

The forecasting power of short-term options*

Arthur Bóök[†] Juan F. Imbet[‡] Martin Reinke[§] Carlo Sala[¶]

September 3, 2022

Abstract

We propose robust option-implied measures of conditional volatility, skewness and kurtosis based upon quantiles and expectiles inferred from weekly options on the S&P 500. All quantities are by construction forward-looking and estimated non-parametrically through a novel robust and arbitrage-free natural smoothing spline technique that produces quick to estimate volatility smiles. We find that some of the option-implied robust indicators exhibit short-, medium- and long-term predictive ability for the U.S. equity risk premium and higher moments, both in- and out-of-sample, which outperform equal indicators inferred from historical returns.

Keywords: Option pricing, Volatility smile, Quantiles, Weekly options, Forecasting.

JEL classification: G10, G13, G14, G17.

*We are grateful for helpful comments and suggestions from seminar participants at the Finance Forum 2022, FMARC 2022, MAF 2022, the XXII Workshop on Quantitative Finance, the Munich Finance Day 2022, the FMARC21 and the UIB. There is an accompanying web-app for the paper which illustrates the estimation results. The app can be accessed here: <https://share.streamlit.io/mreinke1/app-quantile-expectile/main/app-quantile-expectile.py>

[†]Department of Financial Management and Control, Universitat Ramon Llull, ESADE, Avenida de Torreblanca 59, 08172 Sant Cugat, Barcelona, Spain; E-mail: arthuremilalexis.book@alumni.esade.edu.

[‡]Université Paris-Dauphine, PSL Research University, DRM, Pl. du Maréchal de Lattre de Tassigny, 75016 Paris, France; E-mail: juan.imbet@dauphine.psl.eu.

[§]Institute for Finance and Banking, LMU Munich School of Management, Ludwigstrasse 28, 80539 Munich, Germany; E-mail: reinke@lmu.de.

[¶]Department of Financial Management and Control, Universitat Ramon Llull, ESADE, Avenida de Torreblanca 59, 08172 Sant Cugat, Barcelona, Spain; E-mail: carlo.sala@esade.edu. Financial support from the AGAUR - SGR 2017-640 grant and from the Spanish Ministry of Science and Innovation - PID2019-1064656GB-I00/AEI / 10.13039/501100011033 are gratefully acknowledged.

1 Introduction

Option contracts are forward-looking financial assets that embed the investors' expectations for different future states of the world (strike prices) at different time horizons (time-to-maturity). These expectations naturally contain the investors' risk preferences, which determine the investors' required expected rate of return for any asset pricing framework. As a consequence, functions of option market data might contain valuable information for forecasting future returns. Using weekly options written on the S&P 500, we derive short-term (weekly) option-implied quantiles and expectiles, and check their forecasting power through different conditional and forward-looking robust indicators. More precisely, after calibrating the option prices through a novel natural smoothing spline technique, that produces arbitrage-free and robust volatility smiles, we infer option-implied conditional quantiles and expectiles non-parametrically. From them, we calculate different robust option-implied measures of conditional dispersion (volatility), asymmetry (skewness), and flatness (kurtosis) and test their forecasting power. We find that some of these indicators have short-, medium- and long-term forecasting power, both in- and out-of-sample, that is not present in the same indicators once estimated using realized returns in the underlying asset.

Being extremely noisy, assessing the investors' beliefs from option-market data is a delicate task which is highly dependent on the pricing model used in estimation. Moreover, the forecasting accuracy of option-implied quantities also depends on the absence of arbitrage in the related set of option prices. Overall, any option estimation thus requires absence of arbitrage in the option prices used, and correction for the noise that illiquid options could carry in the estimation (Hentschel (2003)).¹ We tackle this issue by proposing a novel approach to estimate arbitrage-free volatility smiles (Section 3). We refer to such estimation approach as the BIRS² approach. The BIRS is made of two steps: first, we perform a cubic interpolation over the implied volatility/delta space that gives a dense pre-estimate of the volatility smile (Section 3.2). Second, we solve a quadratic program enabling us to compute the linear and quadratic terms of a natural spline that yields arbitrage-free points over the full volatility smile (Section 3.3). Importantly, the two steps are fully complementary. On one hand, the implied volatility smile produced in the first step provides a good fit on the original data, but does not ensure the absence of arbitrages. On the other hand, the piece-wise quadratic polynomial estimated in the second step constrains all static arbitrages, but requires an accurate and dense pre-estimation to reduce the error of the quadratic approximation between the spline knots.

¹As documented by Hentschel (2003), implied volatility estimates are noisy and prone to bias, thus compromising results linked to them, unless necessary pre-processing is conducted.

²The name BIRS is due to the surnames of the authors, in alphabetical order.

Combining these two elements, we obtain a robust and arbitrage-free implied volatility smile, which is then converted into prices that we will use to perform different forecasting exercises.

To assess the validity of the BIRS approach, in Section 4 we benchmark it on two popular non-parametric approaches, namely the Positive Convolution Approximation (PCA) of Bondarenko (2003), and the Fast and Stable (FS) approach of Jackwerth (2004). The PCA uses convolutions to estimate the risk-neutral distribution, while the FS is a curve fitting technique on the implied volatility. Our approach is also fully non-parametric, allowing us a fair comparison across models.³ We find that the BIRS approach exhibits better performance in recovering both the implied volatility smile and the observed option prices and from both real and simulated data. More precisely, using real data we test which approach better recovers (has a lower loss function) the implied volatility smile and option prices (Section 4.1). For the simulated data, we perform a Monte Carlo experiment to test which technique performs better in recovering the option-implied cumulative distribution function in period of high and low volatility. For both the exercises, and in both situations (high and low volatility), the BIRS approach exhibits superior performance with respect to the PCA and FS approaches.

Having removed possibly noisy and mispriced data, we test if weekly options possess some forecasting power. Since deep out-of-the-money options might not be traded, we propose quantile- (or expectile-) based robust measures, which allow us to infer moments of the distribution without using the entire distribution. The forecasting analysis is thus divided in three steps. First, from the BIRS estimates, we infer conditional option-implied quantiles (Section 5.2) and option-implied expectiles (Section 5.3). Second, from the estimated quantities, we infer robust coefficients of conditional dispersion, asymmetry and flatness (Section 5.5). Finally, we test their forecasting power on realized moments of S&P 500 returns over different time horizons. Our results show that some robust option-implied indicators have in- and out-of-sample forecasting power on the first two realized moments of S&P returns (Section 6). Option-implied indicators of dispersion perform well in-sample, with some 7 days-ahead R^2 around 2%, and increasing with the time-horizon of the estimates. Results are confirmed out-of-sample for horizons of 26 weeks, where the option-implied conditional quantile and expectile dispersion indicators produce an R_{OOS}^2 (Campbell and Thompson (2007)) using historical indicators as a benchmark of above 18%. We also find that the option-implied conditional quantile and expectile dispersion indicators R_{OOS}^2 in predicting volatility increases up to 52% in one week ahead regressions. Aligned with the longer-term literature, we also document the lack in time-series

³The choice of the PCA and FS as benchmarks is twofold. First, it is due to their ability in recovering different option-implied quantities, demonstrated by their extensive use in the literature for a number of empirical and theoretical applications. Second, despite being all non-parametric, the PCA, the FS and the BIRS approaches are among them fairly different techniques. This diversity produces an heterogeneous and more complete analysis.

forecasting power in the higher moments of the distribution. Our short-term results complement the existing literature which tells us that for predicting the market index return (excess return, risk premiums), only a few implied quantities work: (1) ex-ante variance risk premium (VRP) (Carr and Wu (2009), Bollerslev et al. (2009), and some follow-up papers splitting the VRP into semi-variance risk premiums (for VRP it is relatively short-term predictability, i.e., one to three months), which worked well only from 1996 to 2015-2017, and lost significance afterward; (2) implied variance, esp. for simple returns, which is used to construct the variance-based lower bound of expected excess market returns as in Martin (2017); (3) a combination of higher-order moments forming the generalized lower bound as in Chabi-Yo and Loudis (2020); (4) implied correlation, i.e. Driessen et al. (2009) or Buss and Vilkov (2012). However, there is no robust evidence that higher-order implied moments quantifying distribution asymmetry or tail-fatness predict market returns.⁴

Importantly our results are both economically and statistically robust. Economically it makes sense for the expected return (volatility) to find higher (lower) forecasting power as the time-horizons increases. It is in fact well-known that forecasting the expected returns at short time horizons might be very difficult, due to the different price pressure that could make the price diverge from the fundamental (estimated) one. Differently, the decreasing forecasting power of the market volatility at long time horizon is again economically grounded, due to the well-known clustering property of volatility (Engle (1982) and Bollerslev (1986)). From a statistical viewpoint, we estimate Newey-West standard errors (Newey and West (1987)) with lags equal to the number of weeks between observations.

While historically observed price-paths do not disclose the underlying process from which they are generated, option prices enable the re-creation of the expected price distributions in the form of implied densities (Linn et al. (2017), Barone-Adesi et al. (2020)). Specifically, stock-based indicators are valid only if markets are ergodic and stationary as, due to the impossibility of the model to react quickly to changing market scenarios, it is almost impossible to forecast expected return. This is especially true at shorter time horizons. The non-parametric and option-based nature of the proposed robust coefficients allows us to overcome these issues. By construction, options are forward-looking financial assets, and their liquidity in major markets has become high enough to provide timely estimates with non-parametric models. Most of these estimates follow from the seminal papers of Breeden and Litzenberger (1978) and Banz and Miller (1978), which show how the expected future states of the underlying price can be inferred by differentiating the option prices with respect to their strike prices. These quan-

⁴Recently, Alexiou and Rompolis (2022) condense option implied higher moments into a single score and sort companies based on it into portfolios. They find that their constructed portfolios yield statistically significant return.

tities have been found to provide forward-looking information of predictive value for future realizations of index returns (see, among others, the review of Figlewski (2018), or the survey of Christoffersen et al. (2011) and references therein). Despite the growing literature on the use of option market data for predictions, not much has been said on the forecasting power of short-term option-implied quantiles and expectiles.

The aforementioned robust quantities - conditional dispersion indicators, conditional asymmetry indicators and conditional flatness indicators - are all derived non-parametrically, with the goal to preserve the subtleties of the shapes of the empirical data. Following Barone-Adesi (2016) and Bellini et al. (2018) we derive non-parametric option-implied quantiles and expectiles, respectively. Barone-Adesi (2016) shows how to use option-implied *quantiles* to infer option-implied risk measures. The empirical analyses of Barone-Adesi et al. (2019), Barone-Adesi et al. (2019) and Molino and Sala (2020), show that quantile-based option-implied value at risk (VaR) and conditional value at risk (CVaR) are good alternatives to classical risk measures based on historical returns. In particular, these indicators perform well when mostly needed - that is in periods of high volatility - when the statistical properties of the underlying deviate the most from the past, and cannot be captured using a large amount of historical returns.

Bellini et al. (2018) show how to infer option-implied *expectiles*. Bellini et al. (2021) test some properties of the S&P 500 index option-implied expectiles, confirming its ability in producing sensible risk measures without the need of a huge amount of data. We contribute to these literatures by proposing new option-based robust indicators and testing the out-of-sample performances of them.⁵

In terms of out-of-sample forecasting power Metaxoglou and Smith (2017) show that their monthly State Prices of Conditional Quantiles (SPOQ) have a significative time series forecasting power at long time horizons (18 to 24 months). Conrad et al. (2012) focus on the cross-section by inferring the ex-ante higher moments of the underlying individual securities' risk-neutral distribution and find that they strongly related to future returns. We differ from Metaxoglou and Smith (2017) and Conrad et al. (2012) by focusing at short time-horizon.⁶

We focus at short-term for a an interesting and natural reasons. While due to a lack of liquidity the option literature has always (and correctly) discarded options with short time-to-maturity⁷, weekly options changed the classical paradigm existing in the option literature. The liquidity

⁵Moreover, while some of the cited papers lay the ground to the theoretical foundation on option-implied quantiles and expectiles, while others test some of their properties, none of them consider the importance of a solid volatility smile for the estimation of their quantities.

⁶Another difference with Conrad et al. (2012) is that we focus on the time series, while they focus on the cross-section. Not surprisingly, they find some forecasting power in the higher moments, which is not present in the S&P 500 time series.

⁷Usually, only options up to 10 days have been considered

of these assets has increased remarkably in the recent years and, as explored in Section 2, accounts now for around the 40 to 50% of the entire total liquidity of the S&P 500 Index options. This allows us to perform for the first time a short-term analysis of the forecasting power of option-implied quantiles and expectiles.

Moreover, as aforementioned, forecasting returns at short-time horizons is of great challenge. While it is reasonable to assume that the option-implied information in its entirety might not be fully mapped into the physical measure, this bias is a decreasing function of the time horizon in consideration.⁸ The use of short-term options thus allows us to forecast returns at short time horizon, reducing the mapping problem as much as possible. Having said that, it is worth stressing again how *all* the econometric models and tests proposed in this paper can be extended to *any* time horizon and with no loss of generality.

2 Dataset

The empirical analysis of this paper is based upon data provided by OptionMetrics. We use End-of-week (EOW) Friday weekly S&P 500 index options for the period January 2, 2011 to December 31, 2021. Option prices are recorded given their mid-closing prices (defined as the arithmetic average of bid and ask closing prices) at 15:59 from OptionMetrics we also obtain the term-structure of risk-free zero-coupon interest rates and the S&P 500 (continuously compounded) dividend yield. S&P 500 index weekly options are short-term European options written on the S&P 500 index, cash settled and with a fixed time-to-maturity of seven days (five working days). Weekly options are listed under the root ticker symbol, “SPXW” and are commonly included in SPX (traditional) options chains which are AM settled. Officially introduced by the CBOE on October 28, 2005, weekly options start being actively traded one year later. Due to their increasing adoption from the finance community, the CBOE launched in the subsequent years weekly options for all days of the week.⁹ Nowadays SPXW options

⁸The pricing kernel, defined as the ratio of the risk-neutral over the physical measure, converges to 1 as the time to maturity decreases:

$$S_t = e^{r_t T} q \cdot 1 = e^{R_t T} p \cdot 1 \quad (1)$$

where S_t is the current expected payoff of a primitive contingent claim under the risk-neutral measure, q , and under the physical measure, p , respectively. In a complete and arbitrage-free economy the risk-neutral price grows at the current risk-free rate, r_t , whereas the physical price grows at the current risk-adjusted risk-free rate, R_t . It goes by consequence that as the time to maturity T approaches zero, both quantities converge to the same state price density, S_t ; conversely, when T increases their divergence increases at the rate $(R - r)T$.

⁹As a consequence of the high interest from traders, the CBOE now proposes short-term options with even shorter maturities (up to 2 days options) and written on many other underlyings, like equities (American style), ETFs, ETNs, VIX and other indexes like the Dow Jones Industrial Average or the Russell 2000 Index. While increasing, the liquidity of these products might still not high enough to work with non-parametric models. Please see <http://www.cboe.com/products/weekly-options/available-weeklys>.

constitute more than 40% of the average daily volumes of all options traded on the S&P 500 Index, with an average number of daily options traded that went from below 10,000 per day to above 500,000 per day. One possible reason that justifies the great interest of market practitioners, above all market makers, on weekly options is their capability to provide a not expensive (in terms of delta-hedging) and effective tool to hedge short-term market-wide exposures. In particular, weeklys can provide a hedge to short-dated tail risks, without loading on the unneeded volatility risk present in longer-term options (i.e.: monthly or quarterly). Also, weeklys give traders the possibility to have a more targeted exposure to events, e.g.: economic releases, earnings announcements. In this paper we use the Friday SPXW due to their longer history and higher liquidity. Moreover, given that the liquidity starts being stable from 2011, we drop from the dataset all data before 2011 such that the sample size of our analysis spans the 2011-2021 period. As reported by the CBOE, Fridays SPXW are typically listed on Thursday and expire on Friday of the subsequent week¹⁰ and, for this reason, are also called the end-of-the-week (EOW) weekly options.¹¹ From a forecasting perspective weekly options are interesting financial assets that allow for forward-looking estimates, still limiting the assumption that all investors are neutral with respect to the risk. Moreover, from an academic viewpoint, weekly options allow for an interesting change in the option literature. While in fact short-term options have (correctly) always been discarded from most of empirical studies because of not being liquid enough, weekly options are liquid enough to also work with non-parametric approaches.

3 The BIRS model

In this section we describe how we fit the arbitrage-free implied volatility smile that will then be used as the input for the estimation of all option-implied based quantities of the paper. After cleaning the dataset (Section 3.1), the estimation approach is made of two main parts, among them linked one to another. Specifically, we apply a quadratic program for estimating the linear and quadratic terms of piecewise polynomials (Section 3.3), describing the price of call options as a function of their strikes. This method requires a dense pre-estimate. The pre-estimate is obtained through a cubic spline interpolation on the implied volatility - delta space (Section 3.2). The cubic interpolation requires a set of unique strike prices, and the

¹⁰The same applies to all other weeklies, e.g. Monday SPXW options typically expire on Monday, and Wednesday SPXW options typically expire on Wednesday.

¹¹An exception of this rules are if the exchange is closed on a Friday, and/or if the Friday expiration overlaps with the expiration of monthly or quarterly options. In the former case all options are anticipated and so they expire the first business day immediately prior to that Friday, in the latter case the Friday expiration of weekly options is delayed to the next available Friday. In our case, whenever this happens, we substitute the quotation of the weekly option with the corresponding quotation of the monthly or quarterly option that will expire in the week.

quality of the interpolation is determined by the starting point; we achieve this starting point by cleaning the dataset following the option literature (Section 3.1).

3.1 Data cleaning

As it is common in the option literature, we clean the dataset to remove stale and/or irrational option prices, and we only work with the most liquid assets. To avoid stale prices, we remove all options with zero volume and zero open interest. To avoid possibly irrational prices, we discard all options with a zero bid price and take the mid-prices, defined as the average of the bid and ask prices as our option price in order to compute the specific implied volatility of the option. Still to avoid possibly mispriced options, we discard all options with a very high implied volatility ($> 100\%$). Finally, to work with the most liquid assets, we only select out-of-the-money options, and we convert out-of-the-money puts into in-the-money call options via the put-call parity. Table 1 and table 2 provide summary statistics of the data filter and an overview of the available number of options after the cleaning. This cleaned dataset is then used to estimate the volatility smile.

Data filter	Before	After	in percent
Friday expiration	310421	151715	51.13
Positive open interest	151715	104205	31.32
Positive volume	104205	73323	29.64
Positive Bid price	73323	73323	0.00
OTM call	73323	57304	21.85
OTM put	57304	49559	13.52
Starting in 2011	49559	49335	0.45
Implied Volatility higher 100%	49335	49143	0.39
Implied Volatility not calculated	49143	49125	0.04

Table 1: **Summary statistic data cleaning steps:** The table summarizes the data filter presented in Section 3.1

	mean	std	min	25%	50%	75%	max
# Calls available	27.39	14.85	4.00	17.50	22.00	33.00	124.00
# Puts available	79.64	40.26	14.00	51.00	72.00	98.00	211.00
Strike calls	2824.06	940.40	1125.00	2110.00	2810.00	3390.00	4980.00
Strike puts	2577.23	830.36	800.00	1930.00	2515.00	3080.00	4695.00
S&P 500	2813.16	912.42	1123.53	2079.36	2760.17	3348.44	4697.96
Risk-free rate	0.01	0.01	0.00	0.00	0.00	0.01	0.02
Dividend yield	0.02	0.00	0.00	0.02	0.02	0.02	0.03

Table 2: **Summary statistics data used for estimation:** The table provides an overview of the data used to estimate option implied quantiles and expectiles between 2011 and 2021.

3.2 Cubic smoothing spline

Once the raw data are cleaned, we interpolate to obtain a pre-estimate of the volatility surface. For this, we follow and modify the approach of Bliss and Panigirtzoglou (2002) who themselves combine and extend the approaches of Shimko (1993), Malz (1997a), Malz (1997b) and Campa et al. (1998). Specifically, we interpolate the implied volatility curve over the implied volatility/delta space through a cubic smoothing spline, where the smoothing spline is the function f that solves

$$\min_{\theta} \lambda \sum_{i=1}^n w_i (\sigma_i - \hat{\sigma}_i(\delta_i, \theta))^2 + (1 - \lambda) \int_{-\infty}^{\infty} f''(x, \theta)^2 dx, \quad (2)$$

where θ is a set of parameters of the cubic spline (knot points and component polynomial parameters), λ a smoothing parameter, w_i the weights of the spline, σ_i the implied volatility, and $\hat{\sigma}_i(\delta_i, \theta)$ is the fitted implied volatility at δ_i . Differently from Bliss and Panigirtzoglou (2002), we optimally set the smoothing parameter λ at each iteration, and we fix $w_i = 1$.¹² The role of λ is to determine the goodness-of-fit of the fitted spline and its smoothness trade-off, where smoothness is determined by the integrated squared second derivative of the implied volatility function. For $\lambda = 0$, the function is the variational, or natural, cubic spline interpolant. For $\lambda = 1$, the function is the least-squares straight-line fit to the data. Here is how we set these two parameters. First, instead of fixing λ , we search for the optimal balance between having a smooth curve and being close to the given data. Our optimization approach chooses a default value for λ which is usually close to $1/(1 + h^3/6)$, where h is the average spacing of the data sites. Having set λ the calculation of the smoothing spline is the resolution of a linear system,

¹² Bliss and Panigirtzoglou (2002) set λ equal to 0.99 or 0.9999 and $w_i = \nu$, where ν is the Black and Scholes greek letter defined as the first derivative of the option price with respect to the option volatility. While their choice was justified to obtain the best fit possible for the option-implied probability density function, here the focus is on achieving the closest possible empirical fit of the volatility smile.

whose coefficient matrix has the form $\lambda \cdot A + (1 - \lambda) \cdot B$, with the matrices A and B depending on the data sites x . The default value of λ makes $\lambda * \text{Tr}(A) = (1 - \lambda) * \text{Tr}(B)$. Second, the role of w_i is to determine how much weight to give to the i th option's squared fitted implied volatility error. While the central part of the distribution would benefit the most by setting the w_i equal to greek letter vega (as did by Bliss and Panigirtzoglou (2002)), here we set it equal to one, as we are not only interested at-the-money. Finally, the choice of using the option delta as the independent variable comes from Malz (1997a) and guarantees more stability in the interpolation.¹³

Figure 1 depicts the fit of the implied volatility smile estimated using the presented approach. The figure represents 9 days, picked at random, over the implied volatility/strike space in our analysis, where the grey dots are the market prices of the option and the continuous black line the result of the proposed interpolation.

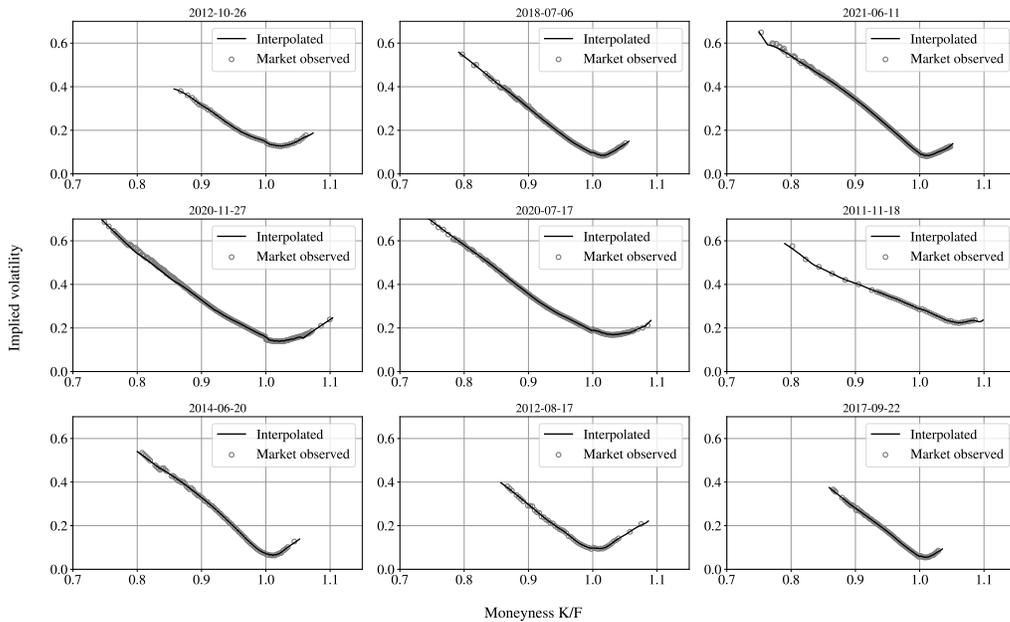


Figure 1: **Implied volatility smile:** obtained solving Equation 2 with $w_i = 1$ and optimally choosing $\lambda \in [0, 1]$.

Having generated a set of fixed-maturity implied volatilities across a grid of deltas, we use an option pricing model, e.g.: the Black and Scholes (1973) pricing model, to convert delta and the implied volatilities into European option prices over the prices/strike space. It is worth noticing that the delta and price conversion through the Black and Scholes model is just a convenient choice that does not impose log-normality, or does not presume that the pricing

¹³Alternatively the implied volatility function can also be expressed as a function of strike or of moneyness.

model correctly prices options.

3.3 Quadratic program to remove the remaining arbitrages

The presented pre-smoother might still be contaminated by arbitrages, above all into the tails of the distribution. For example, given the sample of this paper, 17.2% of the option prices are still contaminated by arbitrages which, in our case, leads to a cumulative density function curve (CDF) that is either negative, greater than 1, or not strictly increasing with respect to the strike (see Figure 2).¹⁴

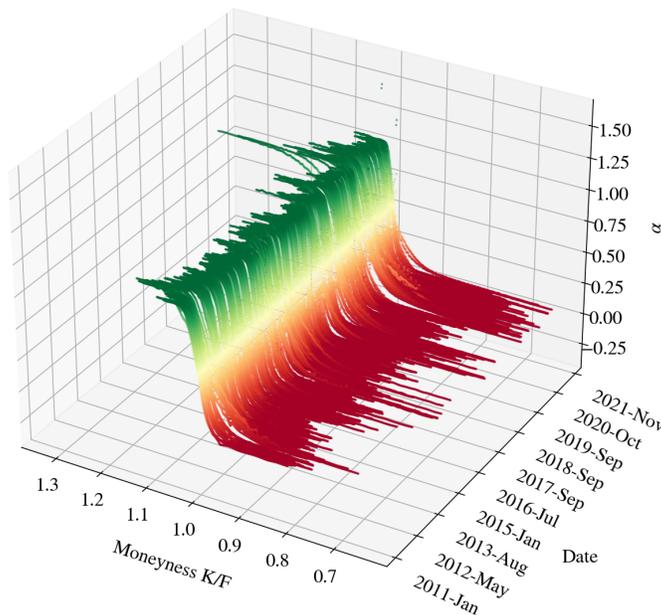


Figure 2: **Option-implied quantiles surface:** The figure shows the time series of the option-implied quantiles estimated as described in Section 5. The figure depicts the time-line on the x-axis, the option-implied quantiles on the y-axis and the strike prices on the z-axis.

To fix any possibly remaining arbitrage on the interpolated smile, we apply the Fengler (2009) quadratic program to the cleaned and interpolated smoother.¹⁵ As a main advantage, the proposed approach does not change anything that does not have arbitrages, and removes

¹⁴How to estimate the quantile-based CDF depicted in Figure 2 and the relative analysis on the importance of having no arbitrages in the tails are both presented in Section 5.

¹⁵As documented in Fengler (2009) and Green and Silverman (1994), the choice of the initial estimator is flexible i.e.: any two-dimensional non-parametric smoother such as a local polynomial estimator or a thin plate spline are valid candidates.

any still possibly contaminated price present in the smile. As detailed in Appendix A, at each point t we solve the following quadratic program:

$$\min_x -y^T x + \frac{1}{2}x^T Bx \quad (3)$$

$$\text{subject to } A^T x = 0 \quad (4)$$

Where y is a vector of size $2n$ containing the n observed implied volatilities and n zeros, and x is a vector of parameters specifying the natural cubic spline. Please note that while the original Fengler (2009) approach is to fit the implied volatility surface, in this paper we do not deal with the dimension of time, fitting only the implied volatility smile. More precisely, focusing on short-term options, we fix $T = 7$ such that the final output of the quadratic program is applied to the call option price/strike plane. The quadratic program in Equation 3 is convex, thus solvable within polynomial time (see the quadratic optimization part of Floudas and Visweswaran (1995)) and naturally casts a cubic smoothing spline in it (Green and Silverman (1994)). The convexity property follows from the strict positive-definiteness of B and guarantees the uniqueness of the solution. Moreover, shape constraining the spline smoother guarantees the optimal rate of convergence in shape-restricted Sobolev classes (Mammen and Thomas-Agnan (2002)). A snapshot of our estimation procedure is in Figure 3, that depicts the S&P 500 options time series of weekly implied-volatility smile for the period 2011-2021.

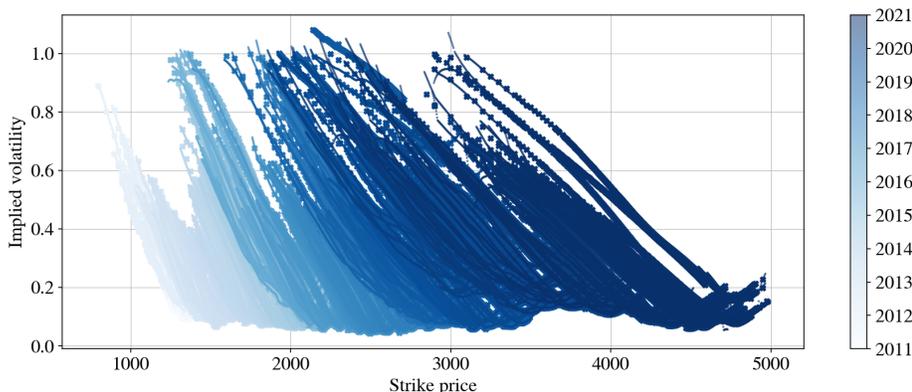


Figure 3: **Implied-volatility smile:** for the S&P 500 index weekly options estimated with the approach presented in Section 3. The figure depicts the strike prices on the x axis, the implied-volatility on the y axis while the older-to-more-recent data of the time series are depicted with a lighter-to-darker color.

As a final way to test the robustness of the estimation approach, and the models ability to accommodate different market scenarios, Figure 4 depicts the implied volatility smile for low

and high volatility days of the time series, respectively. Specifically, having the implied-volatility smile on the y-axis and the strike prices on the x-axis, the top (bottom) panel of Figure 4 represents a day with low (high)¹⁶ volatility, which translates in an implied-volatility smirk (smile).

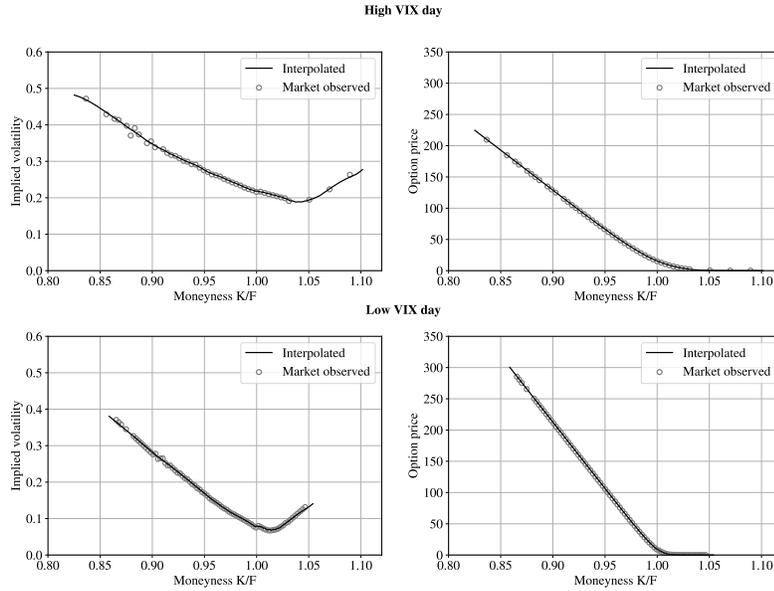


Figure 4: **Implied-volatility smile:** for two single days, both estimated with the approach presented in Section 3. The upper (lower) panel depicts the estimated implied volatility and the obtained option prices for a date in the sample on which the VIX was above (below) its historical level. Both panels of the figure depict the moneyness (K/F) on the x axis.

4 Alternative estimation approaches

In this section we first quickly recall the two estimation approaches that are used to validate and compare our estimation approach with the existing (huge) literature, namely the PCA of Bondarenko (2003) and the FS of Jackwerth (2004).¹⁷ Then, we compare the PCA and the FS models with the BIRS approach proposed in this paper. To do it, we check their capability of

¹⁶The low (high) volatility day is March 21, 2014 (January 15, 2016) where the VIX index closed at 15% (27.02%).

¹⁷We are aware that this literature is incredibly vast and still growing. As such, to select the best alternative approach to benchmark our proposal is a very difficult (if not impossible) task. To do it, in this paper we place our focus on non-parametric models only, and we select the PCA and FS for the reasons presented in the Introduction. Indeed, and not being the goal of the analysis, we do not pretend to provide any exhaustive and/or complete comparison.

recovering the implied volatility and option market prices using observed real data (Section 4.1).

The PCA of Bondarenko (2003) is a non parametric approach that starts from a set of admissible densities and optimally chooses the one that best fits a given cross-section of empirically observed option prices. The set of admissible densities excludes all implausible densities that are economically meaningless (e.g.: discontinuous functions) and is obtained through the convolution of some arbitrary density function and a fixed kernel.

The FS of Jackwerth (2004) is the “fast and stable” approach which is related to the original approach of Jackwerth and Rubinstein (1996) and Jackwerth (2000).¹⁸ The main objective of the “fast and stable” approach is to find a smooth risk-neutral distribution which also explains the observed option prices by first estimating an optimal smooth implied volatilities curve, and then infer the risk-neutral distribution applying Breeden and Litzenberger (1978). Both the PCA and the FS approaches are described in Appendix B and Appendix C, respectively.

4.1 Real data experiment

As a first test we use observed real market data from a cross section of option prices to empirically validate the BIRS approach. Formally, we check its capability of recovering the implied volatility and option market prices using observed real data, and we compare it with the PCA and FS approaches. For each approach, we follow the literature and compute two loss measures used to measure the accuracy of the estimates namely, the root mean squared error (RMSE) and the mean absolute error (MAE):

$$\text{RMSE} = \left[\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2 \right]^{1/2}$$

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|$$

where $i = 1, \dots, N$ denotes the total number of observations, y_i the observed values, and \hat{y}_i the interpolated ones. In clockwise order, the four panels of Figure 5 depict the entire time series of the RMSE price, RMSE implied volatilities, MAE price, and MAE implied volatilities for the three approaches presented, respectively. Placing on the x-axis the timeline of the analysis, and on the y-axis the value of the loss function, the figures plot with a black continuous line our approach (BIRS), in gray the FS approach and in light gray the PCA approach.

¹⁸While there is a substantial change from the approach in Jackwerth and Rubinstein (1996) to the subsequent version in Jackwerth (2004), the only “difference” between Jackwerth (2004) and Jackwerth (2020) is a rewriting of the scaling factor in the fit-smoothness objective function used for the optimization.

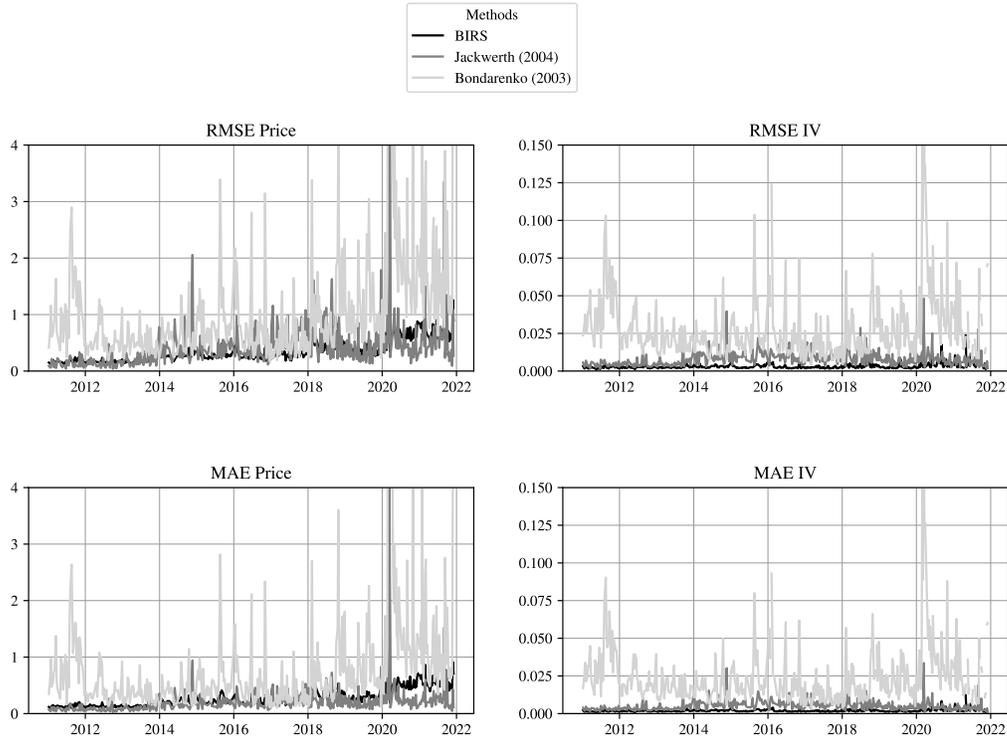


Figure 5: **Time series RMSE and MAE:** Time series plots of the calculated RMSE and MAE for the different methods to estimate risk-neutral densities (risk-neutral distributions). The left hand side of the graphs display the mean dollar price deviations on a daily base within the sample period. The right hand side of the graphs display the respective deviations in terms of implied volatility.

From the figure it clearly emerges that for the entire time series in consideration, and for both prices and implied volatilities, the BIRS approach produces much lower loss measures, once compared with the FS and PCA approaches. While Figure 5 depicts graphically the entire time series of RMSE and MAE, Table 3 summarizes the above findings, reporting different summary statistics of the RMSE and MAE metrics, respectively.

	BIRS		Jackwerth (2004)		Bondarenko (2003)	
	Prices	IV	Prices	IV	Prices	IV
RMSE						
mean	0.2909	0.0021	0.2091	0.0053	0.8365	0.0240
std	0.1935	0.0013	0.2398	0.0035	1.3751	0.0198
min	0.0922	0.0009	0.0247	0.0011	0.0997	0.0034
25%	0.1694	0.0015	0.1031	0.0031	0.3105	0.0124
50%	0.2336	0.0018	0.1661	0.0044	0.5175	0.0183
75%	0.3680	0.0022	0.2581	0.0066	0.9170	0.0291
max	2.7562	0.0193	4.1889	0.0335	19.8959	0.2164
MAE						
mean	0.3769	0.0032	0.4337	0.0081	1.1571	0.0286
std	0.3123	0.0025	0.4426	0.0050	1.5793	0.0226
min	0.1210	0.0012	0.0392	0.0017	0.1614	0.0040
25%	0.2209	0.0022	0.2090	0.0046	0.4811	0.0151
50%	0.2962	0.0026	0.3385	0.0069	0.7731	0.0224
75%	0.4611	0.0035	0.5560	0.0103	1.3052	0.0346
max	5.3922	0.0368	6.7896	0.0482	22.8134	0.2342

Table 3: **RMSE and MAE:** The table summarizes the root mean squared error (RMSE) and the mean absolute error (MAE). The RMSE and MAE are calculated on a daily basis. Both measures are calculated based on observed and interpolated option prices (Prices) and the corresponding Black-Scholes implied volatility (IV).

Overall the BIRS approach clearly outperforms both the PCA and the FS approaches in terms of IV. Results holds for the mean, and at all confidence levels. In terms of prices, the BIRS approach is always the best performing model, but for some values of the RMSE (mean and some percentiles). Focusing on the RMSE,¹⁹ the BIRS approach has a maximum value of 2.7562 (0.0193) for the prices (IV), which is lower than 4.1889 (0.0335) and above all the 19.8959 (0.2164) of the FS and PCA, respectively. For the other values, the BIRS approach is comparable to the FS approach and always better than the PCA approach. Importantly, and for both prices and IVs, the BIRS approach always displays the highest stability in estimation, proxied by the lowest standard deviation. As a consequence of this stability, the max RMSE and MAE of the BIRS approach are lower than both the FS and PCA. These results (and specifically the higher stability of the BIRS approach) are the consequence of both the interpolation and the quadratic problem presented in the previous sections, which lead to a lower presence of possibly mispriced assets in estimation. This is especially true into the tails of the distribution, where lies most of the option prices noise.

¹⁹Which by construction has values that are always greater or equal to the MAE.

5 Option-implied quantiles and expectiles

In this section we first quickly recall the concepts of quantiles and expectiles (Section 5.1). Then we explain how to infer option-implied quantiles (Section 5.2) and option-implied expectiles (Section 5.3) from traded option prices. While quantiles are well-known statistical quantities that have been heavily used for different applications (see, among others, the survey of Koenker et al. (2017)), expectiles are less known in finance, perhaps due to their lower interpretability. Despite being less intuitive than quantiles, expectiles have interesting and possibly even superior statistical properties than quantiles, once used as risk management tool (Bellini and Di Bernardino (2017)).

5.1 Quantiles and Expectiles

Given a random variable X on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$, a quantile q_α determines the value of X such that the probability of the variable being less than or equal to that value equals α . It follows that for a probability density function $f(x)$ the quantile q_α splits the distribution in two parts, which integrals are of size α and $1 - \alpha$. Formally, if X is equipped with a continuous and strictly monotonic distribution function $F_X(x) := P(X \leq x)$, then q_α is the unique solution of the equation:

$$F_X(q_\alpha) = \alpha \tag{5}$$

Any quantile functions satisfies (see e.g. Föllmer and Schied (2016)):

$$q_\alpha^-(X) \leq q_\alpha(X) \leq q_\alpha^+(X), \text{ for each } \alpha \in (0, 1),$$

where $q_\alpha^-(X)$ and $q_\alpha^+(X)$ are the left and right quantiles, respectively defined by:

$$\begin{aligned} q_\alpha^-(X) &= \inf\{t \in \mathbb{R} \mid F_X(t) \leq \alpha\} \\ q_\alpha^+(X) &= \sup\{t \in \mathbb{R} \mid F_X(t) \geq \alpha\}. \end{aligned}$$

Equivalently, the left and right quantiles of the random variable X are defined as the minimizer of the asymmetric *linear* loss function:

$$[q_\alpha^-(X), q_\alpha^+(X)] = \operatorname{argmin}_{x \in \mathbb{R}} \mathbb{E}[\alpha(X - x)_+ + (1 - \alpha)(X - x)_-], \quad \alpha \in (0, 1).$$

In a financial context quantiles are (among others) often used for Value-at-Risk (Jorion (2007)) or quantile regressions (Koenker (2005)).

First introduced by Newey and Powell (1987), expectiles are a one-parameter family of coherent risk measure defined as the minimizers of an asymmetric *quadratic* loss function. Formally, the expectile e_θ of a random variable $X \in L^2(\Omega, \mathcal{F}, P)$ is defined as:

$$e_\theta(X) := \arg \min_{x \in \mathbb{R}} \mathbb{E} [\theta(X - x)_+^2 + (1 - \theta)(X - x)_-^2], \quad \theta \in (0, 1),$$

where $(\cdot)_+ = \max(\cdot, 0)$, $(\cdot)_- = \min(\cdot, 0)$. Expectiles combine the concept of “expectation” and “quantiles” and are the asymmetric generalization of the mean, being that for $\theta = 1/2$, we have $e_\theta(X) = \mathbb{E}(X)$. Expectiles can also be conveniently defined for any $X \in L^1(\Omega, \mathcal{F}, P)$ as the unique solution of the first order condition:

$$\theta \mathbb{E}[(X - e_\theta(X))_+] = (1 - \theta) \mathbb{E}[(X - e_\theta(X))_-]. \quad (6)$$

Statistically, quantiles and expectiles share many similar properties but differ substantially in one aspect. While quantiles determine the value of X such that the *probability* of the variable being less than or equal to that value equals a given level α , expectiles are linked to the properties of the *expectation* of the random variable X , conditional on X being into the tail of the distribution. Moreover, notice that for a given distribution function F^{20} , the values of θ and α are related by the following formula (see e.g. Yao and Tong (1996)):

$$\theta = \frac{\alpha q_\alpha - \int_{-\infty}^{q_\alpha} x dF(x)}{\mathbb{E}(X) - 2 \int_{-\infty}^{q_\alpha} x dF(x) - (1 - 2\alpha)q_\alpha} \quad (7)$$

For example, if $X \sim U[-a, a]$, then $q_\alpha = 2\alpha a - a$, $\theta = \alpha^2 / (2\alpha^2 - 2\alpha + 1)$ and for $\alpha = 1\%, 5\%, 10\%, 25\%, 50\%$ the values of θ are 0.01%, 0.27%, 1.2%, 10%, 50%. In a financial context expectiles are related to but still different from, what is generally referred in risk management as the Conditional Value at Risk (CVaR); as such contain information about what to expect when the random variable attains a value beyond the quantile (VaR) (Taylor (2008)).

5.2 Option-Implied Quantiles

From the fundamental theorems of asset pricing, given a dynamically complete and arbitrage-free finite economy, the value of a European put option corresponds to the present value of the expected payoff under the risk-neutral measure (see, e.g. Ross (1976)):

$$P_{t,T} = e^{-r_t(T-t)} \int_0^K (K - S_T, 0)^+ f(S_T) dS_T \quad (8)$$

²⁰It is worth noticing that this approach does not apply in this paper, being that both quantiles and expectiles will be inferred non-parametrically from options.

where r_t is the continuously compounded interest rate, $(K - S_T, 0)^+$ is the option payoff and $f(S_T)dS_T$ the risk-neutral density function. Given this economy, and from the seminal papers of Breeden and Litzenberger (1978) and Banz and Miller (1978), the cumulative distribution function can be inferred from European option prices by taking the first order derivatives of the option price with respect to the strike price:

$$\frac{\partial P_{t,T}}{\partial K} = \frac{\partial[e^{-r_t(T-t)} \int_0^K (K - S_T, 0)^+ f(S_T)d(S_T)]}{\partial K} \quad (9)$$

$$= e^{-r_t(T-t)} \int_0^K f(S_T)dS_T \quad (10)$$

$$= e^{-r_t(T-t)} \alpha_{t,T} \quad (11)$$

With no loss of generality, the same holds using European call options:

$$\frac{\partial C_{t,T}(K)}{\partial K} = \frac{\partial[e^{-r_t(T-t)} \int_K^\infty (S_T - K, 0)^+ f(S_T)d(S_T)]}{\partial K} \quad (12)$$

$$= -e^{-r_t(T-t)} \int_K^\infty f(S_T)dS_T \quad (13)$$

$$= -e^{-r_t(T-t)} (1 - \alpha_{t,T}) \quad (14)$$

Since the results above are attained by integrating continuous functions, practical application with discrete strike distances between price observations demand an approximate solution. The literature offers many approaches, both parametric and non-parametric, to solve this problem. Parametric approaches rely on a pricing model and enjoy known partial derivatives. For example, following the Black and Scholes pricing model the current risk-neutral probability of the underlying being in-the-money at expiration are $N(-d_2) = \mathbb{P}(S_T < K)$ and $N(d_2) = \mathbb{P}(S_T > K)$ for put and call option, respectively. Although quick and elegant, the definition of a pricing function inherently requires an assumption of the underlying price process, and is for this reason neglected in this paper.²¹ Alternatively, among the non-parametric approaches, finite differences-based models enable results without defining a-priori a price-process. Following Barone-Adesi and Elliott (2007), we propose a finite differences-based scheme that is free of any first-order error caused by the changes of the implied volatility across strike prices, and that eliminates the first-order error that arises from the Taylor expansion of the derivative. Formally, for three consecutive European put option prices $P_{t,T}(K_{i+1}) > P_{t,T}(K_i) > P_{t,T}(K_{i-1})$ with strike prices

²¹For example, through the proposed Black and Scholes pricing model we would inherently impose log-normality to our final results.

$K_{i+1} > K_i > K_{i-1}$ the finite-difference equivalent of Equation 11 is defined as:

$$\hat{\alpha}_{t,T}^{\text{Put}}(K_i) = e^{-r_t(T-t)} \left[\lambda \left(\frac{P_{t,T}(K_{i+1}) - P_{t,T}(K_i)}{K_{i+1} - K_i} \right) + (1 - \lambda) \left(\frac{P_{t,T}(K_i) - P_{t,T}(K_{i-1})}{K_i - K_{i-1}} \right) \right] \quad (15)$$

with $\lambda = \frac{(K_{i+1} - K_{i-1})}{(K_i - K_{i-1})}$ such that $\lambda = 0.5$ in presence of equidistant strike prices $K_{i+1} - K_i = K_i - K_{i-1}$.²² Similarly to the VIX Index, and to maximize the liquidity of put and call options, option implied quantiles are computed joining Equations (11) and (14). More precisely, we will use European put options to compute the portion of $\hat{\alpha}_{t,T}$ whose strike prices are lower than the current forward value, and European call options to compute the portion of $\hat{\alpha}_{t,T}$ whose strike prices are higher than the current forward price Figure 6 depicts the time series of the S&P 500 weekly option-implied quantile curves for the period 2011-2021.

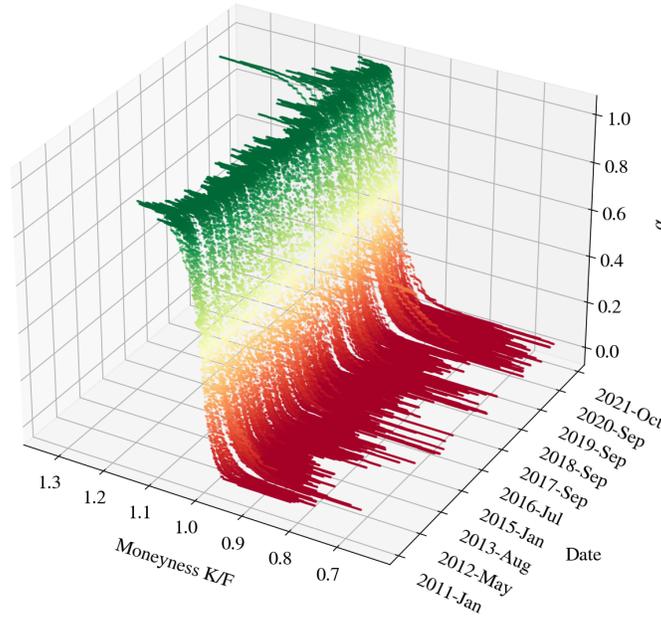


Figure 6: **Option-implied quantiles surface:** The figure shows the time series of the option-implied quantiles estimated as described in Section 5. The figure depicts the time-line on the x-axis, the option-implied quantiles/expectiles on the y-axis and the moneyness defined as K/F on the z-axis.

The option-implied curve is in green (red) whenever its value is greater (smaller) than 0.5. The option-implied curve shifts accordingly with the evolution of the S&P 500. For example,

²²With no loss of generality, the same scheme can be also applied to European Call options to estimate Equation (14).

the overall uptrend of the market is visible, with some abrupt fall (like in Summer 2018 and the end/beginning part of 2015/2016). Moreover, comparing Figure 6 with Figure 2 the role of the quadratic program implemented is now recognizable. In particular, while Figure 6 depicts a smooth curve with no arbitrage violations, Figure 2 has many irregularities present in different areas of the curve.

Dropping one dimension, Figure 7 shows the time series of the option-implied quantiles with the price evolution of the S&P 500 superimposed. The figure depicts in blue the time series of the S&P 500 index and the entire daily curve of the option-implied quantiles with all values above (below) the underlying price in green (red).

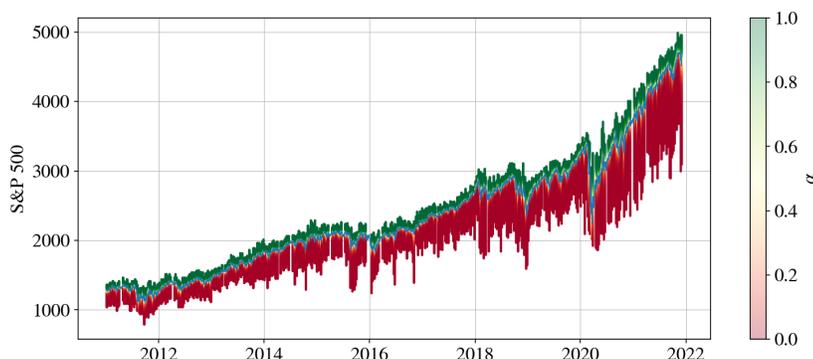


Figure 7: Option-implied quantiles: The figure shows the time-series of the option-implied quantile for the weekly options written on the S&P 500 index for the period 2011-2021 with the S&P 500 index superimposed (in blue). The figure depicts in green (red) all option-implied quantile values that are above (below) the daily price of the S&P 500 Index.

From the figure it emerges how the time series of option-implied quantiles closely mimics the underlying, both in calm and in turbulent periods. In particular, during low volatility periods the dispersion of the daily option-implied quantile values is much smaller compared to more volatile days, characterized by stronger up or down spikes whenever the market goes up or down. As mentioned in Section 1, unless in presence of stationary and ergodic market, this ability of following the market so closely would not be possible using historical data and/or parametric estimation approaches.

5.3 Option-Implied Expectiles

Departing from Equation (6) and still assuming a dynamically complete and arbitrage-free finite economy, the $\theta_{t,T}$ -expectile is the unique strike price \bar{K} that solves equation:

$$\begin{aligned}\theta_{t,T}\mathbb{E}[(X - e_\theta(X))_+] &= (1 - \theta_{t,T})\mathbb{E}[(X - e_\theta(X))_-] \\ \theta_{t,T}C_{t,T}(\bar{K}) &= (1 - \theta_{t,T})P_{t,T}(\bar{K}).\end{aligned}\tag{16}$$

where $\mathbb{E}[(X - e_\theta(X))_-]$ and $\mathbb{E}[(X - e_\theta(X))_+]$ are the payoff of a European call and put option with underlying X and strike price $e_\theta(X)$, respectively. To empirically extract the inverse option-implied expectile $\hat{\theta}_{t,T}(K_i)$ at a generic level K_i , we reorder the above equation such that:

$$\hat{\theta}_{t,T}(K_i) := \frac{P_{t,T}(K_i)}{P_{t,T}(K_i) + C_{t,T}(K_i)}.\tag{17}$$

Whenever necessary, and to exploit the higher liquidity of European put (call) options for strike prices lower (greater) than the underlying, we resort to the put-call parity to infer possible missing prices for a given strike price. More precisely, if for a given day and strike price only the price of a European put or call option is available, the other price can be determined through the put-call parity so that the option-implied expectile is estimated from solely European call options:

$$\theta_{t,T}(K_i) = \frac{C_{t,T}(K_i) - S_t e^{-q_t(T-t)} + K_i e^{-r_t(T-t)}}{2C_{t,T}(K_i) - S_t e^{-q_t(T-t)} + K_i e^{-r_t(T-t)}}\tag{18}$$

or solely from European put options:

$$\theta_{t,T}(K_i) = \frac{P_{t,T}(K_i)}{2P_{t,T}(K_i) + S_t e^{-q_t(T-t)} - K_i e^{-r_t(T-t)}}\tag{19}$$

where q_t is the continuously compounded dividend rate. As for the option-implied quantiles, also this method is fully non-parametric and data driven. Figure 8 depicts the time series of the S&P 500 weekly option-implied expectile curves estimated with Equation 17.

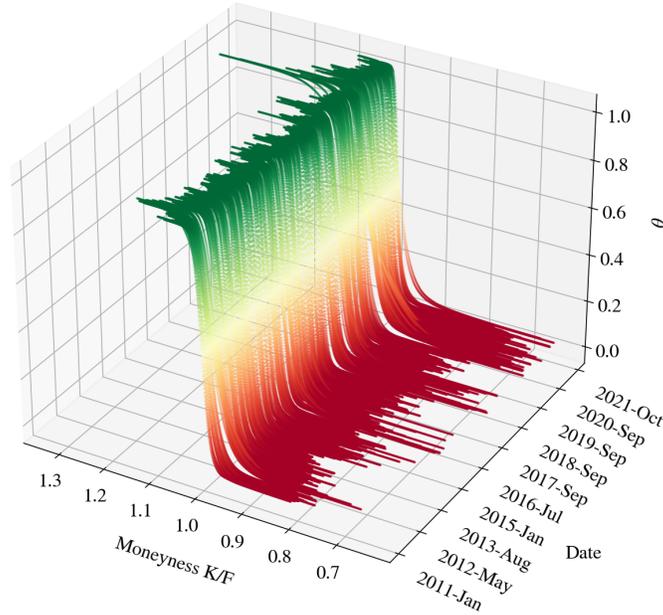


Figure 8: **Option-implied expectiles surface:** The figure shows the time series of the option-implied expectiles estimated as described in Section 5. The figure depicts the time-line on the x-axis, the option-implied quantiles/expectiles on the y-axis and the moneyness defined a K/F on the z-axis.

As expected, the option-implied expectile curve is both similar to the option-implied quantile curve and less noisy, above all into the tails of the distribution. This is a direct consequence of the two different estimation methods, the one used to estimate option-implied expectiles being less numerically intensive than the one used for quantiles. Figure 9 shows the two-dimensional time series of the option-implied expectiles with the price evolution of the S&P 500 superimposed, and confirms both the ability of the time series to track the underlying and the lower amount of noise, once compared with the option-implied quantile time series.

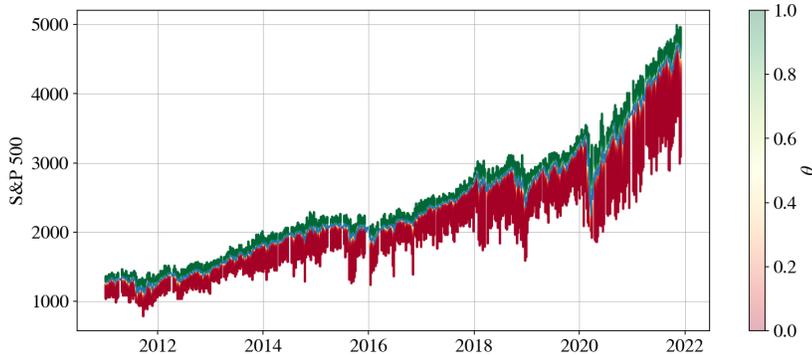


Figure 9: **Option-implied expectiles:** The figure shows the time-series of the option-implied expectile for the weekly options written on the S&P 500 index for the period 2011-2021 with the S&P 500 index superimposed (in blue). The figure depicts in green (red) all option-implied expectile values that are above (below) the daily price of the S&P 500 Index.

5.4 Monte Carlo experiment

Having defined option-implied quantiles and expectiles, we now propose a further validation test to compare the BIRS approach with the PCA and FS models presented in Section 4. While in Section 4.1 we compared the BIRS approach with the PCA and FS ones using real market data, in this Section we compare it using simulated data through a Monte Carlo experiment. The rationale behind the experiment is to check the capability of the models to recover a quantity knowing the parametric form of the quantity itself. For the experiment, we follow Bondarenko (2003) and test the validity of the proposed pricing models during days with high and low volatility. A high (low) volatility day in our sample corresponds to a VIX value above (below) its historical average. To infer the risk-neutral distribution, we depart from the observed weekly option prices on the specific date and we calibrate a mixture of three log-normal distributions.

	Mean	Std	Weight
LN1	7.14	0.04	0.33
LN2	7.16	0.02	0.26
LN3	7.17	0.02	0.41

Table 4: **High VIX day log-normal parameter estimates:** The table summarizes the results of the fitted mixture of log-normal model with three log-normals. LN1, LN2 and LN3 represent the specific log-normal distribution used. The parameter represent a day in the sample when the VIX was above its historical average.

In order to estimate the parameter of the mixtures of log-normal distributions, we rely on the Nelder-Mead algorithm. To initialize the algorithm, we set the starting value for the mean, $\mu_{\text{initial guess}}$, equal to the natural logarithm of the current spot S&P 500 level on the specific date. The initial values for the standard deviation are set to $\sigma_{\text{initial guess}} = [0.03, 0.05, 0.3]$. Based on the estimated parameter, we generate an artificial set of European call and put options, corresponding to the available strike prices on the specific date. To include possible option market microstructures, such as a the bid/ask spreads, we test the performance of each method under two noise specifications. First, as a starting point, we follow Rompolis (2010) and add to the generated option prices a uniformly distributed noise term between -0.025 and 0.025. Secondly, to account for the different liquidity that an option can have depending on its moneyness (which results in higher bid/ask spreads), we follow Ait-Sahalia and Duarte (2003) and add to the generated option prices a noisy term that varies with the option moneyness equal to $0.5 \times (\text{Price}_{\text{ask}} - \text{Price}_{\text{bid}}) \times \text{Liquidity factor}$ where $\text{Liquidity factor} = 1 + \left(\frac{2}{0.2}\right) * \left|\left(\frac{K}{F_t} - 1\right)\right|$ and F_t represents today's forward value. Finally, we evaluate each method based on the root mean integrated squared error (RMISE), and split this error measure up into two parts, the root integrated squared bias (RISB), referred to as 'bias', and the root integrated variance (RIV), referred to as 'variability', which allows us to examine the stability and accuracy of each method separately (Bondarenko (2003)):

$$\begin{aligned}
\text{RMISE} &= \frac{1}{\sqrt{\int f(x)^2 dx}} \sqrt{E \left[\int (\widehat{f(x)} - f(x))^2 dx \right]} \\
\text{RISB} &= \frac{1}{\sqrt{\int f(x)^2 dx}} \sqrt{\int (E[\widehat{f(x)}] - f(x))^2 dx} \\
\text{RIV} &= \frac{1}{\sqrt{\int f(x)^2 dx}} \sqrt{\int (E[(\widehat{f(x)} - E[\widehat{f(x)}])^2]) dx}
\end{aligned} \tag{20}$$

while

$$\text{RMISE}^2 = \text{RISB}^2 + \text{RIV}^2$$

Table 5 summarizes the results for a random day in which the VIX was above its historical average.

	BIRS			Jackwerth (2004)		Bondarenko (2003)	
	LN 3	QAlpha	EAlpha	QAlpha	EAlpha	QAlpha	EAlpha
No noise							
RMISE	1.1341	0.9806	1.0069	1.1365	1.1590	1.2746	1.2930
Bias	0.8019	0.8061	0.8078	0.8205	0.8225	0.8334	0.8339
Variability	0.8019	0.5585	0.6010	0.7864	0.8165	0.9645	0.9882
Uniform noise							
RMISE	1.1341	0.9807	1.0069	1.1363	1.1589	1.2746	1.2930
Bias	0.8019	0.8060	0.8078	0.8205	0.8225	0.8334	0.8339
Variability	0.8019	0.5586	0.6011	0.7860	0.8164	0.9644	0.9882
Moneyness noise							
RMISE	1.1341	0.9801	1.0062	1.1357	1.1582	1.2739	1.2923
Bias	0.8019	0.8061	0.8079	0.8205	0.8226	0.8334	0.8340
Variability	0.8019	0.5576	0.5997	0.7851	0.8154	0.9635	0.9871

Table 5: **High VIX:** The table summarizes the root mean integrated squared error (RMISE). ‘Bias’ refers to the root integrated squared bias and ‘Variability’ to the root integrated variance as calculated in equation 20. The ‘Uniform noise’ specification refers to the case when we added uniformly distributed noise terms between $[-0.025, 0.025]$ to the artificial generated option prices as in Rompolis (2010). The results for the category ‘Moneyness noise’ specify the case where the noise term for deeper out-of-the money options increases as in Ait-Sahalia and Duarte (2003).

Examining the RMISE across the different methods, we infer that the BIRS approach has a lower RMISE, once compared with the PCA and the FS, respectively, both the uniform and the moneyness noises. It is noteworthy to see that increasing the complexity of the noise specification does not necessarily lead to a higher RMISE for the BIRS.²³

5.5 Constructing option-implied quantile and expectile based measures

As a follow-up to Figures 6, 8, 7 and 9, we further investigate if the estimated option-implied quantiles or option-implied expectiles have some forecasting power. For the regression analysis, the vectors of option-implied quantiles and expectiles for each weekly option chain are compressed to descriptive statistics that highlight the attributes of the price densities. We therefore use the ranges between the option-implied quantiles or option-implied expectiles of certain orders as a robust descriptive statistics of the variability of the estimated distributions.

Specifically for $\alpha, \theta > 1/2$, we define a robust and conditional Quantile Dispersion (QD)

²³Low VIX days confirm the superiority of the model, once compared with the PCA and FS approaches. Results are presented in the On-line appendix.

indicators and Expectile Dispersion (ED) indicators as:

$$\text{QD}_{t,T} = \hat{q}_{t,T}^\alpha - \hat{q}_{t,T}^{1-\alpha} \quad (21)$$

$$\text{ED}_{t,T} = \hat{e}_{t,T}^\theta - \hat{e}_{t,T}^{1-\theta} \quad (22)$$

and we set the confidence intervals α and θ equal to 75%, 90%, and 95%, them being the most commonly investigated.²⁴ To bypass the arbitrary choice of how to set the confidence intervals, we propose the integrated versions of the above indicators, which summarize the entire information for all α and θ :

$$\text{QD}_{t,T}^{\int_{0.5}^1} = \int_{0.5}^1 (\hat{q}_{t,T}^\alpha - \hat{q}_{t,T}^{1-\alpha}) d\alpha \quad (23)$$

$$\text{ED}_{t,T}^{\int_{0.5}^1} = \int_{0.5}^1 (\hat{e}_{t,T}^\theta - \hat{e}_{t,T}^{1-\theta}) d\theta. \quad (24)$$

Being defined as the difference between two quantiles (or expectiles), these dispersion indicators are alternative measures of market variability.

Following Hinkley (1975) and Ghysels et al. (2016), we then propose a robust alternative to the third moment. Formally, we define the robust option-implied coefficient of quantile asymmetry (QA) and expectile asymmetry (EA) as:

$$\text{QA}_{t,T} = \frac{[q_{t,T}^\alpha - q_{t,T}^{0.5}] - [q_{t,T}^{0.5} - q_{t,T}^{1-\alpha}]}{[q_{t,T}^\alpha - q_{t,T}^{1-\alpha}]} \quad \text{for } \alpha \in (0.5, 1) \quad (25)$$

$$\text{EA}_{t,T} = \frac{[e_{t,T}^\theta - e_{t,T}^{0.5}] - [e_{t,T}^{0.5} - e_{t,T}^{1-\theta}]}{[e_{t,T}^\theta - e_{t,T}^{1-\theta}]} \quad \text{for } \theta \in (0.5, 1) \quad (26)$$

Again, we evaluate the QA and EA, setting the confidence levels at 0.75, 0.9, and 0.95 where, for a confidence level equal to 0.75, we retrieve the Bowley (1920) statistic. Once more to avoid the arbitrary choice of setting the confidence sets, we propose the integrated version of the QA and EA:

$$\text{QA}_{t,T}^{\int_{0.5}^1} = \frac{\int_{0.5}^1 [(q_{t,T}^\alpha - q_{t,T}^{0.5}) - (q_{t,T}^{0.5} - q_{t,T}^{1-\alpha})] d\alpha}{\int_{0.5}^1 [q_{t,T}^\alpha - q_{t,T}^{1-\alpha}] d\alpha} \quad \text{for } \alpha \in (0.5, 1) \quad (27)$$

$$\text{EA}_{t,T}^{\int_{0.5}^1} = \frac{\int_{0.5}^1 [(e_{t,T}^\theta - e_{t,T}^{0.5}) - (e_{t,T}^{0.5} - e_{t,T}^{1-\theta})] d\theta}{\int_{0.5}^1 [e_{t,T}^\theta - e_{t,T}^{1-\theta}] d\theta} \quad \text{for } \theta \in (0.5, 1) \quad (28)$$

By construction, the above measures are bounded between -1 and 1, revolve around the median and indicate negative (positive) asymmetry whenever they are below (above) 0, while being

²⁴Being strike prices not in a continuum, each range is estimated using the strike prices closest to each bound.

different than the third moment of returns. As emphasized in Ghysels et al. (2016), these asymmetry indicators are robust to outliers and can be computed at various time horizons. Differently from Ghysels et al. (2016), our measures are short-term option-implied coefficients of asymmetry, instead of long-term stock-based coefficients of asymmetry. As such, our measures are based upon forward-looking information, rather than constructing the forecast on historical return data.

Finally, as a proxy for the fourth moment, we follow Ruppert (1987) and Ammann and Feser (2019) and propose a robust and conditional measure of quantile flatness (QF) and expectile flatness (EF):

$$\text{QF}_{t,T} = \frac{[q_{t,T}^{\alpha} - q_{t,T}^{1-\alpha}]}{[q_{t,T}^{\omega} - q_{t,T}^{1-\omega}]} \quad \text{for } \alpha \in (0.75, 1) \text{ and } \omega = 0.7 \quad (29)$$

$$\text{EF}_{t,T} = \frac{[e_{t,T}^{\theta} - e_{t,T}^{1-\theta}]}{[e_{t,T}^{\omega} - e_{t,T}^{1-\omega}]} \quad \text{for } \theta \in (0.75, 1) \text{ and } \omega = 0.7 \quad (30)$$

where we select as starting point the value of the tail $\omega = 0.7$ and again select 0.75, 0.9 and 0.95 as the confidence interval. Again, and to avoid an arbitrary choice for the confidence, we also calculate an alternative more generic version that considers the entire distribution:

$$\text{QF}_{t,T}^{\int_{0.75}^1} = \frac{\int_{0.75}^1 [q_{t,T}^{\alpha} - q_{t,T}^{1-\alpha}] d\alpha}{[q_{t,T}^{\omega} - q_{t,T}^{1-\omega}]} \quad \text{for } \alpha \in (0.75, 1) \text{ and } \omega = 0.7 \quad (31)$$

$$\text{EF}_{t,T}^{\int_{0.75}^1} = \frac{\int_{0.75}^1 [e_{t,T}^{\theta} - e_{t,T}^{1-\theta}] d\theta}{[e_{t,T}^{\omega} - e_{t,T}^{1-\omega}]} \quad \text{for } \theta \in (0.75, 1) \text{ and } \omega = 0.7 \quad (32)$$

The summary statistics of all calculated option-implied robust measures are collected in Table 6 while the BIRS time series are depicted in Figure 10.²⁵

²⁵The time series of the PCA and FS approach are in the Online Appendix.

	mean	std	min	25%	50%	75%	max
QD(0.75)	0.02	0.02	0.01	0.01	0.02	0.03	0.15
QD(0.9)	0.05	0.03	0.01	0.03	0.04	0.06	0.25
QD(0.95)	0.06	0.04	0.02	0.04	0.06	0.07	0.31
QD(Integral)	0.02	0.01	0.01	0.01	0.01	0.02	0.08
QA(0.75)	-0.10	0.28	-0.91	-0.29	-0.11	0.05	0.81
QA(0.9)	-0.21	0.15	-0.66	-0.31	-0.22	-0.11	0.34
QA(0.95)	-0.26	0.12	-0.69	-0.35	-0.27	-0.18	0.12
QA(Integral)	-0.19	0.17	-0.61	-0.31	-0.21	-0.08	0.27
QF(0.75)	1.37	0.31	1.00	1.18	1.31	1.47	3.21
QF(0.9)	2.83	0.71	1.69	2.44	2.66	2.95	9.21
QF(0.95)	3.83	1.02	2.29	3.25	3.54	4.05	11.91
QF(Integral)	0.67	0.17	0.42	0.57	0.62	0.70	1.91
ED(0.75)	0.02	0.01	0.01	0.01	0.01	0.02	0.10
ED(0.9)	0.04	0.02	0.01	0.02	0.03	0.04	0.20
ED(0.95)	0.05	0.03	0.02	0.03	0.04	0.05	0.24
ED(Integral)	0.01	0.01	0.00	0.01	0.01	0.01	0.06
EA(0.75)	-0.09	0.05	-0.29	-0.12	-0.09	-0.05	0.09
EA(0.9)	-0.16	0.05	-0.38	-0.20	-0.16	-0.13	-0.01
EA(0.95)	-0.21	0.06	-0.42	-0.25	-0.21	-0.17	-0.06
EA(Integral)	-0.18	0.05	-0.36	-0.21	-0.17	-0.14	-0.03
EF(0.75)	1.30	0.04	1.17	1.28	1.30	1.32	1.42
EF(0.9)	2.66	0.09	2.42	2.61	2.65	2.72	3.01
EF(0.95)	3.67	0.18	3.09	3.54	3.64	3.80	4.38
EF(Integral)	0.65	0.04	0.53	0.62	0.64	0.67	0.80

Table 6: **Summary statistics BIRS**: The table provides summary statistics of the option implied measures.

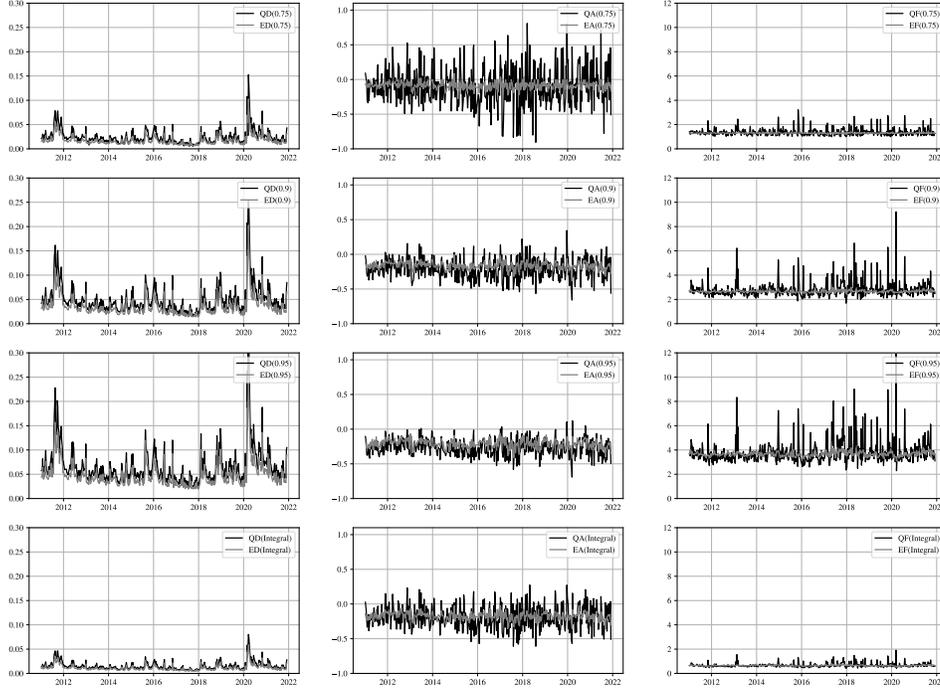


Figure 10: **Time series of the BIRS-based robust option-implied measures:** The figure illustrates the time series of the four robust option-implied measures. The figure depicts vertically the robust option-implied measures of conditional dispersion (left), asymmetry (center), and flatness (right) and horizontally considering four different confidence intervals, 0.75, 0.9, 0.95 and Integral, respectively.

From both the tables and the figures it emerges the higher numerical stability of the expectile-based robust measures and of those quantities that are closer to the center of the distribution. As summarized by the min, max and mean values as well as by all intervals and coherent with equation (7), quantile-based indicators are always larger than expectile-based ones. Once more, the higher stability of the option-implied expectiles once compared with the option-implied quantiles, reflect the different estimation approaches.²⁶ The higher stability of the indicators that are closer to the center of the distribution, reflects the higher presence of option market data at the main part of the distribution. Moving away from the center reflects the scarcity of

²⁶As mentioned in Section 5 option-implied quantiles are in fact estimated with three contiguous prices, while option-implied expectiles only need two points in estimation.

the data in the tails and the resulting lower stability of the indicator.

After the estimation of the robust option-implied quantities, our subsequent goal of our investigation is to analyze if these robust indicators (QD, ED, $QD^{J_{0.5}^1}$, $ED^{J_{0.5}^1}$, QA, EA, $QA^{J_{0.5}^1}$, $EA^{J_{0.5}^1}$ and finally QF, EF and $QF^{J_{0.5}^1}$ and $EF^{J_{0.5}^1}$) have some forecasting power. More generally, we investigate the economic gains from exploiting option-implied dispersion, asymmetry and flatness in the distribution of returns at short, medium and long time period. It is worth noticing that to promote stability and stationarity, the option-implied quantile curve and the option-implied expectile curves are computed as their implied deviations from the forward price: $R_{t,t+\tau}^{CS} = \log\left(\frac{S_{t,t+\tau}^{CS}}{S_f}\right)$, where $S_{t,t+\tau}^{CS}$ is the index price at time $t + 1$ at a confidence set (CS) determined by α or τ for quantiles and expectiles, and S_f is the forward price of the index. For example, to estimate the QD at a given level α we consider, not the difference in dollar value of the index at the level α and $1 - \alpha$, but its return difference at level α and $1 - \alpha$.

To benchmark the performances of our option-implied indicators, the same analysis is also repeated with the same indicators, but this time inferred from historical returns. More precisely, we estimate again the aforementioned indicators using realized quantiles inferred from a 30 days rolling window of past historical returns, in order to be used as a benchmark model in computing out-of-sample R^2 s (Campbell and Thompson (2007)).²⁷ The null hypothesis of this study is that the forward-looking information of risk-neutral quantities derived from options can not predict future returns. To test this hypothesis, QD, ED, QA, EA, EF, and QF, are used as regressors X for predicting $r_{t,t+\tau}^A$, the τ -days ahead log-return adjusted for the risk-free interest rate. The analysis is first performed in-sample through a simple linear regression fitted to each feature, X , with $r_{t,t+\tau}^A$ as dependent variable:

$$\hat{r}_{t,t+\tau}^A = \alpha + \beta X_t + \varepsilon_{t+\tau} \quad (33)$$

For each factor we evaluate its in-sample performance (R^2), slope β for a short ($\tau = 7$ days), medium ($\tau = 60$ days) and long ($\tau = 180$ days) time horizon. Moreover, following (Campbell and Thompson (2007)), we compare the obtained results with a similar exercise using historical data to compute the measures of asymmetry:

$$R_{OS}^2 = 1 - \frac{\sum_{t=1}^T (r_t - \hat{r}_{(t-w,t)})^2}{\sum_{t=1}^T (r_t - \tilde{r}_{(t-w,t)})^2} \quad (34)$$

²⁷The same analysis has also been performed for the expectile equivalent of the same indicators (ED, $ED^{J_{0.5}^1}$, ECA and $ECA^{J_{0.5}^1}$). Due to the similarity in the final results and to save space, results are not reported but are available upon request to the authors.

where $\hat{r}_{(t-w,t)}$ is estimated with Equation 33 using the forward-looking measures of asymmetry between $t-w$ and time t , where w is set at 36 weeks, and \tilde{r} is estimated similarly using historical data.²⁸ The comparison of these two quantities gives rise to a benchmark-adjusted R_{OS}^2 . For positive values of R_{OS}^2 , the proposed model, $\hat{r}_{(t-w,t)}$, is performing better than the benchmark using historical data, \tilde{r} , and vice versa for negative values.

Finally, we test if our option-implied indicators have some forecasting power in predicting future volatility, skewness, and kurtosis, at short, medium and long-term horizon through a univariate regression:

$$\widehat{\sigma^2}_{t,t+\tau} = \alpha + \beta X_t + \varepsilon_{t,t+\tau} \quad (35)$$

$$\widehat{\mathbb{S}}[r_{t,t+k}] = \alpha + \beta X_t + \varepsilon_{t,t+\tau} \quad (36)$$

$$\widehat{\mathbb{K}}[r_{t,t+k}] = \alpha + \beta X_t + \varepsilon_{t,t+\tau} \quad (37)$$

To account for potential time-series correlation caused by the overlapping observations in computing compounded moments, we estimate Newey-West standard errors (Newey and West (1987)) with lags equal to the number of weeks between observations. Finally we estimate the same specifications but using a quantile-regression instead of a standard linear model in order to estimate the median instead of the conditional average.

6 Forecasting exercise

In this section we discuss our main findings related to the predictive power of the robust option-implied measures presented in Section 5.5 with respect to risk premium, volatility, and higher realized moments.²⁹ Tables 7 and 8 present the results of regressing the option-implied measures on realized returns, volatility, skewness and kurtosis. We provide 54 specifications for estimating each one of the realized moments, varying the probability α in the quantiles and expectiles, as well as the horizon in the prediction. We vary the parameter α between 0.75 and 0.95 to study how the predictive power changes as we approach the tails of the distribution, and we vary the forecast horizon between one week and 26 weeks to study the short-, medium- and long-run forecast power. Finally, since the realized moments of the distribution are likely to contain extreme values in short horizons, in Tables 9 and 10 we repeat the analysis using quantile

²⁸To lighten the notation for the out-of-sample analysis, we drop the time orientation $_{t,t+\tau}$. Nevertheless, the analysis is again performed for the short, medium and long time horizon.

²⁹For reason of space and to facilitate the reading of the analysis we only present the findings related to the approach presented in this paper, and compare it with historical-based benchmarks. Nevertheless, it is worth noticing that the *entire* empirical analysis has been performed also for the PCA and FS approaches presented in Section 4 and are available upon request from the authors. Results have been omitted because the PCA and FS approaches produces among them overall very similar results, but always inferior to the ones generated with the approach presented in this paper, thus confirming Sections 4.1 and 5.4.

regressions to estimate the median value of the conditional distribution. We find that the option-implied indicators of dispersion, QD and ED, have the highest predictive power on the realized moments. Moreover, these measures are more robust across the different specifications. We concentrate our discussion of the results mostly on these two measures during the remaining of the section. The other robust measures, QA, EA, QF and EF show small or zero forecasting power, thus confirming the absence of time-series forecasting power of the implied high moments, also for weekly options.

Table 7 shows that, QD and ED, the robust option-implied measures of dispersion, predict expected returns for horizons of 9 and 26 weeks, and that QD exhibits short term predictability for an α of 0.75. Consistent with the literature on return predictability (e.g. Cochrane (2008)), the magnitude of the point estimates, t-statistics, and in-sample R^2 s increase with the time horizon. Point estimates are statistically significant at the 1% level for horizons of 9 and 26 weeks. Point estimates and in-sample R^2 s decrease as we approach the tails of the distribution, which suggests that the predictive power of the measures of asymmetry is reduced when we take into account observations in the tails of the distribution.

Our results are also economically significant and, as expected, improve at short-time. Specifically, for short horizons, taking the $\alpha = 0.75$ specification as benchmark, a one standard deviation (0.02) increase in the QD measure, is related to an increase of 0.5% (4.6%) in average market returns over one (26) week(s). Using the ED measure, we find that the same increase of a one standard deviation (0.01) for an $\alpha = 0.75$ is related to an increase of (3.43%) in expected returns over horizons of one (26) week(s). The specifications using 26 weeks as horizons provide positive out-of-sample R^2 s when using historical measures of asymmetry as a benchmark varying between 12 and 16 percent. We find that the robust measures of asymmetry and flatness do not predict expected returns or volatility in our sample.

Table 7: Linear Predictive Model

$$y_{t,t+k} = a + bx_t + \varepsilon_{t,t+k}$$

$y_{t,t+k}$		$\mathbb{E}[r_{t,t+k}]$										$\sigma[r_{t,t+k}]$							
k		1w			9w			26w			1w			9w			26w		
x_t	α	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2
QD	0.75	0.25** (2.2)	2.63	-2.65	1.1*** (5.06)	9.79	-2.91	2.29*** (7.04)	22.57	14.91	0.33*** (9.41)	43.72	36.52	0.19*** (7.85)	20.63	14.22	0.07*** (2.6)	3.82	1.81
	0.9	0.03 (0.36)	-0.03	-5.36	0.52*** (4.37)	8.21	-4.77	1.2*** (7.03)	23.33	15.35	0.2*** (7.79)	56.9	50.71	0.11*** (7.9)	25.91	19.43	0.04*** (3.26)	5.76	3.12
	0.95	0.02 (0.27)	-0.11	0.12	0.4*** (4.18)	7.67	-5.96	0.94*** (6.75)	23.2	12.3	0.15*** (7.38)	56.15	47.38	0.09*** (7.19)	26.15	19.3	0.03*** (3.2)	5.81	3.09
ED	0.75	0.13 (0.53)	0.16	-10.03	1.53*** (4.71)	8.65	-3.92	3.43*** (7.32)	23.39	16.07	0.56*** (7.88)	56.69	52.09	0.31*** (8.16)	25.07	19.15	0.12*** (3.22)	5.54	3.13
	0.9	0.06 (0.5)	0.11	-7.34	0.76*** (4.58)	8.43	-4.74	1.72*** (7.11)	23.24	16.14	0.28*** (7.92)	56.89	50.05	0.15*** (8.18)	25.2	17.53	0.06*** (3.26)	5.67	2.91
	0.95	0.04 (0.44)	0.04	-1.52	0.57*** (4.4)	8.09	-5.56	1.3*** (6.62)	22.7	12.72	0.21*** (7.71)	56.54	47.36	0.12*** (7.63)	25.21	17.15	0.05*** (3.16)	5.69	2.65
QA	0.75	-0.0 (-0.12)	-0.22	-0.34	0.01 (1.26)	0.03	-6.07	0.01 (0.81)	-0.1	2.03	0.0 (0.25)	-0.2	-3.19	0.0 (0.02)	-0.22	-1.19	0.0 (1.15)	0.04	0.42
	0.9	0.01 (0.94)	0.35	-0.35	0.03 (1.48)	0.42	-1.93	0.0 (0.14)	-0.21	0.08	-0.01 (-1.52)	1.63	-7.63	-0.0 (-1.1)	1.1	-5.73	-0.0 (-0.41)	-0.16	-0.09
	0.95	0.02 (1.08)	0.75	-2.18	0.04 (1.38)	0.61	-1.2	0.02 (0.46)	-0.16	-1.95	-0.01* (-1.71)	2.57	-4.05	-0.01 (-1.32)	1.73	-1.59	-0.0 (-1.51)	0.53	-10.75
EA	0.75	0.05 (1.46)	0.92	-0.61	0.07 (1.38)	0.17	-5.91	-0.01 (-0.08)	-0.22	2.56	-0.02* (-1.71)	1.71	-2.29	-0.01 (-1.0)	0.59	-1.89	-0.01 (-0.98)	0.09	0.37
	0.9	0.04 (1.1)	0.58	0.34	0.08 (1.32)	0.4	-2.5	0.06 (0.59)	-0.07	0.08	-0.03** (-2.03)	3.47	-8.13	-0.02 (-1.06)	1.27	-8.71	-0.01 (-1.55)	1.52	-1.21
	0.95	0.04 (1.22)	0.68	-1.17	0.09 (1.55)	0.72	-2.18	0.12 (1.43)	0.68	-1.2	-0.02* (-1.81)	1.97	-5.37	-0.01 (-0.96)	0.63	-4.63	-0.01* (-1.68)	1.59	-12.76
QF	0.75	-0.0 (-0.02)	-0.22	0.02	-0.0 (-0.82)	-0.14	-0.21	-0.01** (-1.99)	0.09	0.31	-0.0 (-1.21)	0.07	1.3	-0.0 (-0.72)	-0.18	0.31	0.0 (0.35)	-0.21	-0.26
	0.9	-0.01 (-1.49)	3.19	-0.97	-0.0 (-1.48)	0.1	-0.9	-0.0 (-1.17)	0.0	1.61	0.0 (0.61)	0.46	-3.42	0.0 (0.57)	-0.06	0.34	0.0 (1.47)	0.48	1.72
	0.95	-0.0 (-1.49)	2.82	2.43	-0.0* (-1.84)	0.39	-3.41	-0.01** (-2.14)	0.31	1.07	0.0 (0.21)	-0.15	-13.87	0.0 (0.14)	-0.21	0.87	0.0 (1.13)	0.34	0.63
EF	0.75	-0.03 (-1.04)	-0.0	0.6	0.03 (0.39)	-0.19	-0.47	0.03 (0.32)	-0.2	0.03	0.01 (0.91)	-0.14	0.35	-0.0 (-0.29)	-0.2	0.05	-0.0 (-0.13)	-0.22	0.33
	0.9	-0.03** (-2.16)	0.8	1.43	-0.05 (-1.57)	0.44	-1.28	-0.1 (-1.51)	1.29	1.69	-0.02*** (-3.59)	3.06	7.68	-0.01** (-1.99)	2.16	8.21	-0.0 (-0.22)	-0.2	1.8
	0.95	-0.01** (-2.04)	0.9	3.39	-0.04* (-1.92)	1.39	-4.31	-0.08*** (-1.99)	4.34	0.85	-0.01*** (-4.01)	8.47	2.97	-0.01** (-2.25)	5.06	12.55	-0.0 (-0.72)	0.03	1.99

Table 8: Linear Predictive Model

$$y_{t,t+k} = a + bx_t + \varepsilon_{t,t+k}$$

$y_{t,t+k}$		$\mathbb{S}[r_{t,t+k}]$									$\mathbb{K}[r_{t,t+k}]$								
k		1w			9w			26w			1w			9w			26w		
x_t	α	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2
QD	0.75	-0.3 (-0.23)	-0.21	-0.21	4.98 (1.54)	1.32	-2.64	7.23* (1.68)	4.35	6.27	-2.38** (-2.36)	0.69	1.93	-6.51 (-0.78)	0.24	-6.6	-11.13 (-0.7)	0.39	0.61
	0.9	-0.35 (-0.53)	-0.18	0.15	1.89 (1.13)	0.62	1.31	3.11 (1.33)	2.98	4.48	-1.26** (-2.5)	0.75	1.16	-2.98 (-0.65)	0.14	-5.77	-3.01 (-0.33)	-0.05	0.93
	0.95	-0.34 (-0.64)	-0.16	0.39	1.43 (1.07)	0.56	0.36	2.44 (1.31)	2.94	2.28	-0.95** (-2.33)	0.67	1.05	-2.5 (-0.68)	0.19	-6.75	-2.82 (-0.39)	0.02	-0.33
ED	0.75	-1.18 (-0.62)	-0.16	-0.09	6.02 (1.3)	0.83	-2.84	9.57 (1.45)	3.5	5.89	-3.57** (-2.53)	0.74	1.72	-9.35 (-0.73)	0.22	-6.65	-10.08 (-0.38)	0.01	0.45
	0.9	-0.67 (-0.69)	-0.15	0.21	2.91 (1.25)	0.74	1.36	4.76 (1.45)	3.39	5.08	-1.82** (-2.55)	0.75	1.26	-4.95 (-0.75)	0.26	-5.93	-5.51 (-0.41)	0.05	0.99
	0.95	-0.6 (-0.8)	-0.12	0.46	2.15 (1.2)	0.68	0.45	3.61 (1.44)	3.34	2.64	-1.38** (-2.49)	0.75	1.27	-4.09 (-0.8)	0.35	-6.79	-4.89 (-0.48)	0.15	-0.23
QA	0.75	0.04 (0.47)	-0.17	-0.74	-0.02 (-0.2)	-0.21	0.13	0.14** (2.2)	0.4	1.27	-0.08 (-1.19)	0.1	-2.05	-0.04 (-0.17)	-0.21	-0.59	-0.26 (-0.92)	-0.11	-3.5
	0.9	0.14 (0.83)	-0.06	-0.69	0.09 (0.43)	-0.17	2.21	0.3 (1.6)	0.49	1.62	-0.14 (-1.21)	0.07	-1.0	-0.01 (-0.01)	-0.22	-2.32	-1.25 (-1.54)	0.49	2.88
	0.95	0.14 (0.69)	-0.11	2.0	0.15 (0.52)	-0.13	2.5	0.29 (1.15)	0.26	-1.36	-0.14 (-0.99)	-0.0	-0.17	0.36 (0.6)	-0.13	-9.06	-1.41 (-1.5)	0.41	-1.26
EA	0.75	0.42 (0.88)	-0.07	-1.06	0.44 (0.58)	-0.1	0.45	1.32* (1.81)	1.32	1.32	-0.48 (-1.35)	0.15	-0.42	0.65 (0.42)	-0.17	-0.43	-5.19 (-1.56)	1.12	-1.64
	0.9	0.82* (1.8)	0.41	-0.76	1.0 (1.16)	0.46	2.47	1.3 (1.43)	1.39	1.05	-0.37 (-1.06)	0.02	-0.53	1.61 (0.92)	0.09	-1.58	-5.41 (-1.23)	1.36	3.67
	0.95	0.83** (2.12)	0.69	2.24	0.98 (1.29)	0.7	1.96	0.92 (1.11)	0.93	-2.15	-0.25 (-0.83)	-0.07	0.05	1.94 (1.3)	0.41	-7.77	-3.35 (-0.85)	0.64	-1.03
QF	0.75	0.0 (0.02)	-0.22	-0.83	-0.01 (-0.09)	-0.22	-0.17	-0.08 (-1.35)	0.03	1.1	-0.05 (-0.83)	-0.06	-0.49	-0.45** (-2.09)	0.68	1.49	-0.21 (-1.0)	-0.13	2.52
	0.9	-0.02 (-0.75)	-0.14	0.09	-0.05 (-1.19)	0.17	-0.95	-0.09*** (-3.06)	1.42	7.97	-0.03 (-1.12)	0.06	0.95	-0.15 (-1.44)	0.33	-1.2	0.22 (1.58)	0.3	-8.3
	0.95	-0.01 (-0.55)	-0.17	0.33	-0.03 (-0.95)	0.11	-2.75	-0.07*** (-2.75)	1.67	8.44	-0.01 (-0.8)	-0.09	1.03	-0.1 (-1.39)	0.29	0.13	0.17 (1.59)	0.39	-8.53
EF	0.75	0.46 (0.67)	-0.12	-1.19	0.54 (0.76)	-0.12	0.23	1.17*** (3.48)	0.45	2.02	-0.48 (-0.98)	-0.01	-0.68	-4.32** (-2.46)	0.9	1.19	-4.39*** (-3.33)	0.31	3.05
	0.9	-0.0 (-0.01)	-0.22	0.33	-0.28 (-0.62)	-0.05	-0.07	-0.48 (-1.04)	0.48	8.79	-0.15 (-0.81)	-0.09	0.29	-1.78* (-1.75)	0.97	-0.18	-0.66 (-0.34)	-0.15	-6.16
	0.95	-0.08 (-0.57)	-0.15	0.22	-0.23 (-0.86)	0.24	-2.5	-0.4 (-1.45)	1.84	10.35	-0.01 (-0.07)	-0.22	0.37	-0.77 (-1.49)	0.72	0.64	-0.22 (-0.2)	-0.18	-5.91

Inference using linear regression models can be misleading if residuals are not spherical. We repeat our analysis, using quantile regressions in which the right hand side variables provide variation in the estimated conditional median of the distribution, rather than the conditional mean. Table 9 shows the results of estimating the median realized returns for the same set of horizons discussed above. We confirm that the measures of asymmetry QD and ED predict median returns for all horizons up to 26 weeks, and that all results are once more both economically and statistically significant at the 1% level. Consistent with the results obtained in predicting average returns, point estimates, t-statistics, and in sample R^2 s increase with the time horizon. However point estimates are significant at the 1% level across all specifications. This suggests that our measures of dispersion have a better performance in capturing the variation of the median of these moments both for short and long horizons. The same one standard deviation increase (0.02) for an $\alpha = 0.75$ for QD is related to an increase in the median return of 0.66% (one week), 1.8% (nine weeks) and 3.2% (26 weeks). For ED, this increase (0.01) is related to an increase of 0.46% (one week), 1.82 (nine weeks) and 3.08% (26 weeks) respectively. Moreover, we find that using a quantile specification out-of-sample R^2 s are positive for horizons of 26 weeks suggesting that the robust option-implied measures of asymmetry QD and ED better predict average and median market returns.

Table 9: Quantile Predictive Model

$$y_{t,t+k} = a + bx_t + \varepsilon_{t,t+k}$$

$y_{t,t+k}$		$E[r_{t,t+k}]$									$\sigma[r_{t,t+k}]$								
k		1w			9w			26w			1w			9w			26w		
x_t	α	$b/t(b)$	R^2	R_{Oos}^2	$b/t(b)$	R^2	R_{Oos}^2	$b/t(b)$	R^2	R_{Oos}^2	$b/t(b)$	R^2	R_{Oos}^2	$b/t(b)$	R^2	R_{Oos}^2	$b/t(b)$	R^2	R_{Oos}^2
QD	0.75	0.33*** (5.47)	2.14	-2.57	1.26*** (8.83)	7.43	-3.81	2.19*** (9.98)	10.37	13.53	0.31*** (27.58)	28.37	39.21	0.18*** (18.71)	16.29	21.04	0.08*** (9.34)	5.02	3.16
	0.9	0.15*** (4.93)	1.37	-4.55	0.62*** (8.58)	6.66	-4.61	1.08*** (9.54)	10.8	17.59	0.17*** (28.24)	30.12	50.18	0.1*** (21.13)	19.33	24.76	0.05*** (11.72)	6.79	5.3
	0.95	0.11*** (4.49)	1.2	-3.78	0.47*** (8.29)	6.41	-5.49	0.82*** (9.18)	10.78	13.13	0.13*** (28.36)	29.55	45.17	0.08*** (21.03)	18.96	23.13	0.04*** (10.04)	6.52	5.62
ED	0.75	0.46*** (5.09)	1.34	-5.38	1.82*** (8.92)	6.83	-3.81	3.08*** (9.62)	10.99	14.73	0.5*** (30.23)	30.47	54.24	0.29*** (21.76)	18.77	24.39	0.13*** (10.27)	6.66	4.68
	0.9	0.21*** (4.71)	1.23	-4.05	0.91*** (8.83)	6.59	-4.14	1.58*** (9.77)	10.92	18.82	0.24*** (29.0)	30.37	52.64	0.15*** (21.76)	18.98	22.51	0.07*** (10.32)	6.86	5.87
	0.95	0.13*** (3.76)	1.06	-2.04	0.69*** (8.77)	6.27	-5.04	1.18*** (9.49)	10.66	13.34	0.18*** (28.46)	30.02	47.08	0.11*** (20.55)	19.2	21.16	0.05*** (11.31)	6.95	6.14
QA	0.75	0.0 (0.25)	0.01	0.19	-0.0 (-0.0)	0.0	-4.42	0.02 (1.36)	0.25	3.62	0.0 (0.23)	0.04	-0.06	0.0** (2.23)	0.44	-0.54	0.0 (0.08)	0.0	-1.13
	0.9	0.0 (0.07)	0.0	0.16	-0.0 (-0.15)	0.01	-0.81	0.0 (0.03)	0.0	-2.18	-0.0 (-0.75)	0.12	-4.57	0.0 (0.26)	0.0	-4.52	-0.0 (-1.18)	0.15	4.74
	0.95	0.0 (0.55)	0.03	-0.1	-0.01 (-0.63)	0.06	-0.89	0.0 (0.08)	0.0	-4.46	-0.0** (-2.26)	0.72	-1.28	0.0 (0.35)	0.0	0.22	-0.0* (-1.85)	0.39	0.28
EA	0.75	0.03* (1.67)	0.33	1.06	-0.02 (-0.35)	0.02	-4.42	0.03 (0.38)	0.01	4.51	-0.01** (-2.56)	0.81	1.09	0.01 (1.42)	0.12	-2.16	-0.0 (-1.31)	0.17	-1.65
	0.9	0.02 (1.24)	0.26	1.93	-0.02 (-0.45)	0.03	-0.8	0.0 (0.07)	0.0	-0.96	-0.01*** (-2.81)	1.03	-3.34	0.0 (0.24)	0.01	-5.11	-0.01** (-2.1)	0.7	4.08
	0.95	0.02 (1.35)	0.29	1.82	0.01 (0.17)	0.01	-0.34	0.03 (0.6)	0.04	-3.47	-0.01** (-2.59)	0.96	-2.26	-0.0 (-0.35)	0.0	-0.27	-0.01*** (-2.62)	0.86	0.84
QF	0.75	0.0 (1.32)	0.19	-1.34	0.0 (0.04)	0.0	-0.57	-0.0 (-0.45)	0.07	-2.11	-0.0 (-0.98)	0.11	0.41	0.0 (0.54)	0.0	-0.22	0.0 (0.44)	0.03	0.05
	0.9	-0.0 (-1.34)	0.12	-1.74	0.0 (0.34)	0.01	-3.26	-0.01** (-2.07)	0.19	-1.83	-0.0* (-1.65)	0.36	0.85	-0.0 (-0.12)	0.01	0.09	0.0* (1.85)	0.27	0.52
	0.95	-0.0 (-1.44)	0.18	1.24	0.0 (0.02)	0.0	-5.33	-0.01* (-1.88)	0.25	4.23	-0.0*** (-3.23)	0.82	-1.38	-0.0* (-1.96)	0.22	4.08	0.0 (1.15)	0.18	0.04
EF	0.75	-0.0 (-0.12)	0.0	-0.89	0.06 (0.98)	0.12	-1.03	0.06 (0.65)	0.06	-2.5	0.0 (0.52)	0.08	-0.06	0.0 (0.57)	0.04	-0.08	-0.0 (-0.52)	0.02	-0.19
	0.9	-0.02 (-1.55)	0.39	0.17	-0.03 (-1.05)	0.11	-2.32	-0.06 (-1.6)	0.4	-1.87	-0.01*** (-5.62)	3.76	5.08	-0.0** (-2.06)	0.61	2.84	-0.0 (-0.41)	0.01	1.5
	0.95	-0.01** (-2.35)	1.09	-0.4	-0.03*** (-2.41)	0.84	-3.94	-0.06*** (-3.24)	1.31	4.48	-0.01*** (-6.64)	6.31	7.46	-0.0*** (-2.75)	1.08	8.28	-0.0 (-1.07)	0.07	1.14

Table 10: Quantile Predictive Model

$$y_{t,t+k} = a + bx_t + \varepsilon_{t,t+k}$$

$y_{t,t+k}$		$\mathbb{S}[r_{t,t+k}]$									$\mathbb{K}[r_{t,t+k}]$								
k		1w			9w			26w			1w			9w			26w		
x_t	α	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2	$b/t(b)$	R^2	R_{oos}^2
QD	0.75	0.3 (0.14)	0.0	-1.02	5.28** (2.43)	0.84	-5.89	7.7*** (4.28)	2.53	4.98	-2.17 (-1.3)	0.49	1.54	-3.93 (-0.94)	0.16	-5.53	-20.53*** (-3.18)	1.2	-6.48
	0.9	-0.08 (-0.08)	0.01	0.16	1.41 (1.27)	0.4	-0.42	3.36*** (3.59)	1.82	-1.74	-1.19 (-1.37)	0.4	-0.57	-2.02 (-0.94)	0.14	-7.88	-11.01*** (-3.3)	0.89	-2.17
	0.95	-0.21 (-0.25)	0.03	0.77	1.02 (1.16)	0.39	-0.69	2.79*** (3.78)	1.8	-2.53	-0.92 (-1.36)	0.42	-0.98	-1.67 (-0.99)	0.19	-10.41	-8.87*** (-3.4)	1.15	-0.8
ED	0.75	-0.25 (-0.08)	0.01	-0.88	4.83 (1.52)	0.49	-4.49	11.42*** (4.32)	2.29	2.67	-3.45 (-1.41)	0.49	0.95	-6.3 (-1.03)	0.15	-6.1	-33.22*** (-3.53)	1.18	-6.49
	0.9	-0.41 (-0.26)	0.02	0.41	1.89 (1.18)	0.42	0.08	5.79*** (4.33)	2.2	-0.15	-1.75 (-1.41)	0.52	-0.11	-3.2 (-1.03)	0.17	-8.33	-16.53*** (-3.5)	1.32	-2.53
	0.95	-0.4 (-0.33)	0.03	0.46	1.3 (1.06)	0.37	-0.51	4.39*** (4.3)	2.14	-0.77	-1.33 (-1.41)	0.53	0.19	-2.39 (-1.01)	0.2	-10.91	-13.3*** (-3.7)	1.54	-0.88
QA	0.75	0.1 (0.89)	0.15	-1.64	-0.11 (-0.96)	0.1	3.91	0.12 (1.16)	0.36	2.9	-0.17* (-1.95)	0.66	-2.73	0.33 (1.48)	0.22	3.32	-0.93** (-2.56)	0.85	-5.24
	0.9	0.27 (1.25)	0.39	-0.85	0.2 (0.87)	0.07	2.65	0.15 (0.78)	0.14	1.28	-0.31* (-1.89)	0.92	0.19	0.51 (1.2)	0.23	-1.07	-1.52** (-2.17)	0.85	0.58
	0.95	0.19 (0.74)	0.18	2.6	0.31 (1.16)	0.18	-0.64	0.09 (0.4)	0.02	4.84	-0.33 (-1.64)	0.85	-0.45	0.73 (1.43)	0.42	-4.17	-1.43* (-1.67)	0.6	-0.03
EA	0.75	0.35 (0.53)	0.13	-0.37	0.95 (1.4)	0.31	2.42	0.82 (1.38)	0.35	-0.01	-0.81 (-1.58)	0.54	0.61	1.66 (1.29)	0.33	2.29	-4.63** (-2.16)	0.59	-3.65
	0.9	0.71 (1.12)	0.36	-1.31	1.26* (1.96)	1.13	1.9	0.5 (0.87)	0.18	-0.52	-0.61 (-1.23)	0.45	2.63	1.66 (1.34)	0.47	-0.89	-3.32 (-1.61)	0.37	0.11
	0.95	0.86 (1.63)	0.62	3.63	1.33** (2.49)	1.37	-4.41	0.18 (0.38)	0.05	2.35	-0.51 (-1.21)	0.29	1.74	1.91* (1.83)	0.65	-3.67	-1.84 (-1.06)	0.14	-0.51
QF	0.75	0.01 (0.12)	0.0	-1.06	-0.06 (-0.61)	0.08	3.26	-0.15 (-1.61)	0.52	0.72	-0.04 (-0.55)	0.06	0.15	-0.39** (-1.98)	0.51	2.89	-0.0 (-0.0)	0.0	1.45
	0.9	-0.02 (-0.38)	0.07	-0.2	-0.05 (-1.11)	0.48	2.1	-0.1** (-2.58)	1.68	22.6	-0.01 (-0.2)	0.0	2.76	-0.18** (-2.1)	0.63	3.91	0.44*** (3.07)	0.71	-7.27
	0.95	-0.02 (-0.52)	0.1	1.23	-0.04 (-1.2)	0.42	-1.03	-0.07** (-2.57)	1.77	16.87	-0.01 (-0.26)	0.01	1.09	-0.13** (-2.22)	0.7	1.36	0.28*** (2.82)	0.82	-7.65
EF	0.75	0.23 (0.26)	0.01	-0.61	-0.29 (-0.31)	0.02	2.6	1.12 (1.39)	0.18	3.65	-0.71 (-1.01)	0.37	-0.17	-2.33 (-1.33)	0.3	2.64	-7.01** (-2.44)	0.96	1.33
	0.9	-0.4 (-1.14)	0.36	-1.28	-0.74** (-2.03)	0.82	1.12	-0.41 (-1.3)	0.34	24.67	-0.29 (-1.04)	0.43	0.81	-1.52** (-2.21)	0.75	3.34	-2.41** (-2.11)	0.51	-6.24
	0.95	-0.24 (-1.4)	0.54	0.15	-0.32* (-1.79)	0.9	-2.21	-0.28* (-1.79)	0.75	21.81	-0.13 (-0.98)	0.33	-0.14	-0.5 (-1.49)	0.57	1.32	-1.0* (-1.75)	0.3	-6.94

In the second half of Tables 7 and 10 we present the results of using linear and quantile univariate regressions in predicting realized volatility. We find that QD and ED, the robust measures of conditional dispersion, positively predict realized volatility for horizons up to 26 weeks, and that these results are economically and statistically significant. Contrary to how point estimates behave in predicting expected returns, we find that the point estimates, t-statistics and in sample R^2 s in predicting volatility decrease with the time horizon. This result is consistent with the existence of volatility clusters (Engle (1982); Bollerslev (1986)). We find that a one standard deviation increase in QD for an $\alpha = 0.75$ is associated with an increase in volatility of 0.62% for a one week horizon, 0.4% for a nine week horizon, and 0.16% for a 26 week horizon. Similarly for the ED measure, these estimations are 0.5 %, 0.29% and 0.13% equivalently. We find that point estimates decrease as we approach the tails of the distribution in the specifications, however the in-sample R^2 remains in average unchanged. Finally, we find that for horizons up to 9 weeks, the out-of-sample R^2 is large and positive, varying from 13% up to 52%. This evidence suggests that forward-looking measures of asymmetry perform better in predicting realized volatility for short horizons.

We find that for longer horizons, the forward-looking measures perform better at estimating the median of the realized volatility distribution as suggested by the positive and large out-of-sample R^2 s in the last column compared to the baseline linear specification. In fact the out of sample R^2 s in the quantile regressions that aim to predict volatility at longer horizons are two or three times larger than in the linear specification. This is due to the fact that the realized estimation of volatility over large horizons have extreme realizations in the data.

We proceed to study the predictive power of the robust forward-looking measures in explaining higher moments. Tables 8 and 10 present the results of predicting the realized skewness and kurtosis of the market portfolio for short and medium run horizons. We find that for horizons of 26 weeks the QD and ED measures positively predict skewness. These results are statistically significant at the 5% level for a nine week horizon and 1% level for a 26 week horizon. Point estimates, t-statistics and in sample R^2 s increase with the horizon used in the regressions. For the benchmark case, and setting $\alpha = 0.75$, we find that a one standard deviation increase in QD is associated with an increase in the expected realized skewness of 10.6 (nine weeks), and 15.2 (26 weeks). Equivalently, for the ED measure a one standard deviation increase is associated with an increase in the expected realized skewness of 22.8 for 26 weeks.

We find some smaller evidence that the measure of asymmetry EF negatively predicts the average of the realized kurtosis as we approach the tails of the distribution for longer horizons. This relation is significant at the 1% level. There is some evidence that the quantities QA and EA negatively predict the median of the realized kurtosis of the conditional distribution for short horizons but the magnitude of the coefficients in the regression shifts sign when looking at

horizons of nine weeks. This suggests a possible median reversion for longer horizons, predicting a higher probability of extreme events at longer horizons.

Finally, we find that two of the measures calculated, QD (QF) negatively (positively) predict the average kurtosis for horizons of 26 weeks. Additionally, we find that ED and EF negatively predict the median of realized kurtosis. The point estimates are negative for the QD and the ED, which become less negative as we approach the tails.

Since the existence of extreme observations hampers the ability to predict expected returns (Alti and Titman (2019)), in unreported results we drop observations after 2019 spanning the COVID crisis. Our results are also economically significant. For short horizons, taking the $\alpha = 0.75$ specification as benchmark, a one standard deviation (0.01) increase in the QD measure, is related to an increase of 0.27% (1.67%) in average market returns over one (26) week(s). Using the ED measure, we find that the same increase for an $\alpha = 0.75$ is related to an increase of 0.39% (2.39%) in expected returns over horizons of one (26) week(s). The specifications using one and 26 weeks as horizons provide positive out-of-sample R2s when using historical measures of asymmetry as a benchmark. For a one week horizon we obtain out-of-sample R2s between 0.28 and 1.68 for the QD measure and between 0.84 and 1.92 for the ED measure. For horizons of 26 weeks these out-of-sample R2s vary between 6.81 and 10.52 for the QD measure and between 6.62 and 9.4 for the ED measure. Although we obtain in-sample R2s as large as 5.42 for the regressions using nine week horizons, out-of-sample R2s are low, suggesting that forward-looking measures do not necessarily always beat the historical measures for predicting expected returns. With respect to the option-implied measures QA and EA we find that they can negatively predict expected returns for long horizons, although significant at the 5% and 10% level, the economic magnitude is small.

7 Conclusion

Option market data are forward-looking but noisy and highly non-linear financial assets. A proper empirical analysis of option-implied quantities thus requires solving these two problems to propose meaningful (arbitrage-free) option-based indicators. In this paper we propose the BIRS, a novel approach to fit the volatility smile that combines a cubic spline interpolation with a quadratic program that produces arbitrage-free final estimates, and we use it to estimate option-implied quantiles and expectiles which allows us to infer conditional robust indicators of risk and return. To reduce as much as possible the fact that the option-market data impose risk-neutrality we work with short term (weekly) options and investigate if the option-implied indicators have predictive information content. The proposed quantities appear to have some forecasting power - not shared by the equivalent quantities once inferred from backward-looking

historical returns, both in- and out-of-sample at short, medium and long time horizon and for both returns and volatility.

Bibliography

- Ait-Sahalia, Y. and J. Duarte (2003). Nonparametric option pricing under shape restrictions. *Journal of Econometrics* 116(1-2), 9–47.
- Alexiou, L. and L. S. Rompolis (2022). Option-implied moments and the cross-section of stock returns. *Journal of Futures Markets* 42(4), 668–691.
- Alti, A. and S. Titman (2019). A dynamic model of characteristic-based return predictability. *Journal of Finance* 74(6), 3187–3216.
- Ammann, M. and A. Feser (2019). Robust estimation of risk-neutral moments. *Journal of Futures Markets* 39(9), 1137–1166.
- Banz, R. W. and M. H. Miller (1978). Prices for state-contingent claims: Some estimates and applications. *The Journal of Business* 51(4), 653–672.
- Barone-Adesi, G. (2016). VaR and CVaR implied in option prices. *Journal of Risk and Financial Management* 9(1), 2.
- Barone-Adesi, G. and R. J. Elliott (2007). Cutting the hedge. *Computational Economics* 29(2), 151–158.
- Barone-Adesi, G., M. A. Finta, C. Legnazzi, and C. Sala (2019). Wti crude oil option implied var and cvar: An empirical application. *Journal of Forecasting* 38(6), 552–563.
- Barone-Adesi, G., N. Fusari, A. Mira, and C. Sala (2020). Option market trading activity and the estimation of the pricing kernel: A bayesian approach. *Journal of Econometrics* 216(2), 430–449.
- Barone-Adesi, G., C. Legnazzi, and C. Sala (2019). Option-implied risk measures: An empirical examination on the S&P 500 index. *International Journal of Finance & Economics* 24(4), 1409–1428.
- Bellini, F. and E. Di Bernardino (2017). Risk management with expectiles. *The European Journal of Finance* 23(6), 487–506.
- Bellini, F., L. Mercuri, and E. Rroji (2018). Implicit expectiles and measures of implied variability. *Quantitative Finance* 18(11), 1851–1864.
- Bellini, F., E. Rroji, and C. Sala (2022). Implicit quantiles and expectiles. *Annals of Operations Research* 313(2), 733–753.

- Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy* 81(3), 637–654.
- Bliss, R. and N. Panigirtzoglou (2002). Testing the stability of implied probability density functions. *Journal of Banking and Finance* 26, 381–422.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31(3), 307–327.
- Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected stock returns and variance risk premia. *Review of Financial Studies* 22(11), 4463–4492.
- Bondarenko, O. (2003). Estimation of risk-neutral densities using positive convolution approximation. *Journal of Econometrics* 116(1), 85–112.
- Bowley, A. L. (1920). *Elements of statistics. Vol. 2.* PS King.
- Breedon, D. and R. Litzenger (1978). Prices of State-Contingent Claims Implicit in Option Prices. *The Journal of Business* 51(4), 621–651.
- Buss, A. and G. Vilkov (2012). Measuring equity risk with option-implied correlations. *Review of Financial Studies* 25, 3113–3140.
- Campa, J. M., P. K. Chang, and R. L. Reider (1998). Implied exchange rate distributions: Evidence from OTC option markets. *Journal of International Money and Finance* 17(1), 117–160.
- Campbell, J. Y. and S. B. Thompson (2007). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies* 21(4), 1509–1531.
- Carr, P. and L. Wu (2009). Variance risk premiums. *Review of Financial Studies* 22(3), 1311–1341.
- Chabi-Yo, F. and J. Loudis (2020). The conditional expected market return. *Journal of Financial Economics* 137(3), 752–786.
- Christoffersen, P., K. Jacobs, and B. Y. Chang (2011). Forecasting with option implied information. *Handbook of Economic Forecasting*.
- Cochrane, J. (2008). The dog that did not bark: A defense of return predictability. *Review of Financial Studies* 21(4), 1533–1575.

- Conrad, J., R. Dittmar, and E. Ghysels (2012). Ex ante skewness and expected stock returns. *Journal of Finance* 68(1), 85–124.
- Driessen, J., P. J. Maenhout, and G. Vilkov (2009). The price of correlation risk: Evidence from equity options. *The Journal of Finance* 64, 1377–1406.
- Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* 50(4), 987–1007.
- Fengler, M. R. (2009). Arbitrage-free smoothing of the implied volatility surface. *Quantitative Finance* 9(4), 417–428.
- Figlewski, S. (2018). Risk-neutral densities: A review. *Annual Review of Financial Economics* 10, 329–359.
- Floudas, C. A. and V. Visweswaran (1995). Quadratic optimization. In *Handbook of Global Optimization*, pp. 217–269. Springer.
- Föllmer, H. and A. Schied (2016). *Stochastic Finance. An Introduction in Discrete Time* (Fourth ed.). Walter De Gruyter, Berlin.
- Ghysels, E., A. Plazzi, and R. Valkanov (2016). Why invest in emerging markets? The role of conditional return asymmetry. *Journal of Finance* 71(5), 2145–2192.
- Green, P. J. and B. W. Silverman (1994). *Nonparametric Regression and Generalized Linear Models. Monographs on Statistics and Applied Probability*, Volume 58. Chapman and Hall/London.
- Hentschel, L. (2003). Errors in implied volatility estimation. *The Journal of Financial and Quantitative Analysis* 38(4), 779–810.
- Hinkley, D. V. (1975). On power transformations to symmetry. *Biometrika* 62(1), 101–111.
- Jackwerth, J. (2000). Recovering risk aversion from option prices and realized returns. *Review of Financial Studies* 13(2), 433–451.
- Jackwerth, J. (2004). Option-implied risk-neutral distributions and risk aversion. *Charlottesville, VA: Research Foundation of AIMR*.
- Jackwerth, J. (2020). What do index options teach us about covid-19? *The Review of Asset Pricing Studies* 10(4), 618–634.

- Jackwerth, J. and M. Rubinstein (1996). Recovering probability distributions from option prices. *Journal of Finance* 51(5), 1611–1631.
- Jorion, P. (2007). *Value at Risk - The New Benchmark for Managing Financial Risk* (3rd ed.). The McGraw-Hill Companies.
- Koenker, R. (2005). *Quantile Regression*. Cambridge University Press.
- Koenker, R., V. Chernozhukov, X. He, and L. Peng (2017). *Handbook of Quantile Regression*. Chapman and Hall/CRC.
- Linn, M., S. Shive, and T. Shumway (2017). Pricing kernel monotonicity and conditional information. *Review of Financial Studies* 31(2), 493–531.
- Malz, A. M. (1997a). Estimating the probability distribution of the future exchange rate from options prices. *Journal of Derivatives* 5(2), 18–36.
- Malz, A. M. (1997b). Option-implied probability distributions and currency excess returns. *Staff Reports no. 32, Federal Reserve Bank of New York*.
- Mammen, E. and C. Thomas-Agnan (2002). Smoothing splines and shape restrictions. *Scandinavian Journal of Statistics* 26(2).
- Martin, I. (2017). What is the expected return on the market? *The Quarterly Journal of Economics* 132(1), 367–433.
- Metaxoglou, K. and A. Smith (2017). Forecasting Stock Returns Using Option-Implied State Prices. *Journal of Financial Econometrics* 15(3), 427–473.
- Molino, A. and C. Sala (2020). Forecasting value at risk and conditional value at risk using option market data. *Journal of Forecasting* 40(7), 1190–1213.
- Newey, W. and J. Powell (1987). Asymmetric least squares estimation and testing. *Econometrica* 55(4), 819–847.
- Newey, W. K. and K. D. West (1987). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55(3), 703–708.
- Rompolis, L. S. (2010). Retrieving risk neutral densities from european option prices based on the principle of maximum entropy. *Journal of Empirical Finance* 17(5), 918–937.
- Ross, S. (1976). Options and efficiency. *The Quarterly Journal of Economics* 90, 74–89.

- Ruppert, D. (1987). What is kurtosis? An influence function approach. *The American Statistician* 41(1), 1–5.
- Shimko, D. (1993). The bounds of probability. *Risk* 6, 33–37.
- Taylor, J. W. (2008). Estimating value at risk and expected shortfall using expectiles. *Journal of Financial Econometrics* 6(2), 231–252.
- Yao, Q. and H. Tong (1996). Asymmetric least squares regression estimation: A nonparametric approach. *Journal of Nonparametric Statistics* 6(2-3), 273–292.

Appendix

A The Quadratic Problem of Fengler (2009)

Without going into any mathematical detail, in this section we recall the Fengler (2009) quadratic program used to obtain a smooth and arbitrage-free pricing surface. The interested reader is referred to Fengler (2009) for further details. We first present the main steps and then repropose the same procedure in a summarized schematic approach.

For a sample of strikes and European call option prices $\{u_i, y_i\}, u_i \in [a, b]$:

$$\min_{\hat{g}} \sum_{i=1}^n w_i [y_i - g(u_i)]^2 + \lambda \int_a^b [f''(v)]^2 dv \quad i = 1, \dots, n \quad (38)$$

with strictly positive weights $w_i > 0$ and smoothing parameter $\lambda > 0$. The minimizer \hat{g} represents a globally arbitrage-free European call price function and needs to be twice differentiable. Moreover, assuming y_i to be a European call option with strike prices $a = u_0, \dots, u_n = b$ a function g is a cubic spline if on each sub-interval $(a, u_1), (u_2, u_3), \dots, (u_n, b)$ is a cubic polynomial and is twice differentiable $\mathcal{C}([a, b])$ such that:

$$g(u) = \sum_{i=0}^n 1\{[u_i, u_{i+1}]\} s_i(u) \quad (39)$$

where u_i are the knots of the spline. Green and Silverman (1994) shows that an alternative and convenient approach to represent the above cubic polynomial is given by the value second derivative representation of the natural cubic spline, which allows one to cast the optimization in 38 as a quadratic problem. Setting $g_i = g(u_i)$ and $\gamma_i = g''(u_i)$ and defining $g = (g_1, \dots, g_n)^T$ and $\gamma = (\gamma_2, \dots, \gamma_{n-1})^T$ with $\gamma_1 = \gamma_n = 0$, the natural spline can be completely specified by g and γ . Not all possible vectors of g and γ give rise to valid solutions, and the sufficient and necessary conditions for a valid solution are formulated via the matrices Q and R such that

$$\min_x -y^T x + \frac{1}{2} x^T B x \quad (40)$$

$$\text{subject to } A^T x = 0 \quad (41)$$

where $A = (Q, -R^T)$ and B is strictly positive definite.

In summary: At each day t we solve the quadratic program:

$$\min_x -y^T x + \frac{1}{2}x' Bx \quad (42)$$

subject to the following constraints:

- $A^T x = 0$
- $\lambda_i \geq 0$ (to guarantee convexity)
- $\frac{g_2 - g_1}{h_1} - \frac{h_1}{6} \gamma_2 \geq -e^{-\int_t^T r_s ds}$ (to guarantee the price function to be non-increasing in price)
- $-\frac{g_n - g_{n-1}}{h_{n-1}} - \frac{h_{n-1}}{6} \gamma_{n-1} \geq 0$ (again to guarantee the price function to be non-increasing in price)
- $g_1 \leq e^{-\int_t^T q_s ds} S_t$ (to prevent arbitrages)
- $g_1 \leq e^{-\int_t^T q_s ds} S_t - e^{-\int_t^T r_s ds} u_1$ (to prevent arbitrages)
- $g_n \geq 0$ (to prevent zero or negative prices)

where: $x - (g^T, \gamma^T)^T$. To prevent calendar arbitrages, the same quadratic program could also be applied to volatility surfaces by replacing the fifth constrained with:

$$g_i^j < e^{-\int_{t_j}^{t_{j+1}} q_s ds} g_i^{j+1}, \quad \text{for } i = 1, \dots, n \quad (43)$$

for t_j with $j = 1, \dots, m$. For this paper $j = 1$, fixed at one week.

B The PCA of Bondarenko (2003)

Let L^d be the set of all possible probability densities, of which $L_1(-\infty, \infty)$ denotes the set of all nonnegative functions that integrate to one. For the convolution, we start by fixing a kernel function $\phi(x) \in L^d$, which can be scaled with a bandwidth h to form a new density $\phi(x)_h := 1/h\phi(x/h)$. The role of h is to control for the smoothness of the densities, thus playing a crucial role in the final estimation. Given a fixed $\phi(x)_h$ we define an approximating set of all functions g as the set of admissible densities, W_h , which is obtained as the convolution of $\phi(x)_h$ and another positive density function u :

$$W_h := \{g \in L^d | g = \theta_h * u \text{ for some } u \in L^d\} \quad (44)$$

where g and f are some integrable function such that:

$$f * g := \int_{-\infty}^{\infty} f(x-y)g(y)dy \quad (45)$$

Then, the risk-neutral distribution is inferred by selecting a density $\hat{f}(x)$ from the set of admissible densities that optimally fits an empirically observed cross section of put option prices p_i with a finite set of strike prices $x_1 < \dots < x_n$:

$$\min_{\hat{f} \in W_h} \sum_{i=1}^n (p_i - D^{-2}\hat{f}(x_i))^2 \quad (46)$$

where $D^{-2}g(x) := \int_{-\infty}^x (\int_{-\infty}^y)g(z)dz)dy$ represents the second integral of function $g(x)$.

While exposed in continuous form, the presented optimization problem is solved numerically through a discretization of the admissible sets:

$$W_h^{\Delta z} := \left\{ g \in L^d \mid g(x) = \sum_{j=-\infty}^{\infty} a_j \phi_h(x - z_j) a_j \geq 0 \quad \sum_{j=-\infty}^{\infty} a_j = 1 \right\} \quad (47)$$

where Δz is the grid step used to determine an equally spaced grid so that for $j = 0, \pm, \dots$ the grid is defined over the real line as $z_j = \Delta z j$. It is worth noticing that $W_h^{\Delta z} \subset W_h$, and that the distance among the two sets is determined by the grid step Δz . Given a discretized admissible set, the discretized convolution is achieved on the equally spaced grid through a mixture of Dirac delta functions and the basis density:

$$f = \phi_h * u \text{ where } u(x) = \sum_{j=-\infty}^{\infty} a_j \delta(x - z_j) a_j \geq 0 \quad \sum_{j=-\infty}^{\infty} a_j = 1 \quad (48)$$

Throughout our analysis we follow Bondarenko (2003) and set the bandwidth $h = 0.95h_0$, the grid step $\Delta z = 0.5h$ and the basis density equal to the standard normal distribution:

$$\theta(x) = n(x) := \frac{1}{\sqrt{2\pi}} e^{-x^2/2} \quad (49)$$

Finally, as the optimization problem is a minimization with respect to the observed put option prices, we convert all call prices into put prices via the put-call parity.

C The FS approach of Jackwerth (2004)

The FS approach recovers the risk-neutral distribution in three phases. First, we collect the Black-Scholes option-implied volatilities $\bar{\sigma}_i$ for all available strike prices $i = 1, \dots, I$. To do it, we discretize the price grid $S_j = S_0 + j\Delta$ for $j = 1, \dots, J$ for Δ equal to the difference between two adjacent strike prices, so that the price grid coincides with the strike grid. Secondly, we minimize the objective function:

$$\min_{\sigma_i} \sum_{j=0}^J (\sigma_j'')^2 + \frac{(J+1)\lambda}{I\Delta^4} \sum_{i=0}^I (\sigma_i - \bar{\sigma}_i)^2 \quad (50)$$

where λ is a smoothing parameter that determines the trade off between smoothness and goodness of fit, σ_i and $\bar{\sigma}_i$ are the implied-volatility and the observed implied-volatility associated with strike price $i = 1, \dots, I$, respectively. The same holds for σ_j for $j = 1, \dots, J$. The second derivative σ_j'' is numerically approximated by $\sigma_j'' = (\sigma_{j-1} - 2\sigma_j + \sigma_{j+1})/\Delta^2$, where delta is the difference between two adjacent strike prices. Variables σ_j, σ_i are selected so that the curvature of the volatility curve is minimized and the estimated volatility curve agrees with the observed volatilities. The role of the penalty is to regulate the trade-off between the smallest squared second-order derivative of the implied volatility curve (smoothness), and the minimum of the sum of the squared errors, which is the distance between the estimated and the implied volatilities (fit). Finally, after solving for the optimal volatility function, these quantities are used to compute the Black and Scholes option prices and infer the risk-neutral distribution by differentiating twice the obtained option prices with respect to the strike prices. For our empirical applications (Section 4), while we fix $\Delta = 2.5$, we follow Jackwerth (2004) and calibrate λ such that the inferred risk-neutral distributions are positive and smooth. Results are summarized in table 11.

	mean	std	min	25%	50%	75%	max
Trade-off term	0.000188	0.000352	0.000000	0.000004	0.000016	0.000113	0.001491

Table 11: **Summary statistics trade-off term:** The table provides summary statistics on the trade-off term $(J+1)\lambda/(I\Delta^4)$ as in Jackwerth (2004).