing performance is generated by the XGB model. The second best-performing model on all time periods is a RF that produces almost 70% lower out-of-sample  $RMSE$  than the Black-Scholes model over the first sub-sample. We can also conclude that the regime shift which took place in the second and third sub-samples to a certain extent limited learning and generalization abilities of non-parametric models. Surprisingly, the pricing accuracy of the Black-Scholes model improved during the pandemic, although it was still inferior to that of the RF and XGB models. Being a pre-specified non-linearity, the Black-Scholes model showed less sensitivity to a regime shift than the competing non-parametric models. This evidence is in accord with

<sup>7</sup> Based on the referee's recommendation, we also consider the extreme gradient boosting (XGB) model (Chen and Guestrin 2016) that is known for higher forecast accuracy relative to the RF model. We are extremely grateful to Zied Ftiti (Lead Guest Editor) and the three anonymous referees for their excellent comments and suggestions.

**Table 2** COMPARISON OF PREDICTIVE MODELS FOR THE SUB-PERIODS OF DATA

RMSE Statistic	pre-COVID-19	COVID-19	post-COVID-19
FF-ANN	0.0330	0.0652	0.0574
SVR	0.0363	0.0753	0.1264
RF	0.0209	0.0466	0.0541
XGB	0.0031	0.0361	0.0153
BS	0.0670	0.0479	0.0580

The data range is divided into three sub-samples: (1) January 1, 2020 to February 19, 2020 (pre-COVID-19): 12989 observations, (2) February 20, 2020 to March 23, 2020 (COVID-19 market crash): 19547 observations, and (3) March 24, 2020 to June 15, 2020 (post-COVID-19 recovery): 38909 observations. The last 10% of each sub-sample is held out-of-sample, then, the root mean-squared prediction error (*RMSE*) of each model is calculated on the testing data and reported in the table. The models that are considered are (1) Feedforward Artificial Neural Network (FF-ANN), (2) Support Vector Regression (SVR), (3) Random Forest (RF), Extreme Gradient Boosting (XGB), and (4) the Black-Scholes (BS) model

Kukolj et al. (2012) for the 2008 crisis. Another interesting finding is that, except for the XGB model, the pricing ability of the models diminished during the post-COVID-19 recovery period. As it could also be observed in Table 1, the uncertainty appeared to have been greater post-COVID-19 than during the actual crash, which impeded the option pricing accuracy.

Next, we will study the relative importance of inputs in the RF model. For such a purpose, we estimate the Gini indices (based on the average impurity decrease) while pricing options on each sub-sample of data. Table 3 displays the estimates for the inputs across the three sub-periods. It can be observed that the moneyness of options ( $S_t/K$ ) was among the top two most important inputs on each market regime. In fact, moneyness is the input that contributed to option pricing the most during extreme market volatility (COVID-19 and post-COVID-19 periods). Interestingly, during a steady market regime, options prices were the most sensitive to implied volatility, while this input was the second most important predictor during the post-COVID-19 recovery. We can conclude that when the market is relatively stable or on an upswing, implied volatility is important for option pricing. In contrast, during a market downswing, market participants react to excessive movements by focusing more on open interest (and even maturity) when determining option prices. This evidence of distressed behavior is a departure from the traditional no-arbitrage principle in option pricing where market-makers are able to perfectly hedge their inventories. Apparently, when the market is collapsing and is in a state of panic, open interest changes affect option demand imbalances, which, in turn, impacts option prices substantially.

Following another referee's suggestion, we supplement the XAI analysis with the Shapley additive explanations (SHAP) values for our set of inputs across the three subsamples (Lundberg and Lee 2017). The advantage of this method lies in its ability to explain the output of machine learning models. Essentially, the SHAP values show how much each predictor contributes (positively or negatively) to the output. Another benefit of the SHAP approach is its local interpretability where each observation is assigned its own set of SHAP values. First, we will estimate the SHAP values across the training data sets for the most accurate models (RF and XGB) by taking the mean absolute value of each feature.

Table 4 demonstrates the key role that traditional inputs, moneyness ( $S_t/K$ ) and time-to-maturity ( $\tau$ ), play in option pricing across all three market regimes. The fact that moneyness is the most informative input is consistent with the Gini indices from Table 3. However, according to the SHAP measures, implied volatility is not as important as the Gini indices suggested, but is consistently the third most important input. Also, when compared to Table 3,

**Table 3** FEATURE IMPORTANCE IN THE RF MODEL FOR THE SUB-PERIODS OF DATA

Gini index (# nodes)	pre-COVID-19	COVID-19	post-COVID-19
$\sigma_{IV}$	<b>278.23</b> (982)	13.70 (5972)	<b>85.08</b> (3310)
$S_t/K$	<b>256.69</b> (3316)	<b>107.02</b> (12955)	<b>219.59</b> (11355)
$\tau$	4.03 (674)	13.84 (5578)	7.86 (2832)
$OPI$	18.95 (805)	<b>22.24</b> (2862)	15.75 (1973)
$VOL$	11.39 (381)	5.97 (2130)	7.89 (1476)

The features (inputs) are as follows: implied volatility ( $\sigma_{IV}$ ), moneyness ( $S_t/K$ ), time-to-maturity ( $\tau$ ), open interest ( $OPI$ ), and volume ( $VOL$ ). The data range is divided into three sub-samples: (1) January 1, 2020 to February 19, 2020 (pre-COVID-19): 12989 observations, (2) February 20, 2020 to March 23, 2020 (COVID-19 market crash): 19547 observations, and (3) March 24, 2020 to June 15, 2020 (post-COVID-19 recovery): 38909 observations. Each cell in the table is a Gini index obtained from the RF model, based on the average impurity decrease, with the number of nodes using that attribute indicated in the parentheses (# nodes). The two inputs with the largest feature importance are bolded

**Table 4** SHAP IN THE RF AND XGB MODELS FOR THE SUB-PERIODS OF DATA

SHAP values	pre-COVID-19		COVID-19		post-COVID-19	
	RF	XGB	RF	XGB	RF	XGB
$\sigma_{IV}$	0.00005	0.00448	<b>0.00058</b>	0.01164	0.00002	0.00601
$S_t/K$	<b>0.10613</b>	<b>0.10649</b>	<b>0.08014</b>	<b>0.08294</b>	<b>0.14497</b>	<b>0.14691</b>
$\tau$	0.00012	<b>0.00810</b>	0.00047	<b>0.01464</b>	0.00008	<b>0.01514</b>
$OPI$	<b>0.00112</b>	0.00011	0.00018	0.00066	<b>0.00010</b>	0.00027
$VOL$	0.00010	0.00011	0.00048	0.00010	0.00005	0.00015

The inputs listed in the first column are as follows: implied volatility ( $\sigma_{IV}$ ), moneyness ( $S_t/K$ ), time-to-maturity ( $\tau$ ), open interest ( $OPI$ ), and volume ( $VOL$ ). The data range is divided into three sub-samples: (1) January 1, 2020 to February 19, 2020 (pre-COVID-19): 12989 observations, (2) February 20, 2020 to March 23, 2020 (COVID-19 market crash): 19547 observations, and (3) March 24, 2020 to June 15, 2020 (post-COVID-19 recovery): 38909 observations. Each cell in the table is the SHAP value, estimated across the training data sets (90% of observations) by the RF and XGB models and expressed as the mean absolute value of each feature. The two inputs with the largest importance are bolded

Table 4 suggests a diminished impact of open interest during the COVID-19 market crash. Nevertheless, relative to the other two sub-periods, it can be observed that open interest affects the output the most during the crash. As explained before, this shows that the markets experienced a greater degree of distress at this time.

For the sake of direct comparison between the Gini indices estimated by the RF and XGB methods, next, we show their normalized values side by side in Table 5. Clearly, the evidence for the RF model mirrors Table 3, while that for the XGB model is slightly different from the SHAP values in Table 4. More precisely, the Gini indices for the XGB model during the pre-COVID-19 period suggest that the implied volatility input was more important than time-to-maturity, which, in fact, reinforces the results for the RF model (i.e., the first column of Table 5). All other Gini values for the XGB model in Table 5 resemble the corresponding ones in Table 4. In all, we conclude that the XAI interpretation, as it may be expected, depends to a certain extent on the underlying non-linear modeling method (e.g., RF vs. XGB) as well as on the XAI technique that is applied (e.g., Gini vs. SHAP). When an input is highly dominant, its influence and interpretation is robust across methods, but, when an input is relatively

**Table 5** FEATURE IMPORTANCE IN THE XGB AND RF MODELS FOR THE SUB-PERIODS OF DATA

Gini indices	pre-COVID-19		COVID-19		post-COVID-19	
	RF	XGB	RF	XGB	RF	XGB
$\sigma_{IV}$	<b>0.10313</b>	<b>0.00070</b>	0.01201	0.00078	<b>0.05808</b>	0.00019
$S_t/K$	<b>0.88418</b>	<b>0.99909</b>	<b>0.96194</b>	<b>0.99764</b>	<b>0.93385</b>	<b>0.99884</b>
$\tau$	0.00086	0.00018	0.00693	<b>0.00137</b>	0.00490	<b>0.00072</b>
$OPI$	0.01175	0.00003	<b>0.01387</b>	0.00017	0.00266	0.00015
$VOL$	0.00008	0.00001	0.00525	0.00005	0.00050	0.00010

The inputs listed in the first column are as follows: implied volatility ( $\sigma_{IV}$ ), moneyness ( $S_t/K$ ), time-to-maturity ( $\tau$ ), open interest ( $OPI$ ), and volume ( $VOL$ ). The data range is divided into three sub-samples: (1) January 1, 2020 to February 19, 2020 (pre-COVID-19): 12989 observations, (2) February 20, 2020 to March 23, 2020 (COVID-19 market crash): 19547 observations, and (3) March 24, 2020 to June 15, 2020 (post-COVID-19 recovery): 38909 observations. Each cell in the table is a Gini index, estimated across the training data sets (90% of observations) by the RF and XGB models and normalized on the [0, 1] interval. The two inputs with the largest importance are bolded

uninformative, there exist differences in interpretation that we believe are complementary in nature as they provide further insights into the phenomenon of interest.

Figure 1 displays the SHAP summary plots for the XGB model that explain how each input contributes to the output and in which direction. This plot utilizes the training data and ranks features vertically in descending order according to the magnitude of their SHAP values. The color of each dot represents the impact direction of the feature on the model output where blue (red) rectangles represent inverse (direct) relationship. For all observations, their location along the x-axis shows whether the impact is associated with a higher or lower prediction. What stands out in Fig. 1 is the strong dominance of the moneyness input and its positive effect on options prices on all market regimes. The other inputs have a much weaker contribution which can be deduced from the high concentration of their SHAP values around zero. Noteworthy, it appears that implied volatility and time-to-maturity display mostly positive relationships, while open interest and volume exhibit negligible negative impacts.

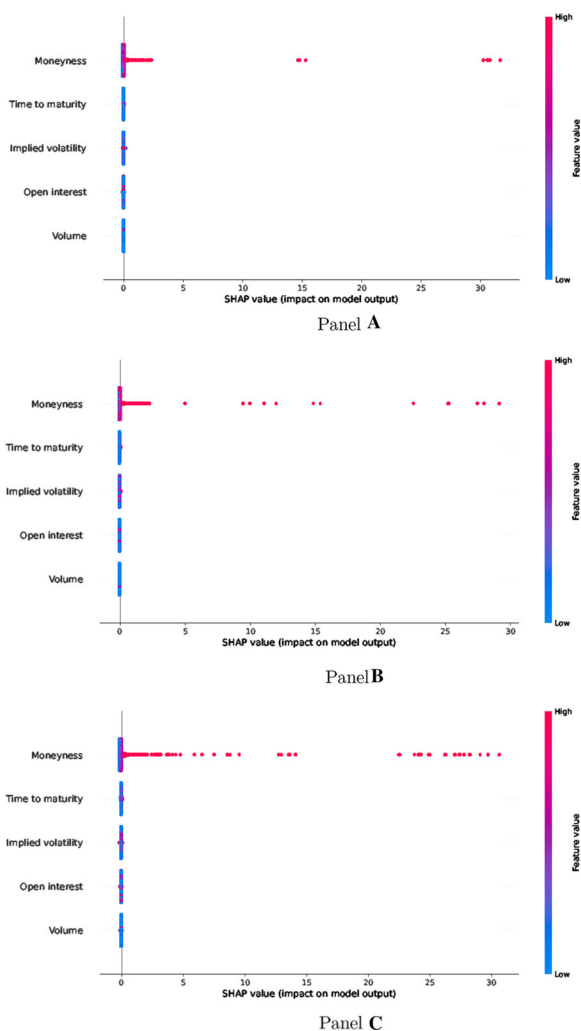
## 5.2 Clustering and feature importance

In this subsection, we perform optimal clustering of options on each market regime. By this, we provide a deeper analysis of the pricing accuracy and the relevance of model's inputs. We will also present an economic interpretation of our findings. Our first goal is to estimate the optimal number of clusters for each of the three time periods. Figure 2 plots the values of the  $DB$  index for various cluster choices. Clearly, the optimal number of clusters for the pre-COVID-19 market is six, while it is five for both COVID-19 and post-COVID-19 data. We conjecture that the pandemic changed the market microstructure to a more homogeneous trading behavior of the investors that originated in the convergence of their beliefs. Put differently, a change in the risk aversion of an average investor and an increased concentration thereof caused a reduction in the market “sub-regimes” (or clusters).

To locate the optimal clusters and their centers on the pre-COVID-19 data, we apply the RONN algorithm. For the visual illustration only, we reduce the five-dimensional space to a two-dimensional projection and display it in Fig. 3. The cluster boundaries are clearly defined and the clusters are reasonably spaced, further attesting to the validity of our approach. It

**Fig. 1** SHAP SUMMARY PLOTS.

Notes: In Panel A, the SHAP feature contributions are shown for the pre-COVID-19 period (January 1, 2020 to February 19, 2020). In Panel B, the SHAP feature contributions are shown for the COVID-19 market crash period (February 20, 2020 to March 23, 2020). In Panel C, the SHAP feature contributions are shown for the post-COVID-19 recovery period (March 24, 2020 to June 15, 2020). The aggregate SHAP values for each instance-input combination and their relationship to the output (option price) is denoted by red (direct relationship) or blue (inverse relationship) dots. The underlying model in all three panels is estimated with the XGB algorithm




is worthwhile to stress that the first two principal components cumulatively preserve about 95% of the original data variance in reduced space. When de-normalized to their original coordinates, the cluster centers for inputs ( $S_t/K$ ,  $\sigma_{IV}$ ,  $\tau$ ,  $VOL$ ,  $OPI$ ) are located in the original input space as follows:

- **Cluster 1**  $\rightarrow$  (1.22, 0.16, 1.71, 93.78, 2018.09);
- **Cluster 2**  $\rightarrow$  (0.90, 0.15, 2.87, 40.08, 131.75);
- **Cluster 3**  $\rightarrow$  (1.17, 0.10, 0.10, 2333.03, 50609.87);
- **Cluster 4**  $\rightarrow$  (1.00, 0.10, 0.16, 322.90, 2384.70);
- **Cluster 5**  $\rightarrow$  (1.03, 0.14, 0.83, 197.83, 3948.50);
- **Cluster 6**  $\rightarrow$  (1.03, 0.10, 0.16, 1419.17, 18355.29).

From the coordinates' values of the cluster centers, we identify cluster 1 as deep-ITM, long-term options, while cluster 2 are OTM, very long-term options. Cluster 3 refers to deep-ITM, short-term options with high volume and open interest. Cluster 4 is at-the-money,






**Fig. 2** SELECTION OF THE OPTIMAL NUMBER OF CLUSTERS. Notes: The selection is performed for three sub-samples: (1) January 1, 2020 to February 19, 2020 (pre-COVID-19): 12989 observations (dot-dashed line), (2) February 20, 2020 to March 23, 2020 (COVID-19 market crash): 19547 observations (solid line), and (3) March 24, 2020 to June 15, 2020 (post-COVID-19 recovery): 38909 observations (dashed line). The optimal number of clusters (horizontal axis) is where the Davies-Bouldin criteria (vertical axis) are at the minimum value

short-term options and cluster 5 is NTM options that are long-term. Finally, cluster 6 is NTM, short-term options with high volume and open interest. Of interest is to test the pricing accuracy of the RF model on each cluster and assess how clustering changes the importance of inputs. The pricing exercise and feature importance estimation will be run in-sample, followed by a comparative study across clusters.

Table 6 lists the Gini indices as well as the pricing errors ( $RMSE$ 's) for the six pre-COVID-19 clusters of data. In terms of pricing accuracy, we find that it was easier to price NTM, short-term options with relatively lower volatility and higher volume (as in clusters 4 and 6).  $RMSE$ 's appeared to have increased with maturity and moneyness. The largest pricing errors were produced in clusters 5 and 1, i.e., for long-term options with lower volume.

As for the relative importance of inputs, the Gini indices in Table 6 reveal that moneyness was always among the two top ranked features. In clusters 1 and 2, however, implied volatility emerged as the most important input. This is in line with the fact that clusters 1 and 2 contained options with the highest volatility of all clusters. Therefore, when pricing options in a cluster with large implied volatility, the relative importance of the volatility input increases. We also infer that open interest becomes a dominant input in clusters 3 and 6, i.e., for the options with high open interest. It is extraordinary to report that some options in the market were sensitive to demand-based shocks already in the pre-COVID-19 period. This surprising evidence was not observable in Table 3 because it examined all option contracts together in a single cluster. Lastly, both volume and maturity inputs do not seem informative in option pricing.

The optimal clusters and their centers for the COVID-19 period (reduced to a two-dimensional projection) are shown in Fig. 4. In this case, the first two principal components preserved in total about 90% of the original data variance in reduced space. The cluster cen-



**Fig. 3** PRE- COVID- 19: CLUSTERS AND THEIR CENTERS. Notes: The clusters are found by using the RONN algorithm for the pre-COVID-19 data (January 1, 2020 to February 19, 2020: 12989 observations). We reduce the five-dimensional space to a two-dimensional projection where PC#1 and PC#2 denote the first two principal components. Numbers 1-6 indicate the location of cluster centers in the reduced space


**Table 6** PRE- COVID- 19: FEATURE IMPORTANCE ACROSS CLUSTERS

Gini index (# nodes)	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
$\sigma_{IV}$	<b>712.57</b> (259)	<b>0.43</b> (696)	0.68 (91)	<b>0.11</b> (2994)	<b>83.60</b> (518)	0.02 (1173)
$S_t/K$	<b>639.10</b> (400)	<b>0.25</b> (1381)	<b>7.74</b> (849)	<b>1.12</b> (11571)	<b>229.39</b> (1164)	<b>1.10</b> (4710)
$\tau$	18.06 (44)	0.03 (196)	0.38 (197)	0.04 (3465)	15.40 (265)	0.03 (1527)
$OPI$	81.32 (60)	0.03 (264)	<b>1.97</b> (227)	0.09 (2785)	29.12 (261)	<b>0.11</b> (1258)
$VOL$	34.12 (23)	0.03 (124)	0.68 (131)	0.10 (1668)	34.87 (157)	0.10 (732)
$RMSE$	0.1080	0.0410	0.0135	0.0038	0.1619	0.0076

The inputs listed in the first column are as follows: implied volatility ( $\sigma_{IV}$ ), moneyness ( $S_t/K$ ), time-to-maturity ( $\tau$ ), open interest ( $OPI$ ), and volume ( $VOL$ ). The data is for the pre-COVID-19 time period (January 1, 2020 to February 19, 2020). The Gini indices are obtained from the RF model, based on the average impurity decrease, with the number of nodes using that attribute indicated in the parentheses (# nodes). The two inputs with the largest feature importance in a cluster are bolded.  $RMSE$  denotes the root mean-squared prediction error of the RF model that was tested on in-sample data

ters for inputs ( $S_t/K$ ,  $\sigma_{IV}$ ,  $\tau$ ,  $VOL$ ,  $OPI$ ) in the five-dimensional original space are at the following coordinates:

- **Cluster 1**  $\rightarrow$  (1.02, 0.21, 1.57, 99.38, 2283.53);
- **Cluster 2**  $\rightarrow$  (0.93, 0.24, 0.72, 242.90, 3710.52);
- **Cluster 3**  $\rightarrow$  (0.97, 0.43, 0.10, 3156.97, 35520.08);
- **Cluster 4**  $\rightarrow$  (0.90, 0.21, 2.76, 25.21, 199.81);
- **Cluster 5**  $\rightarrow$  (0.92, 0.36, 0.15, 388.79, 2900.18).




**Fig. 4** COVID- 19 CRASH: CLUSTERS AND THEIR CENTERS. Notes: The clusters are found by using the RONN algorithm for the COVID-19 data (February 20, 2020 to March 23, 2020: 19547 observations). We reduce the five-dimensional space to a two-dimensional projection where PC#1 and PC#2 denote the first two principal components. Numbers 1-5 indicate the location of cluster centers in the reduced space

**Table 7** COVID- 19: FEATURE IMPORTANCE ACROSS CLUSTERS

Gini index (# nodes)	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5
$\sigma_{IV}$	<b>353.46</b> (320)	<b>22.69</b> (1953)	0.09 (1076)	<b>0.33</b> (2272)	0.02 (22365)
$S_t/K$	<b>301.17</b> (840)	<b>33.73</b> (4893)	<b>3.37</b> (2209)	<b>0.49</b> (3665)	<b>0.43</b> (35127)
$\tau$	9.64 (154)	0.93 (1079)	0.07 (851)	0.08 (1083)	0.02 (18052)
$OPI$	29.08 (160)	1.64 (845)	<b>0.72</b> (557)	0.10 (907)	<b>0.05</b> (11878)
$VOL$	2.08 (87)	2.21 (638)	0.13 (502)	0.03 (759)	0.02 (11896)
$RMSE$	0.1689	0.0488	0.0051	0.0102	0.0029

The inputs listed in the first column are as follows: implied volatility ( $\sigma_{IV}$ ), moneyness ( $S_t/K$ ), time-to-maturity ( $\tau$ ), open interest ( $OPI$ ), and volume ( $VOL$ ). The data is for the COVID-19 time period (February 20, 2020 to March 23, 2020). The Gini indices are obtained from the RF model, based on the average impurity decrease, with the number of nodes using that attribute indicated in the parentheses (# nodes). The two inputs with the largest feature importance in a cluster are bolded.  $RMSE$  denotes the root mean-squared prediction error of the RF model that was tested on in-sample data

Thus, cluster 1 represents ITM, long-term options, while cluster 2 are OTM, long-term options with higher open interest. Cluster 3 are OTM, short-term options with high volatility and volume, and very high open interest. Cluster 4 are deep-OTM, long-term options with low volume and open interest. Cluster 5 are deep-OTM, short-term options with high volatility and open interest. The Gini indices and  $RMSE$ 's for the clusters can be found in Table 7. During the COVID-19 period, the highest pricing accuracy was reached in clusters 3 and 5, which are OTM, short-term options with high volatility and open interest. It can be generally stated that longer maturity is detrimental for the pricing accuracy, while other inputs do not have a clear systematic effect on the model's pricing ability across the clusters.



**Fig. 5** POST- COVID- 19 RECOVERY: CLUSTERS AND THEIR CENTERS. Notes: The clusters are found by using the RONN algorithm for the COVID-19 data (March 24, 2020 to June 15, 2020: 38909 observations). We reduce the five-dimensional space to a two-dimensional projection where PC#1 and PC#2 denote the first two principal components. Numbers 1-5 indicate the location of cluster centers in the reduced space

Table 7 also confirms that moneyness is the most useful input in all clusters, followed by the implied volatility and open interest. Similar to the pre-COVID-19 period, the open interest input gained importance in the clusters where high open interest options are concentrated. In light of our findings for the pre-COVID-19 period, the fact that option demand also drove option prices during the COVID-19 crash was not unexpected. Further, volume was relatively more important only in the third cluster that contains very high volume options. Again, the maturity input was of minor relevance for option pricing.

Finally, we plot the optimal clusters and their centers for the post-COVID-19 time period in Fig. 5. Here, the first two principal components cumulatively preserve about 94% of the data variance in new feature space. The cluster centers for inputs ( $S_t/K$ ,  $\sigma_{IV}$ ,  $\tau$ ,  $VOL$ ,  $OPI$ ) are distributed as follows:

- **Cluster 1**  $\rightarrow$  (1.00, 0.27, 0.14, 145.07, 1300.40);
- **Cluster 2**  $\rightarrow$  (0.99, 0.22, 2.62, 38.94, 469.04);
- **Cluster 3**  $\rightarrow$  (0.96, 0.25, 0.27, 783.05, 13851.57);
- **Cluster 4**  $\rightarrow$  (0.92, 0.22, 1.38, 108.25, 2287.64);
- **Cluster 5**  $\rightarrow$  (1.04, 0.27, 0.63, 113.28, 1901.41).

The cluster centers 1-5 above identify the following option types: (1) at-the-money, short-term options with high volatility and open interest, (2) at-the-money, very long-term options with low volume and open interest, (3) OTM, short-term options high volume and open interest, (4) deep-OTM, long-term options with high open interest, and (5) ITM, medium-term options with high open interest. With regard to the pricing performance in the individual clusters, presented in Table 8, the lowest (highest) error was generated in cluster 3 (4). In accord with some of our previous evidence, pricing errors increase with both moneyness and maturity. Nonetheless, the effect from the remaining inputs on the  $RMSE$ 's was not significant enough for a stylized economic interpretation.

**Table 8** POST- COVID- 19: FEATURE IMPORTANCE ACROSS CLUSTERS

Gini index (# nodes)	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5
$\sigma_{IV}$	5.76 (2075)	<b>3.80</b> (1473)	<b>0.14</b> (2448)	<b>23.85</b> (1130)	<b>158.95</b> (1729)
$S_t/K$	<b>144.05</b> (8416)	<b>2.84</b> (4071)	<b>0.82</b> (5320)	<b>19.15</b> (2915)	<b>145.63</b> (5286)
$\tau$	5.65 (1869)	0.24 (819)	0.03 (3064)	1.00 (614)	7.35 (912)
$OPI$	<b>14.23</b> (1663)	0.93 (751)	0.06 (1425)	6.38 (414)	13.12 (916)
$VOL$	6.58 (1068)	0.16 (409)	0.02 (1450)	4.21 (323)	4.12 (715)
$RMSE$	0.0446	0.0101	0.0097	0.0805	0.0592

The inputs listed in the first column are as follows: implied volatility ( $\sigma_{IV}$ ), moneyness ( $S_t/K$ ), time-to-maturity ( $\tau$ ), open interest ( $OPI$ ), and volume ( $VOL$ ). The data is for the COVID-19 time period (March 24, 2020 to June 15, 2020). The Gini indices are obtained from the RF model, based on the average impurity decrease, with the number of nodes using that attribute indicated in the parentheses (# nodes). The two inputs with the largest feature importance in a cluster are bolded.  $RMSE$  denotes the root mean-squared prediction error of the RF model that was tested on in-sample data

In contrast to the COVID-19 period, implied volatility was the most dominant input in three of the five option clusters. Although moneyness was always among the top two inputs in terms of the Gini index, it was the strongest only in clusters 1 and 3. Hence, clustering provides new findings that were not apparent from Table 3 that analyzed all post-COVID-19 options together. In particular, we demonstrate that volatility was a more important input than moneyness during the market recovery. In all likelihood, among other risk factors, volatility risk was the most significant and it commanded greater option prices in the post-COVID-19 period. Furthermore, demand-induced shocks were not as pronounced at that time, as they had been in the first two periods. Open interest was among the two highest Gini indices only in cluster 1, but, it is noteworthy, that it was consistently the third most important input in other clusters. As observed before, volume and time-to-maturity were the least useful for pricing in all clusters.

## 6 Conclusions

This paper tackles the problem of the lack of parametric transparency and economic interpretability in non-parametric option pricing models by utilizing data clustering and XAI. We also recognize that in real-world markets options are hedged imperfectly. Consequently, we depart from the Black-Scholes framework (where prices are determined by no-arbitrage) and expand the model with additional inputs—open interest and volume—that reflect market demand and sentiment. By concentrating on the most recent S&P 500 index options data that include the COVID-19 crash and post-COVID-19 recovery, we provide unique insights into the option market microstructure and the behavior of traders exposed to multiple regime shifts.

First, we perform a comparative option pricing exercise across various model specifications and time periods. The evidence shows that non-parametric techniques easily dominate the Black-Scholes model in a stable, pre-COVID-19 market. On the other hand, in the volatile, pandemic market, the Black-Scholes model delivers pricing accuracy that is comparable to non-parametric models. This is consistent with Garcia and Gençay (2000) and Kukulj et al. (2012) that reported relative success of the Black-Scholes model during regime shifts in

1987 and 2008, respectively. Overall, on all three sub-periods (pre-COVID-19, COVID-19, post-COVID-19), the XGB and RF models produce the most accurate option prices out-of-sample.

We also find that non-parametric option pricing (and the resulting economic interpretation) depends on the nature of non-parametric models and the XAI technique that is applied. More specifically, when an option pricing model's input is highly important (e.g., moneyness), its contribution and interpretation is invariant to model specification and time period. In contrast, when an input is of minor non-linear importance (e.g., volume), its significance may vary. Although such findings could be perceived as a limitation of our study, it is worthwhile to note that, when we combine the evidence from multiple methods, we are able to obtain a deeper grasp of the forces behind the data generating process of option prices. Our findings point to the dominant role played by the moneyness input in option pricing. The other inputs are significantly less important, but the influence of time-to-maturity and implied volatility is frequently present across models and market regimes. An intriguing evidence we report is that certain options were sensitive to demand-based shocks reflected in the open interest input already in the pre-COVID-19 period. Moreover, during the COVID-19 market crash, we observe an increased importance of open interest which suggests an emergence of liquidity and demand shocks during market distress. More specifically, on all time periods, the importance of open interest was greater in option clusters with high implied volatility and open interest (and to a certain extent moneyness). This in general corroborates the evidence of positive relationship between net demand for options and excess implied volatility found in Gârleanu et al. (2009).

The models and methods proposed in this paper help improve our understanding of the dynamics of option prices and their underlying risk factors. Our approach could be generalized to any financial derivative, providing sufficient market information is available. Within the context of stock options, future research efforts might concentrate on explaining idiosyncratic and systematic risk factors inherent in option contracts across various industries (as in, e.g., Ramelli and Wagner, 2020). Another potentially interesting research avenue is to employ high-frequency data where premia for stochastic volatility or jump risk could also contribute to option price movements.<sup>8</sup> It is of interest to policy-makers to measure how much variability in option (or, more generally, asset) prices can be attributed to non-traditional factors such as open interest and other potential gauges of market sentiment, especially at crisis times. Also, ignoring market microstructure effects (i.e., trader behavior) inferred from our clustering method may lead policy-makers to wrong conclusions about the effectiveness of a particular policy aimed at calming volatile markets.

In all, we demonstrate clear benefits of clustering and XAI to our understanding of the option market mechanisms, especially during periods of extreme uncertainty. Our empirical explorations complement the extant literature on the econometrics of option pricing. Specifically, we shed more light on the potential latent variables and complex non-linearities of the pricing model (Garcia et al. 2010).

## References

Agliardi, E., & Agliardi, R. (2009). Fuzzy defaultable bonds. *Fuzzy Sets and Systems*, 160, 2597–2607.

<sup>8</sup> High-frequency market frictions and price jumps are important potential sources of option pricing model risk. One of our model's limitations is working with daily data and, thus, aggregating intraday, high-frequency risk factors.

- Ait-Sahalia, Y., & Lo, A. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance*, 53, 499–547.
- Akhtaruzzaman, M., Boubaker, S., Chiah, M., & Zhong, A. (2021). COVID-19 and oil price risk exposure. *Finance Research Letters*, 42, 101882.
- Akhtaruzzaman, M., Boubaker, S., Lucey, B. M., & Sensoy, A. (2021). Is gold a hedge or a safe-haven asset in the COVID-19 crisis? *Economic Modelling*, 102, 105588.
- Akhtaruzzaman, M., Boubaker, S., & Sensoy, A. (2021). Financial contagion during COVID-19 crisis. *Finance Research Letters*, 38, 101604.
- Andreou, P., Charalambous, C., & Martzoukos, S. (2008). Pricing and trading European options by combining artificial neural networks and parametric models with implied parameters. *European Journal of Operational Research*, 185, 1415–1433.
- Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. *Journal of Finance*, 52, 2003–2049.
- Barndorff-Nielsen, O., & Shephard, N. (2001). Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society B*, 63, 167–241.
- Barone-Adesi, G., Engle, R. F., & Mancini, L. (2008). A GARCH Option Pricing Model with Filtered Historical Simulation. *The Review of Financial Studies*, 21, 1223–1258.
- Bates, D. S. (1991). The crash of '87: Was it expected? The evidence from options markets. *Journal of Finance*, 46, 1009–1044.
- Bates, D. S. (2000). Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*, 94, 181–238.
- Bates, D. S. (2012). U.S. stock market crash risk, 1926–2010. *Journal of Financial Economics*, 105, 229–259.
- Bellalah, M., Hakim, A., Si, K., & Zhang, D. (2020). Long term optimal investment with regime switching: inflation, information and short sales. *Annals of Operations Research* 0, 1–14. On-line.
- Bhat, H., & Kumar, N. (2012). Option pricing under a normal mixture distribution derived from the Markov tree model. *European Journal of Operational Research*, 223, 762–774.
- Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81, 637–659.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45, 5–32.
- Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. J. (1984). *Classification and regression trees*. Monterey, CA: Wadsworth & Brooks.
- Buehler, H., Gonon, L., Teichmann, J., & Wood, B. (2019). Deep hedging. *Quantitative Finance*, 19, 1271–1291.
- Cai, N., & Kou, S. G. (2011). Option pricing under a mixed-exponential jump diffusion model. *Management Science*, 57, 2067–2081.
- Calvet, L. E., Fearnley, M., Fisher, A. J., & Leippold, M. (2015). What is beneath the surface? Option pricing with multifrequency latent states. *Journal of Econometrics*, 187, 498–511.
- Cao, Y., Liu, X., & Zhai, J. (2021). Option valuation under no-arbitrage constraints with neural networks. *European Journal of Operational Research*, 293, 361–374.
- Carr, P., & Wu, L. (2004). Time-changed Lévy processes and option pricing. *Journal of Financial Economics*, 71, 113–141.
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, NY, USA. pp. 785–794.
- Christoffersen, P., Heston, S., & Jacobs, K. (2009). The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well. *Management Science*, 55, 1914–1932.
- Culkin, R., & Das, S. R. (2017). Machine learning in finance: the case of deep learning for option pricing. *Journal of Investment Management*, 15, 92–100.
- Davies, D. L., & Bouldin, D. W. (1979). A cluster separation measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1, 224–227.
- Ding, L. (2018). Human knowledge in constructing AI systems - Neural logic networks approach towards an explainable AI. *Procedia Computer Science*, 126, 1561–1570.
- Driouchi, T., Trigeorgis, L., & So, R. H. Y. (2018). Option implied ambiguity and its information content: Evidence from the subprime crisis. *Annals of Operations Research*, 262, 463–491.
- Eriksson, A., Ghysels, E., & Wang, F. (2009). The normal inverse Gaussian distribution and the pricing of derivatives. *Journal of Derivatives*, 16, 23–38.
- Fan, J., & Mancini, L. (2009). Option pricing with model-guided nonparametric methods. *Journal of the American Statistical Association*, 104, 1351–1372.
- Fulop, A., Li, J., & Yu, J. (2014). Self-exciting jumps, learning, and asset pricing implications. *The Review of Financial Studies*, 28, 876–912.



- Garcia, R., & Gençay, R. (2000). Pricing and hedging derivative securities with neural networks and a homogeneity hint. *Journal of Econometrics*, 94, 93–115.
- Garcia, R., Ghysels, E., & Renault, E. (2010). The Econometrics of Option Pricing. North-Holland, San Diego. volume 1 of *Handbooks in Finance*. pp. 479–552.
- Gârleanu, N., Pedersen, L. H., & Poteshman, A. M. (2009). Demand-based option pricing. *The Review of Financial Studies*, 22, 4259–4299.
- Gaß, M., Glau, K., Mahlstedt, M., & Mair, M. (2018). Chebyshev interpolation for parametric option pricing. *Finance and Stochastics*, 22, 701–731.
- Gençay, R., & Altay-Salih, A. (2003). Degree of mispricing with the Black-Scholes model and nonparametric cures. *Annals of Economics and Finance*, 4, 73–101.
- Gençay, R., & Gibson, R. (2009). Model risk for European-style stock index options. *IEEE Transactions on Neural Networks*, 18, 193–202.
- Gençay, R., & Qi, M. (2001). Pricing and hedging derivative securities with neural networks: Bayesian regularization, early stopping and bagging. *IEEE Transactions on Neural Networks*, 12, 726–734.
- Gradojevic, N. (2016). Multi-criteria classification for pricing European options. *Studies in Nonlinear Dynamics & Econometrics*, 20, 123–139.
- Gradojevic, N., Gençay, R., & Kukolj, D. (2009). Option pricing with modular neural networks. *IEEE Transactions on Neural Networks*, 20, 626–637.
- Gradojevic, N., Gençay, R., & Kukolj, D. (2011). Clustering and classification in option pricing. *Review of Economic Analysis*, 3, 109–128.
- Graf, A. B., & Borer, S. (2001). Normalization in support vector machines. In B. Radig & S. Florczyk (Eds.), *Pattern Recognition* (pp. 277–282). Heidelberg: Springer, Berlin Heidelberg, Berlin.
- Guidolin, M., & Hansen, E. (2016). Pricing S&P 500 index options: A conditional semi-nonparametric approach. *Journal of Futures Markets*, 36, 217–239.
- Harjoto, M. A., Rossi, F., & Paglia, J. K. (2021). COVID-19: stock market reactions to the shock and the stimulus. *Applied Economics Letters*, 28, 795–801.
- He, X. J., & Zhu, S. P. (2016). An analytical approximation formula for European option pricing under a new stochastic volatility model with regime-switching. *Journal of Economic Dynamics and Control*, 71, 77–85.
- Heston, S. L. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The Review of Financial Studies*, 6, 327–343.
- Hutchinson, J. M., Lo, A. W., & Poggio, T. (1994). A nonparametric approach to pricing and hedging derivative securities via learning networks. *Journal of Finance*, 49, 851–889.
- Jang, H., & Lee, J. (2019). Generative Bayesian neural network model for risk-neutral pricing of American index options. *Quantitative Finance*, 19, 587–603.
- Kingma, D.P., & Ba, J. (2015). Adam: A method for stochastic optimization, pp. 1–13. Proceedings of the 3rd International Conference on Learning Representations (ICLR).
- Kukolj, D., Atlagic, B., & Petrov, M. (2006). Data clustering using a reorganizing neural network. *Cybernetics and Systems*, 37, 779–790.
- Kukolj, D., Gradojevic, N., & Lento, C. (2012). Improving non-parametric option pricing during the financial crisis, in: 2012 IEEE Conference on Computational Intelligence for Financial Engineering and Economics, CIFEr 2012 - Proceedings, pp. 93–99.
- Kukolj, D., & Levi, E. (2004). Identification of complex systems based on neural and Takagi-Sugeno fuzzy model. *IEEE Transactions on Systems, Man, and Cybernetics. Part B (Cybernetics)*, 34, 272–282.
- Lento, C., & Gradojevic, N. (2021). S&P 500 index price spillovers around the COVID-19 market meltdown. *Journal of Risk and Financial Management* 14.
- Liu, H., Wang, Y., He, D., & Wang, C. (2020). Short term response of Chinese stock markets to the outbreak of COVID-19. *Applied Economics*, 52, 5859–5872.
- Liu, X., Cao, Y., Ma, C., & Shen, L. (2019). Wavelet-based option pricing: An empirical study. *European Journal of Operational Research*, 272, 1132–1142.
- Lundberg, S.M., & Lee, S.I. (2017). A unified approach to interpreting model predictions, in: Proceedings of the 31st International Conference on Neural Information Processing Systems, Curran Associates Inc., Red Hook, NY, USA. pp. 4768–4777.
- Luo, X., Xiang, K., & Ding, C. (2018). American option pricing under financial crisis. *Applied Stochastic Models in Business and Industry*, 34, 597–606.
- Mazur, M., Dang, M., & Vega, M. (2021). COVID-19 and the March 2020 stock market crash. *Evidence from S&P500. Finance Research Letters*, 38, 101690.
- Melick, W. R., & Thomas, C. P. (1997). Recovering an asset's implied PDF from option prices: An application to crude oil during the gulf crisis. *Journal of Financial and Quantitative Analysis*, 32, 91–115.
- Morelli, G., & Petrella, L. (2021). Option pricing, zero lower bound, and COVID-19. *Risks* 9.

- Probst, P., Wright, M. N., & Boulesteix, A. L. (2019). Hyperparameters and tuning strategies for random forest. *WIREs Data Mining and Knowledge Discovery*, 9, e1301.
- Qi, M., & Maddala, G. S. (1996). Option pricing using artificial neural networks: The case of S&P 500 index call options, in: Refenes, A.P.N., Abu-Mostafa, Y., Moody, J., Weigend, A. (Eds.), *Neural Networks in Financial Engineering: Proceedings of the Third International Conference on Neural Networks in the Capital Markets*, World Scientific, New York. pp. 78–91.
- Ramelli, S., & Wagner, A. F. (2020). Feverish stock price reactions to COVID-19. *The Review of Corporate Finance Studies*, 9, 622–655.
- Renault, E. (2010). Econometrics of option pricing. *Wiley, New York*, 2, 518–528.
- Ruf, J., & Wang, W. (2020). Neural networks for option pricing and hedging: a literature review. *Journal of Computational Finance*, 24, 1–46.
- Sharif, A., Aloui, C., & Yarovaya, L. (2020). COVID-19 pandemic, oil prices, stock market, geopolitical risk and policy uncertainty nexus in the us economy: Fresh evidence from the wavelet-based approach. *International Review of Financial Analysis*, 70, 101496.
- Shevade, S. K., Keerthi, S. S., Bhattacharyya, C., & Murthy, K. R. K. (2000). Improvements to the SMO algorithm for SVM regression. *IEEE Transactions on Neural Networks*, 11, 1188–1193.
- von Spreckelsen, C., von Mettenheim, H. J., & Breitner, M. H. (2014). Real-time pricing and hedging of options on currency futures with artificial neural networks. *Journal of Forecasting*, 33, 419–432.
- Vapnik, V. N. (1999). An overview of statistical learning theory. *IEEE Transactions on Neural Networks*, 10, 988–999.
- Vercammen, J. (2020). Information-rich wheat markets in the early days of COVID-19. *Canadian Journal of Agricultural Economics*, 68, 177–184.
- Wong, H. Y., & Lo, Y. W. (2009). Option pricing with mean reversion and stochastic volatility. *European Journal of Operational Research*, 197, 179–187.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.