Working Paper Series | 66 | 2024

Option-implied bond spread risk

This paper uses bond-option data to gauge market uncertainty and the magnitude of potential yield and spread moves in case of shocks to sovereign debt markets.

Gergely Hudecz
European Stability Mechanism

Edmund Moshammer
European Stability Mechanism

Marco Onofri
European Stability Mechanism

European Stability Mechanism

Option-implied bond spread risk

Gergely Hudecz [1] European Stability Mechanism
Edmund Moshammer [2] European Stability Mechanism
Marco Onofri [3] European Stability Mechanism

Abstract

Government bond yield futures and related option contracts contain information on the asymmetry of interest rate risks. We construct probability distributions of market-implied bond yield expectations up to 90 calendar days ahead between January 2018 and December 2023. We derive daily distributions for German, French, and Italian bond yields as well as bivariate distributions using a copula to analyse tail risks in bond spread movement. We confirm options to be useful in predicting bond yields and spreads when benchmarking against backward-looking models. Furthermore, we find tail spread measures to be correlated with stock market volatility, inflation expectations, monetary policy surprises, and global economic conditions. In the period under scrutiny, the correlation between these indicators and the Italian spread tail is stronger than the one with the French measure. While changes in global economic conditions and central bank asset purchases strongly correlate with the Italian spread tail, these are less relevant for the French one.

Keywords: Financial market, sovereign bond yield, risk premium, euro area, option contract, risk-neutral distribution, probability density function, copula model

JEL codes: G13, G17

[1] g.hudecz@esm.europa.eu
[2] e.moshammer@esm.europa.eu
[3] m.onofri@esm.europa.eu

ISSN 2443-5503
ISBN 978-92-95223-57-8
doi: 10.2852/7067368
EU catalog number: DW-01-24-002-EN-N

Option-implied bond spread risk*

Gergely Hudecz[§], Edmund Moshammer[§] and Marco Onofri[‡]

[§]European Stability Mechanism
[‡]European Stability Mechanism & Ghent University

Abstract

Government bond yield futures and related option contracts contain information on the asymmetry of interest rate risks. We construct probability distributions of market-implied bond yield expectations up to 90 calendar days ahead between January 2018 and December 2023. We derive daily distributions for German, French, and Italian bond yields as well as bivariate distributions using a copula to analyse tail risks in bond spread movement. We confirm options to be useful in predicting bond yields and spreads when benchmarking against backward-looking models. Furthermore, we find tail spread measures to be correlated with stock market volatility, inflation expectations, monetary policy surprises, and global economic conditions. In the period under scrutiny, the correlation between these indicators and the Italian spread tail is stronger than the one with the French measure. While changes in global economic conditions and central bank asset purchases strongly correlate with the Italian spread tail, these are less relevant for the French one.

Keywords: financial market, sovereign bond yield, risk premium, euro area, option contract, risk-neutral distribution, probability density function, copula model

JEL: G13, G17

*The views expressed are those of the authors and do not necessarily reflect those of the European Stability Mechanism. We thank for useful comments and suggestions G. Amisano, J. Baran, N. Hansen, M. Iseringhausen, M. Khuel, Y. Mimir, P. Mueller, E. Vangelista, L. Ricci, R. Strauch, K. Theodoridis, I. Vazouras, and participants at internal seminars at the European Stability Mechanism.
Corresponding author: Marco Onofri; 6a, Circuit de la Foire Internationale, L-1347 Luxembourg

1 Introduction

Anticipating changes in bond yields as well as movements of sovereign spreads represents an important challenge for economists, investors, and policy makers. In this paper, we build on previous work aimed at extracting risk-neutral distributions from option prices,[1] and we present a methodology for using bond options to analyse risks in European sovereign debt markets. Similarly to Bauer and Chernov (2024), we construct risk-neutral distributions implied by European bond options, and we derive the probability distribution of market-implied bond yield expectations on a daily basis. We contribute to the literature by constructing and back-testing a bivariate risk-neutral distribution that allows us to study movements in sovereign spreads.

Option prices can be seen as "bets" placed by market participants on future sovereign bond yields. An option contract gives the buyer the possibility to buy or sell an underlying security (e.g., a sovereign bond) at a certain price (called "strike" price) and at a specified date.[2] As a result, the price of an option depends on what probability the market attaches to reaching the strike price in the future. Therefore, option-implied distributions contain valuable information on the dispersion that characterises expectations of sovereign yields. Furthermore, if returns of two different government bonds are expected to move in opposite directions, the respective option pricing should reflect this. In other words, option prices reflect different expectations about economic developments, and indicate the market-implied probability of certain outcomes to materialise within a given time frame. If all market participants had access to the same information, the same firm expectations about the state of the economy, and shared the same preferences, they would converge on a single

[1] The risk-neutral distribution of an asset's returns implies that all risks have been hedged away and it is free from any investor's risk preferences (see Figlewski, 2018). It has to be noticed that risk-neutral probabilities differ from the physical ones. As discussed in Martin and Shi (2023), risk-neutral probabilities of market crashes are widely used by practitioners due to their appealing simplicity, and do perform relatively well in predicting crashes in individual stocks. However, these measures overstate the true probabilities to the extent that they give too much weight to bad states of the world, thereby generating a bias.

[2] The option contracts analysed in this paper are European Options. Among other differences, owners of American-style options may exercise at any time before the option expires, while European-style options may be exercised only at expiration.

yield, and there would be no reason to trade options. The existence of a distribution of expected yields indicates that this is not the case. Investors might have different preferences, or different information, or there could be disagreement about the expected state of the economy. Furthermore, we find that option-implied distributions are time-varying, and their shape evolves according to the state of the economy. This suggests that bond option pricing evolves over time as investors receive new information, change their views about the economy, or change their preferences based on the expected state of the economy. In a stressed or uncertain environment, disagreement among market participants increases. As a result, the difference between the "bets" on higher and lower yields increases, so the distribution becomes wider.

Our analysis is based on prices of option contracts for Italian, French, and German government bond futures up to 90 calendar days ahead. Based on these, we obtain the cumulative distribution function and the probability density function of bond futures 45 days ahead. We construct option-implied confidence bands using percentiles of these distributions and derive a measure of bond-yield volatility. Both measures show variation over time coinciding with external shocks and changes in the macroeconomic environment. We show that the option-implied volatility is correlated with other known sovereign bond volatility indicators, based on backward-looking models such as a GARCH. The main difference between these measures is in the underlying data, as option prices are responsive to new information and reflect changes possibly occurring within a trading day. This forward-looking feature represents a great advantage which is specific to option-implied metrics. When comparing them to similar measures of uncertainty based on backward-looking models, we do find the first to be more responsive than the latter to changes in economic conditions. Furthermore, their ability to predict both average and tail interest rate realisations is in most of the cases either comparable or higher than the one characterising backward-looking models.

In a next step, we extend the analysis to the multivariate dimension and construct bivariate risk-neutral distributions of Italian and French bond spreads vis-a-vis Germany. The aim of the exercise is to analyse the co-movement of yields across different government bonds, and assess the probability of a change in spreads. We discuss the performance of these bivariate distributions comparing different copulas, and comparing the log-likelihood of the risk-neutral distribution with those of random walk models over time. Furthermore, we use the bivariate distribution to construct (daily) tail spread measures. These are defined as the spread change associated with the 95th percentile of the corresponding distribution. We find that the bivariate risk-neutral distribution performs best in explaining realised bond spreads compared with three different more naive model specifications.

Finally, we use a regression model to analyse the correlations of Italian and French spread tails with monetary surprises as well as high-frequency proxies for macroeconomic and financial conditions. Results of this exercise show that the Italian spread tail measure is more responsive than its French counterpart to inflation expectations and stock market volatility. The estimated coefficients are statistically significant for French spreads, but they are considerably larger when the Italian spread tail is the dependent variable. However, while global economic conditions and central bank purchases strongly correlate with the BTP-Bund spread tail, their correlation with movements in the OAT-Bund one is weaker. These results suggest that the Italian bond spread to Bunds is more volatile than the French one, and an indicator of uncertainty related to spread tail realisations is also more sensitive to changes in macroeconomic and financial conditions. In other words, in case of a deterioration of economic and/or financial conditions, sudden sovereign bond spread increases are more likely to affect Italy. This may be because lower-rated countries can be perceived as more risky. On the other hand, spreads to Bunds of higher-rated countries, such as France, are expected not to react as strongly. The difference in not only the size, but also the significance of coefficients related to monetary policy surprises suggests that the Italian spread tail is relatively more exposed to changes in the level of asset purchases. Furthermore, the higher sensitivity characterising lower-rated spreads could be also explained by a flight-to-safety dynamic. When macroeconomic and financial conditions deteriorate, and the overall level of risk increases, investors substitute assets that are perceived as risky with safer ones, spreads widen, uncertainty rises even more, and the spread tail increases.

The rest of the paper is structured as follows. Section 2 discusses the economic and financial literature that inspired our paper, and Section 3 describes the data and the methodologies. We compare option-implied metrics with their counterparts obtained from backward-looking models in Section 4. Some intuitions about the degree of correlation between the spread tail measures and macroeconomic and financial indicators are provided in Section 5, while Section 6 concludes.

2 Related literature

A large body of literature spanning from theoretical finance to empirical macro-finance inspired our paper.

First, we build on studies that analyse how to derive option-implied distributions and extract information from option contracts (see Christoffersen et al., 2013 for a survey). As discussed in Breeden and Litzenberger (1978) and Banz and Miller (1978), it is possible to extract an option-implied density from European options having continuous strikes. However, options are often not liquid enough, and strikes are not continuous, which makes it necessary to impose more structure on the option-implied density. Depending on the number of restrictions to be imposed, possible approaches span from model-free non-parametric density estimation methods (see Aït-Sahalia and Lo, 2000; Li and Zhao, 2009; Trolle and Schwartz, 2014) to the parametric approach based on Black-Scholes implied volatility functions (see among many others Shimko, 1993; Rubinstein, 1994; Baran and Vorisek, 2020). Our paper adopts the methodology described in Baran and Vorisek (2020), and extends it by constructing bivariate distributions that allow us to account for the correlation and the co-movement of two option-implied distributions.

Our paper also builds on the literature for spread options. Spread options are options on the difference of two underlying asset prices. Different methods such as monte carlo simulation methods are commonly applied to price spread options (Carmona and Durrleman, 2003). In constructing the bivariate distribution, we draw extensively from the results of Berton and Mercuri (2023), who demonstrate practical advantages of using a copula function methodology over monte carlo simulations.

Furthermore, our paper is related to the literature that analyses how to use the information contained in option-implied distributions to address macro-financial questions within a policy framework. In this respect, Mc Manus (1999) analyses Eurodollar options to monitor the market sentiment related to exchange rate evolutions, while several studies analyse the interaction between monetary policy and bond market behaviour (Gürkaynak and Wright, 2012). Among them, a few make use of information contained in option-implied distributions. Puigvert-Gutiérrez and de Vincent-Humphreys (2012) construct a set of probability density functions based on Euribor futures contracts, and discuss their potential use within the European monetary policy framework in the aftermath of the global financial crisis. They suggest that option-implied market expectations of the Euribor-Eonia spread provide a timely and quantitative indication of the degree of uncertainty around the Euribor forward, as well as a precise quantification on how that is distributed over different possible outcomes. In a related paper, Vähämaa (2005), analyses the behaviour of market expectations in the surrounding of monetary policy announcements of the European Central Bank (ECB, hereafter) using option contracts with underlying German bond futures. They measure asymmetries in bond yield expectations using the skewness of the corresponding option-implied distribution and show that these expectations are systematically asymmetric around policy announcements. Specifically, option-implied distributions are positively skewed around tightening episodes, while the opposite applies to monetary policy easing. A similar question has been also addressed by Bauer et al. (2022) and Bauer and Chernov (2024) for the US economy. Bauer et al. (2022) analyse the role of uncertainty about future policy rates in the transmission of monetary policy to financial markets. They quantify monetary policy uncertainty using a model-free measure based on prices of Eurodollar futures and options. Event studies of Federal Open Market Committee (FOMC) announcements allow them to document the underlying drivers of changes in monetary policy uncertainty. Their results show that uncertainty about future policy rates exhibit substantial variation across FOMC announcements. Forward guidance typically lowers uncertainty, and most profound variations result from changes in the way forward guidance is described in FOMC statements. Finally, they document a novel monetary transmission channel by showing that changes in policy uncertainty have sizeable effects on asset prices, and that these are different from those stemming from shifts in expectations. In a related paper, Bauer and Chernov (2024) show that the conditional skewness of Treasury yields can very well predict monetary policy surprises, and does contain information about high-frequency changes of interest rates in the surrounding of monetary policy announcements. They also show that over the last 30 years, the conditional skewness is characterised by a substantial cyclical variation, and it is affected by macroeconomic conditions. In this paper, we do not focus on the interaction between bond-yield market expectations and monetary policy, but we add to the literature that shows how to extract information contained in option contracts by back-testing the bivariate risk-neutral distribution. Furthermore, as an application, we define and present a novel key indicator, the probability of an immanent bond spread increase, that may enrich early warning frameworks.

Our paper contributes to the literature that studies bond market predictability, such as Campbell and Shiller (1991), Cochrane and Piazzesi (2005), Ludvigson and Ng (2009), Cieslak and Povala (2015), and Bauer and Chernov (2024). In this respect, Bauer and Chernov (2024) argue that within this framework, the conditional skewness of Treasury yields is a good predictor that has not been studied yet. They measure it following two approaches, relying either on the realised or the implied skewness. The first is based on realised moments of yield changes, while the second, which is closer to the measure we present, builds on risk-neutral moments implied by Treasury option contracts. Our model comparison analysis supports their finding that the option-implied skewness is a powerful predictor in the context of bond market predictability. Specifically, when comparing the log-likelihood from different models we do find that the risk-neutral model (i.e., the one that accounts for the full interest rate distribution) is the best predictor for bond spread movements. Other more naive models based on purely empirical data, in particular those with symmetric distributions, such as GARCH or random walk, perform worse.

3 Data and Methodology

3.1 Data

To derive the measures proposed in this paper, we build on daily European option contracts traded on EUREX, having Italian, French, and German futures as underlying instruments. These contracts give the holder the right to sell/buy an asset (i.e., the sovereign future) at a specified price (strike), and date (expiry) to the seller of the option. For each day, we collect bulk data for all traded contracts between January 2018 and December 2023, which span a relatively large range of strikes. For each country and day, we collect data covering up to four horizons (i.e., expiring within four months at the latest). For each day and each horizon, we rely on around 70 pairs of put/call prices per country. The liquidity of option contracts seems to be comparable across the three countries, and in all cases the number of trades usually drops while approaching the expiry date. In this case, liquid strikes are mostly concentrated around the forward price of the underlying instrument. Because of the relatively low uncertainty, the resulting option price distribution is narrower and taller than the one based on contracts having a longer horizon.

For the benchmarking analysis discussed in Section 4, we rely on financial data such as sovereign Italian, German, and French 10-year yields as well as overnight index swaps having 5- and 10-year maturities. These data are used either as input to GARCH models when producing alternative volatility measures, or to specify random walk models.

Finally, in the regression analysis[3] discussed in Section 5, we measure stock market volatility using the Eurostoxx 50 VIX index, while global macroeconomic conditions are proxied by the world Citi economic surprise index. This variable is defined as the weighted average of differences between official economic results and forecasts over a moving window (i.e., a positive value signals economic conditions being better than expected). A great advantage of this variable in the context of the analysis presented in this paper, is to capture changes in economic sentiment at high frequency, which are likely to affect market conditions. Furthermore, the 5y5y inflation swap rate measures market-based inflation expectations, while central bank asset purchases are proxied by the quantitative easing shock proposed by Gnewuch (2022).[4] Compared to other instruments, the one proposed by Gnewuch (2022) has the great advantage of being identified by exploiting news about sovereign debt-based asset purchases. As opposed to other proxies for monetary policy conditions, such as high frequency surprises measured in the surrounding of central bank policy announcements, this instrument more accurately captures yield changes which are unrelated to movements in risk-free rates or risk premiums. The tailored focus on government yields and public bond purchases makes it the perfect candidate to study what drives a measure summarising risk and uncertainty that are specific to the sovereign bond market.

Apart from the quantitative easing shock of Gnewuch (2022), which is available on the website of the author, all other data have been provided by Bloomberg Enterprise LP.

3.2 A univariate risk-neutral distribution

Options are often not particularly liquid, and their prices are not continuous. As a result, an (empirical) density function constructed using only the quoted sample might be unstable. To overcome this limitation and augment the granularity of option prices, we prefer to fit the data to a parametric curve using interpolation and extrapolation techniques. By doing

[3] In this exercise, we cut the sample to October 2023, which is the last observation of the QE shock.

[4] As discussed in the paper, positive values indicate more purchases, hence a monetary policy easing.

so, it is possible to obtain a smoother density, without disregarding information from market pricing (see among others, Clews et al., 2000; Hagan et al., 2002; Malz, 2014; O'Donnell et al., 2016; Baran and Vorisek, 2020).

We approach this problem following the methodology described in Baran and Vorisek (2020).[5] Specifically, we first convert option prices to implied volatility using the Black option pricing model (Black and Scholes, 1973). Because of the data limitation discussed above, we increase the price granularity by fitting the volatility smile using the SABR model (see Hagan et al., 2002). Next, we input the interpolated implied volatility in the Black pricing model to recover (smooth and granular) option prices. As discussed in Appendix A, by taking the first and second derivative of the resulting option prices, we obtain the cumulative distribution function (CDF) and the probability density function (PDF) respectively, in terms of strikes.[6] We normalise the distribution, adjusting the values in the distribution so that the sum of all probabilities equals 1 and ensuring it to be continuous, by extrapolating tails beyond the 0.5th and 99.5th percentile to follow an exponential distribution tail.

Option contracts traded on EUREX have fixed expiry dates on monthly frequency. This makes comparisons of prices from one day to another difficult. However, these options are traded several months prior to their expiry date. We use data from the overlapping time spans to construct 45-day constant-horizon distributions. For each day we derive distributions as the linear interpolation of the distributions with expiry date before and after the horizon date.

3.3 A bivariate risk-neutral distribution

In this section we describe how we derive a measure which allows us to draw conclusions on the co-movement of yields across multiple governments, such as the probability of increasing

[5]Appendix B describes in detail the methodology used to derive risk-neutral distributions.

[6]As strikes are not intuitively interpretable, we convert both objects in terms of yields using the forward yield of the Cheapest To Deliver (CTD) bond and its conversion factor. To do this, we first assume that the yield at the 50% is given by the CTD forward. Next we compute price gaps vs the 50%, which we transform into yield gaps using the CTD conversion factor. Finally, given that the yield at the 50% is known, we recover values for all the other percentiles.

bond yield spreads. To this end, we construct a bivariate risk-neutral distribution. We apply the theorem that any multivariate joint distribution can be written in terms of the univariate marginal distribution functions derived from option prices and a copula which describes the dependence structure between the variables (Sklar, 1959). We estimate the copula parameters based on the realised empirical co-movement of changes in daily yields from 2000 to 2023, and test different functional forms, such as Gaussian copula, or Archimedean copulas, and find that the empirical data is best described by a Student-t copula (see Table 1). This model takes into account the correlation of the two marginal distributions and a second parameter to allow for heavier tails. We compute this copula function separately for every horizon and keep it constant for the whole time series. This copula, paired with the time series of marginal distributions, allows us to construct for each day and horizon the bivariate distribution.

4 Benchmarking

In this section we compare the option-implied bond yield and spread risk metrics (OBYR, hereafter) with other established measures commonly used to capture volatility. The two are expected to be highly correlated, with the only difference that the option-implied ones are more forward looking. We start by simply graphically comparing these two sets of metrics. Next, we compare their predicting accuracy in terms of both the average bond yields and tail outcomes.

4.1 Descriptive evidence

Having a long and daily time series of option-implied distributions enables us to infer the range of possible sovereign yields expected by the market within the specified horizon, as well as the probability attached to these expectations. Figure 1 shows option-implied distributions as well as the 5th and the 95th percentiles. In order to gauge the level of uncertainty characterising these expectation over time, it is possible to look at the evolution of these percentiles, thereby constructing option-implied confidence bands. This is shown in Figure 2. The bands appear to be relatively compressed during the first part of the sample, when uncertainty about inflation and interest rates was lower. In contrast, these became wider when inflation picked-up, the ECB started to tighten monetary policy and uncertainty increased. Figure 3 reports a time series indicator of the Bund-yield volatility calculated as the standard deviation of the distribution for Germany over time. This spiked at the outbreak of the pandemic, and started to trend upwards towards the end of 2021.

We compare these measures of volatility with other established volatility measures, as those estimated through a GARCH model in the spirit of Engle (1982) and Bollerslev (1986). Figure 4 provides a comparative analysis of OBYR and GARCH implied volatility of yields across Germany, Italy, and France. The visualisation indicates a robust correlation between these two volatility measures per country. Notably, OBYR demonstrates a tendency to slightly anticipate volatility pickups compared to GARCH. This phenomenon can be attributed to the inherent differences in the methodologies underlying these volatility measures. While GARCH models derive the volatility solely relying on historical data, the option-implied volatility encapsulates forward-looking information derived directly from market participants' expectations and sentiments. Therefore, OBYR inherently incorporates anticipatory elements, allowing it to provide insights into potential volatility movements before they manifest in historical data. Examining specific instances of market turbulence, such as the COVID-19 outbreak in Q1 2020 and the global uncertainty triggered by the Russian invasion of Ukraine in 2022, unveils interesting dynamics in the behaviour of OBYR. Compared to the much shorter anticipation evident in the volatility observed during the onset of the COVID-19 pandemic, markets exhibited clearer leading properties well before the geopolitical tensions surrounding the Russian invasion of Ukraine. This discrepancy underscores the nuanced interplay between market sentiment, forward-looking expectations, and historical data in shaping volatility dynamics. The findings above also hold for alternative volatility measures. The GARCH implied volatility of 3m10y yields (i.e., the market expectation of the 10-year government bond yield 3 months into the future), is virtually indistinguishable from the results reported in Figure 4. Figure 5 compares the standard deviation of OBYR of Germany and the GARCH implied volatility of 10-year and 5-year Overnight Index Swaps. Also here, OBYR demonstrates leading properties, particularly in early 2022.

Next, we use the bivariate distribution to extract the probability of the spread change to be above a given value. Figure 6 shows the estimated probability of the 45 day spread change to be higher than 15bps and 100bps between 2018 and 2023. This visualisation demonstrates a practical application in the context of risk assessment and early warning, providing insights into the likelihood of significant spread movements. The probabilities of spread changes exceeding 100bps offer insights into extreme market events and potential systemic risks that could threaten financial stability. In the case of France, these probabilities remained relatively subdued for the majority of the observed period, indicating a low likelihood of large spread increases. However, for Italy notable spikes in these probabilities were observed during specific events, such as the tumultuous government formation in 2018, the Covid-19 outbreak, and the energy crisis in the second half of 2022. Conversely, the probabilities of spread changes exceeding 15bps provide a more sensitive measure, capturing smaller yet significant fluctuations in spreads. This metric serves as a valuable tool for assessing market sentiment, and emerging trends that may precede larger spread movements. Notably, the comparison between the probabilities of 15bps and 100bps spread increases reveals interesting insights into the relationship between different levels of spread movements. While certain events may trigger smaller spread changes, they may not necessarily escalate into extreme or systemic risks, as evidenced by the varying levels of correlation between these probabilities.

The analysis of spread movement probabilities surrounding the Italian government formation in 2018 serves as a good example. During this period, Italy experienced significant political uncertainty, leading to heightened market volatility and risk aversion.

Already in March 2018 the probability of a 15bps Italian spread increase stood at a record high of 85%, but this dropped to 60% by end-May with the appointment of the new government. Conversely the probability of a 100bps increase in March was still relatively low at 2.5%, but it peaked at 18% with the appointment. This elevated probability reflected market concerns regarding the potential implications of the government formation process for Italy's fiscal stability and sovereign debt sustainability. This case study underscores the importance of incorporating geopolitical events and political risk factors into risk assessment frameworks, particularly in sovereign bond markets. The analysis of spread movement probabilities offers valuable insights into market sentiment, and the transmission of political events to financial markets. By leveraging these insights, policy makers and investors can better understand, anticipate, and manage risks associated with political uncertainty and its impact on sovereign bond spreads.

Figure 7 provides a different means to analyse the bivariate distributions. Instead of asking for the probability of large spread increases, Figure 7 plots the realised spread at the time and the 95th percentile of the predicted spread in 45-days, both for Italy-Germany and France-Germany. The pattern is similar to the one described for the previous plot. Periods characterised by a high probability of spread increases match those where the 95th percentile is significantly higher than the realised spread.

4.2 Model comparison: average and tail predictions

In this section, we present a battery of comparisons to cross-check the performance of the different model specifications. We compare metrics extracted from the risk-neutral densities (OBYR) to their counterparts obtained from models that are based on past information, which we consider as reference benchmark. As long as options are traded everyday, their prices should better incorporate expectations and sudden changes or incoming news arriving even within a trading day. In contrast, backward-looking models reflect information related to the past. As a result, option-implied metrics should be more forward looking, and the fit of OBYR is expected to be better than, or at least as good as, the one consistent with backward-looking models in the majority of the cases.

For univariate densities, we compare the performance of OBYR to those of an *Empirical* model, in which the prediction of rate changes is extracted from the distribution of past rates,[7] GARCH(1,1) (Engle, 1982; Bollerslev, 1986) and two normally-distributed random walk models. For the first one, which we label as *Random Walk*, the mean is equal to the forward yield and its standard deviation is given by the historical standard deviation of rate changes since 2000. In the second version of the random walk model, labelled *Random Walk Plus*, we replace the historical standard deviation with the option-implied one.[8]

We assess the performance across models using two methods. First, for each model, we derive from the corresponding density the likelihood of the observed yield after the set horizon. We aggregate this information computing the sum of log likelihood across dates, and compare performances across models. A model will score high on this metric, if it consistently assigns a high probability to the observed outcome. Particularly in tranquil times, the observed outcome must be close to the mode of the distribution and the standard deviation must be low. In periods of uncertainty, a higher scoring model typically has higher standard deviation, to account for potential tail events having a higher likelihood of occurring. This first metric serves the purpose of evaluating the ability of different models to predict the yield at the specified horizon on average in both tranquil and uncertain times. Second, we also evaluate the ability to predict tail-realisations. To this end, we use the tick loss, which is a standard loss function used to evaluate quantile estimates (see for instance Iseringhausen, 2024). It allows to evaluate how often each model is correct in predicting a yield realisation to be in one of the two tails of the respective distribution. In order to highlight the accuracy of predicting both left and right tails events, we compute this measure for the 5th and the 95th percentiles of each distributions. We rely on data for an expanding

[7] In other words, the predicted outcomes of sovereign rate changes, or the ones associated to a specific percentile, are just those corresponding to the average (or to a percentile) which characterise the distribution of the same interest rate changes over the past 20 years.

[8] We take as a proxy the inter-quartile difference.

window starting on 1 January 2018, and all backward-looking models use pre-sample data from January 2000. The tick loss function takes the following form.

$$TL_p = \frac{1}{T}\sum_{t=1}^{T}\left(y_{t+h} - \hat{Y}^p_{t+h|t}\right)\left(p - \mathbb{I}_{\{y_{t+h}<\hat{Y}^p_{t+h|t}\}}\right) \tag{1}$$

where t is the day, y is the realised rate, $\hat{Y}^p$ is the value predicted by the model associated to the percentile p, and $\mathbb{I}$ is the indicator function.

Figures 8 and 9 summarise the results of the cross-checks performed for the univariate distributions. In Figure 8, the performance comparison reveals that OBYR emerges as the strongest performer across all models for each of the three countries, based on aggregate likelihood. Particularly for Italy, characterised by significant volatility swings, GARCH follows closely (within the 95% confidence band) behind OBYR in terms of performance. Conversely, for France and Germany, the empirical model, with its constant assumed distribution, demonstrates reasonable performance, but still significantly worse than OBYR. Notably, the random walk models exhibit the weakest performance overall, except for Germany, where the *Random Walk Plus*, incorporating the option-implied mean and volatility, notably excels. This observation suggests that, for Germany, the option-implied skewness plays a less significant role, while market sentiment regarding uncertainty remains a pertinent predictor. To test significance we apply bootstrapping-based hypothesis testing (Efron and Tibshirani, 1993). Figure 9 presents a nuanced perspective by focusing on models' ability to predict tail realisations exclusively. In this analysis, OBYR demonstrates a balanced performance in predicting both large yield increases and decreases for Germany and France. However, for Italy, OBYR struggles more in predicting increases. Conversely, the *Random Walk Plus* emerges as a strong contender for predicting decreases, particularly for Italy, indicating that while skewness may be less relevant, knowledge of option-implied volatility remains valuable in forecasting downward movements.

For the multivariate case, we compare the OBYR benchmark model with three alternative model specifications. We compare the predicting performances in terms of the overall log-likelihood. Similarly to the univariate case, this aggregate measure is given by the sum of the log-probabilities extracted from the corresponding density across days. The alternative *Gaussian* model, also relies on option-implied risk-neutral univariate distributions like OBYR, but connects them with a Gaussian-copula instead of a t-copula. The *Random Walk* and *Random Walk Plus* models correspond to the respective models from the univariate test, but applying the same t-copula specification as OBYR to form multivariate distributions. As shown in Figure 10, comparing OBYR and *Gaussian*, OBYR with the t-copula performs significantly better. OBYR also performs significantly better than either of the random walk models, both for Italian and French spreads. To test significance at the 5% level we apply hypothesis testing using bootstrapping (Efron and Tibshirani, 1993). Figures 11 and 12 provide a more granular explanation, by displaying the log likelihood of all models across time, whereas the previous chart showed just the sum per model.

The *Random Walk* model is an interesting comparison in that there are some periods when it performs better, and some periods when it performs worse than OBYR. It performs worse than OBYR in particularly calm periods where it is penalised for its time-constant larger standard deviation, but also in the very volatile period during 2022, when markets were correctly pricing in the uncertainty. Before 2022 several of the "smaller" shocks were better predicted by the more conservative *Random Walk* model. *Random Walk Plus* matches very closely the predictive performance of OBYR for most of the sample, but in the few cases where they are surprised by a shock, *Random Walk Plus* is even more surprised. The main difference between *Random Walk Plus* and OBYR is that the latter also considers the skewness (and further moments) of the risk-neutral distribution. The result of this comparison emphasises the high relevance of the skewness of the option-implied distribution.

5 Spread tail analysis: a simple regression application

In this section we show that macroeconomic and financial variables correlate with the Italian and French spread tail measures. These are defined as the 95th percentile of the respective bivariate distribution of spread changes. Quantifying the correlation between these measures and macroeconomic and financial variables allows us to investigate whether the different risk profiles of Italian and French sovereign bonds translate into different degrees of co-movements with economic and financial indicators.

To this end, we rely on a simple OLS regression model. In order to exploit the high frequency of the option-implied spread tail measures, we perform the analysis using daily data. On the one hand, this is an appealing feature. Sovereign bonds are traded on a daily basis, and to study the relation between sovereign spreads and changes in economic conditions it is important to understand how market participants react to news enriching their information set over the same horizon. On the other hand, this is also a limitation, as it restricts the choices of explanatory variables. For instance, standard slow-moving macroeconomic indicators, such as growth and inflation, cannot be used. To overcome this limitation, we rely on high-frequency financial market proxies. The regression model takes the following form:

$$ST_t = \alpha + \rho ST_{t-1} + \beta_1 \pi_t^{5y5y} + \beta_2 VIX_t + \beta_3 ESI_t + \beta_4 QE_t + \eta_t \qquad (2)$$

where ST_t is the spread tail measure, π_t^{5y5y} is the 5y5y inflation swap rate, VIX_t represents the EuroStoxx50 VIX index, ESI_t is the global economic surprise index computed by Citigroup, and QE_t is the quantitative easing shock from Gnewuch (2022).

We estimate the model described in (2) using the Italian and French spread tail measures as dependent variables. Results are reported in Table 2, where the first column refers to the Italian spread tail, and the second to the French one. The first result that is worth discussing is the difference in the size of the estimated coefficients, with those referred to the BTP-Bund tail being substantially larger than the French ones. This suggests that the BTP spread is more volatile than the OAT one, and its correlation with changes in macroeconomic and financial conditions is stronger. This result is also consistent with the higher probability of BTP-Bund spread changes being above a given threshold (see Figure 10, for instance). The second result that is worth discussing is the difference in the significance of some coefficients. While inflation expectations and stock market volatility strongly correlate with both the BTP and the OAT spread tails, global macroeconomic surprises and central bank asset purchase shocks seem to influence only the Italian tail measure.[9] Albeit having the correct sign, the estimated coefficients related to these two variables turn out not to be significant for the French spread tail measure. This result also shows that the tail risk measure of Italian sovereign bonds is more sensitive to macroeconomic and financial conditions.

It is important to stress that this measure reflects the level of disagreement among financial market participants. On the one hand, regardless of the riskiness of the underlying asset, the disagreement about the state of the economy is expected to rise when uncertainty increases. On the other hand, this increases even more when the underlying asset is characterised by high risk. Investors anticipate that a deterioration in economic conditions could transmit more strongly to risky assets, which are potentially more vulnerable to the surfacing of non-linear risks. The stronger sensitivity of the Italian spread tail to changes in macroeconomic conditions could be explained by the fact that a worsening in the global surprise index signals a higher probability of a recession materialising in the coming months. Similarly, a faster-than-expected reduction of central bank sovereign bond purchases (i.e., a decrease in the QE factor of Gnewuch, 2022) would likely generate strong effects on the Italian yield and spread, while the French ones, being safer, would likely be less affected. In other words, during the period analysed in this work, if someone sneezed, concerns about

[9] A similar result emerges when testing whether coefficients for BTPs and OATs are significantly different from each other. We find those related to central bank purchases and the global economic surprise index to be statistically different at the 5% and 1% confidence level, respectively. On the other hand, coefficients related to inflation expectations are significantly different only at the 10% level, while we cannot reject the null hypothesis of coefficients related to stock market volatility not to be statistically different in the two specifications.

the BTP-Bund spread getting a cold appear to be sizeable. On the other hand, these seemed much milder for the French spread to Germany.

Beyond this mechanism, one could argue that in a crisis environment non-fundamental volatility might also weigh on the liquidity conditions in low rated sovereign bond markets. Investors might decide to rebalance their portfolio towards safer bonds, thereby worsening the liquidity conditions of "riskier" sovereign borrowers. As outlined in Lane (2020),[10] in a stress scenario, the high substitutability across sovereign bond markets in the absence of currency risk might generate self-fulfilling flight-to-safety dynamics and illiquidity in individual sovereign bond markets. In a related paper, Arcidiacono et al. (2024) investigate the extent to which a country's supply of sovereign bonds affects the convenience yield of another country, and identify large spillovers among equally safe bonds. On the other hand, spillovers from safe to riskier countries are weaker, suggesting that these are poor substitutes to hedge against idiosyncratic shocks.

6 Conclusion

In this paper, we rely on option prices to derive measures that allow us to analyse risks in European sovereign bond markets. Specifically, we exploit the information contained in option contracts for Italian, French, and German government bond futures to construct daily option-implied densities. The shape of these distributions varies with the state of the economy and we find these densities to be informative about market-implied expectations of future rate changes. Due to their forward-looking characteristics, we confirm that option prices are useful in predicting bond yields and spreads when benchmarking against backward-looking models.

By constructing a bivariate distribution to capture the correlation of Italian and French option-implied densities with the German one, we show how to use this measure to analyse co-movements across these sovereign bond markets, and assess the probability of a change in

[10]See the ECB blog post by Philip R. Lane "The monetary policy response to the pandemic emergency".

the related spread. Furthermore, we highlight the correlation of Italian and French spread tail measures with stock market volatility, inflation expectations, monetary policy surprises, and global economic conditions. In the period under scrutiny, we find the correlation between these indicators and the Italian spread tail measure to be stronger than the correlation with the French one.

The analysis in this paper offers an overview of how we construct option-implied bond yield and spread expectations as well as policy-relevant applications in the context of risk assessment and early warning. The market-based measures discussed in this paper could be used in further research addressing many questions within the empirical macro-finance literature. As an example, option-implied moments could be employed to analyse the effectiveness of monetary policy transmission mechanisms and their potential non-linear effects depending on skewness (or volatility) being high or low. Alternatively, the measures presented in this paper could be used to adapt the analysis described in Bauer and Chernov (2024) to the euro area economy. For instance, employing a similar methodology would make it possible to estimate the excess bond returns for euro area countries.

References

Y. Aït-Sahalia and A. W. Lo. Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94(1-2):9–51, 2000.

C. Arcidiacono, M. Bellon, and M. Gnewuch. Dangerous liaisons? debt supply and convenience yield spillovers in the euro area. *mimeo*, 2024.

R. W. Banz and M. H. Miller. Prices for state-contingent claims: Some estimates and applications. *Journal of Business*, pages 653–672, 1978.

J. Baran and J. Vorisek. Volatility indices and implied uncertainty measures of European government bond futures. *ESM Working Paper Series, 43, 2020*, 2020.

M. Bauer and M. Chernov. Interest rate skewness and biased beliefs. *The Journal of Finance*, 79(1):173–217, 2024.

M. D. Bauer, A. Lakdawala, and P. Mueller. Market-based monetary policy uncertainty. *The Economic Journal*, 132(644):1290–1308, 2022.

E. Berton and L. Mercuri. An efficient unified approach for spread option pricing in a copula market model, 2023.

F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3):637–654, 1973.

T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.

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

J. Y. Campbell and R. J. Shiller. Yield spreads and interest rate movements: A bird's eye view. *The Review of Economic Studies*, 58(3):495–514, 1991.

R. Carmona and V. Durrleman. Pricing and hedging spread options. *Siam Review*, 45(4): 627–685, 2003.

P. Christoffersen, K. Jacobs, and B. Y. Chang. Forecasting with option-implied information. *Handbook of Economic Forecasting*, 2:581–656, 2013.

A. Cieslak and P. Povala. Expected returns in Treasury bonds. *The Review of Financial Studies*, 28(10):2859–2901, 2015.

R. Clews, N. Panigirtzoglou, and J. Proudman. Recent developments in extracting information from options markets. *Bank of England Quarterly Bulletin*, 2000.

J. H. Cochrane and M. Piazzesi. Bond risk premia. *American Economic Review*, 95(1): 138–160, 2005.

B. Efron and R. J. Tibshirani. *An Introduction to the Bootstrap*. Chapman & Hall/CRC, 1993.

R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the econometric society*, 50(4): 987–1007, 1982.

S. Figlewski. Risk-neutral densities: A review. *Annual Review of Financial Economics*, 10: 329–359, 2018.

M. Gnewuch. Spillover effects of sovereign debt-based quantitative easing in the euro area. *European Economic Review*, 145:104133, 2022.

R. S. Gürkaynak and J. H. Wright. Macroeconomics and the term structure. *Journal of Economic Literature*, 50(2):331–367, 2012.

P. S. Hagan, D. Kumar, A. S. Lesniewski, and D. E. Woodward. Managing smile risk. *The Best of Wilmott*, 1:249–296, 2002.

M. Iseringhausen. A time-varying skewness model for Growth-at-Risk. *International Journal of Forecasting*, 40(1):229–246, 2024.

H. Li and F. Zhao. Nonparametric estimation of state-price densities implicit in interest rate cap prices. *The Review of Financial Studies*, 22(11):4335–4376, 2009.

S. C. Ludvigson and S. Ng. Macro factors in bond risk premia. *The Review of Financial Studies*, 22(12):5027–5067, 2009.

A. M. Malz. A simple and reliable way to compute option-based risk-neutral distributions. *FRB of New York Staff Report*, (677), 2014.

I. Martin and R. Shi. Forecasting crashes with a smile. *CEPR Discussion Paper No. 18524*, 2023.

D. Mc Manus. The information content of interest rate futures options. Technical report, Bank of Canada, 1999.

S. O'Donnell, M. O'Keeffe, et al. Option implied probability density functions: Methodology and use in understanding investor sentiment. *Quarterly Bulletin Articles*, pages 85–94, 2016.

J. M. Puigvert-Gutiérrez and R. de Vincent-Humphreys. A quantitative mirror on the Euribor market using implied probability density functions. *Eurasian Economic Review*, 2:1–31, 2012.

M. Rubinstein. Implied binomial trees. *The Journal of Finance*, 49(3):771–818, 1994.

D. Shimko. Bounds of probability. *Risk*, 6(4):33–37, 1993.

A. Sklar. Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris*, 8:229–231, 1959.

A. B. Trolle and E. S. Schwartz. The swaption cube. *The Review of Financial Studies*, 27 (8):2307–2353, 2014.

S. Vähämaa. Option-implied asymmetries in bond market expectations around monetary policy actions of the ECB. *Journal of Economics and Business*, 57(1):23–38, 2005.

Tables and figures

Table 1: Likelihood of copula-models for explaining historical sovereign bond yield co-movements

Countries	Copula	Log-likelihood $\frac{1}{1000}$	Correlation	Other parameters
Germany-Italy	Gaussian	1.99	0.60	
	Student-t	2.96	0.71	df = 2.53
	Clayton	1.36		$\alpha = 0.98$
	Frank	2.61		$\alpha = 5.91$
	Gumbel	2.69		$\alpha = 1.91$
Germany-France	Gaussian	8.92	0.93	
	Student-t	10.84	0.95	df = 1.65
	Clayton	7.56		$\alpha = 4.64$
	Frank	9.40		$\alpha = 18.69$
	Gumbel	10.10		$\alpha = 4.96$

Note: Estimated on a sample of 45-day changes of daily German and Italian/French government bond yields between 2000 and 2023. Copula marginals refer to empirical option-implied distributions.

Source: ESM calculation based on Bloomberg LP data.

Table 2: Spread tails analysis

	(1) BTP spread tail	(2) OAT spread tail
LagDepVar	0.962***	0.968***
	(98.31)	(137.12)
5y5y Inflation	0.0151***	0.00535***
	(2.85)	(3.47)
VIX EuroStoxx50	0.113*	0.0271**
	(1.86)	(2.00)
Global ESI	-0.0149***	-0.000645
	(-2.75)	(-0.63)
QE shock	-0.923**	-0.165
	(-2.30)	(-1.58)
Constant	-2.181**	-1.005***
	(-2.20)	(-3.70)
N	1535	1535

t statistics in parentheses

* p<.1, ** p<.05, *** p<.01

Source: ESM calculation based on Bloomberg LP data and Gnewuch (2022)

Figure 1: Broad measure of market expected interest rate range

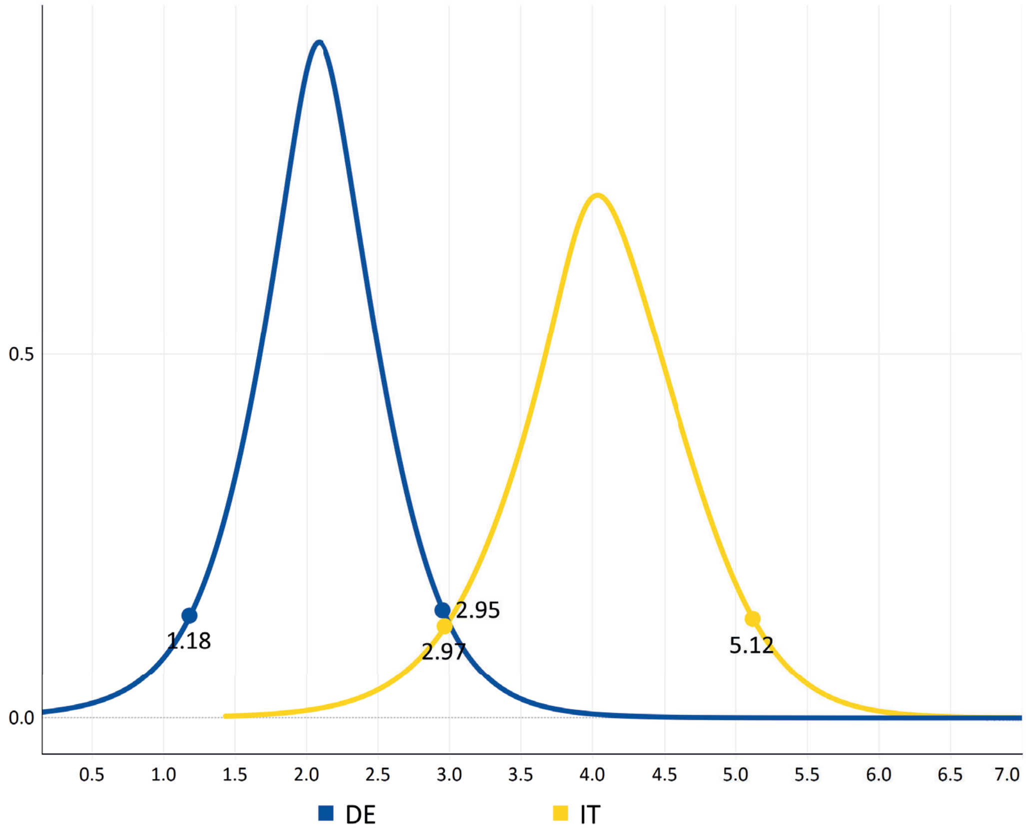

Note: Data for Germany and Italy on 7/3/2023 for 45 days ahead. The dots correspond to the 5th and the 95th percentiles.

Source: ESM calculation based on Bloomberg LP data.

Figure 2: Option-implied confidence bands

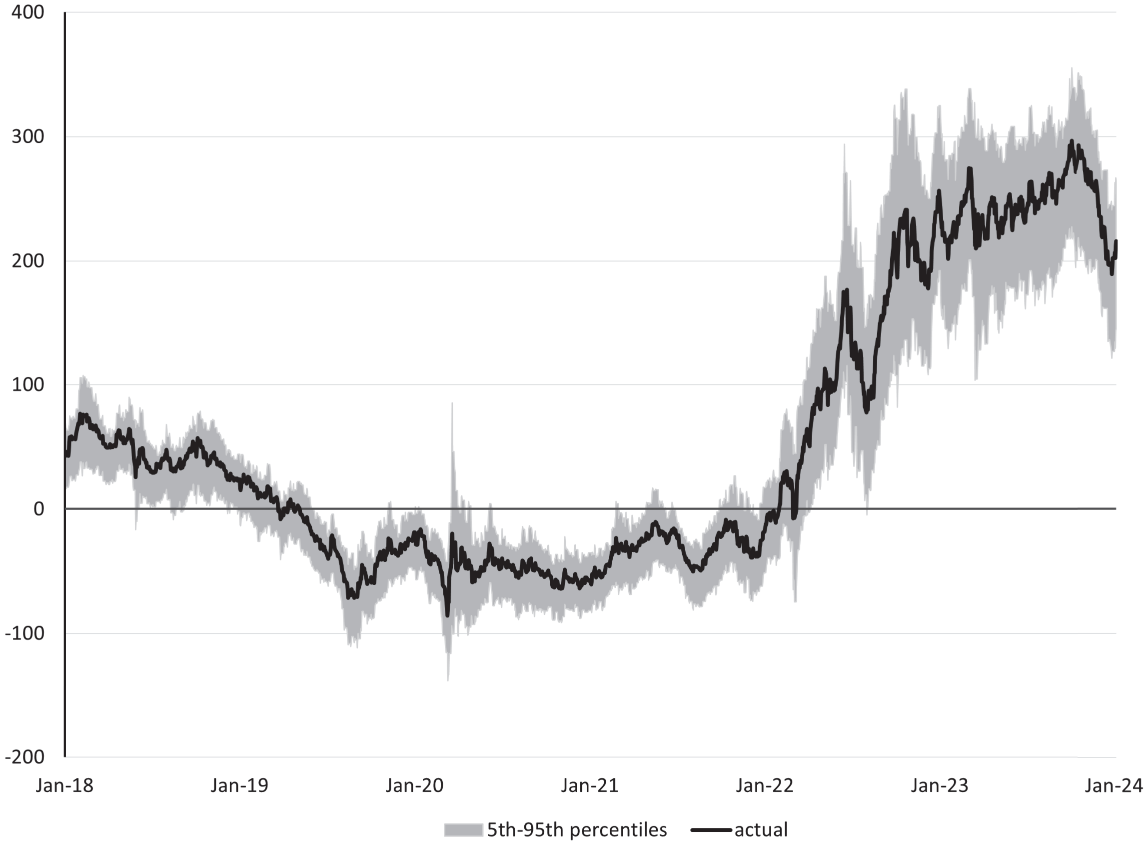

Note: Data for Germany for 45-day-ahead option-implied distribution of bond yields (in basis points). The bands are constructed using the 5th and the 95th percentiles.
Source: ESM calculation based on Bloomberg LP data.

Figure 3: Option-implied standard deviation for Germany

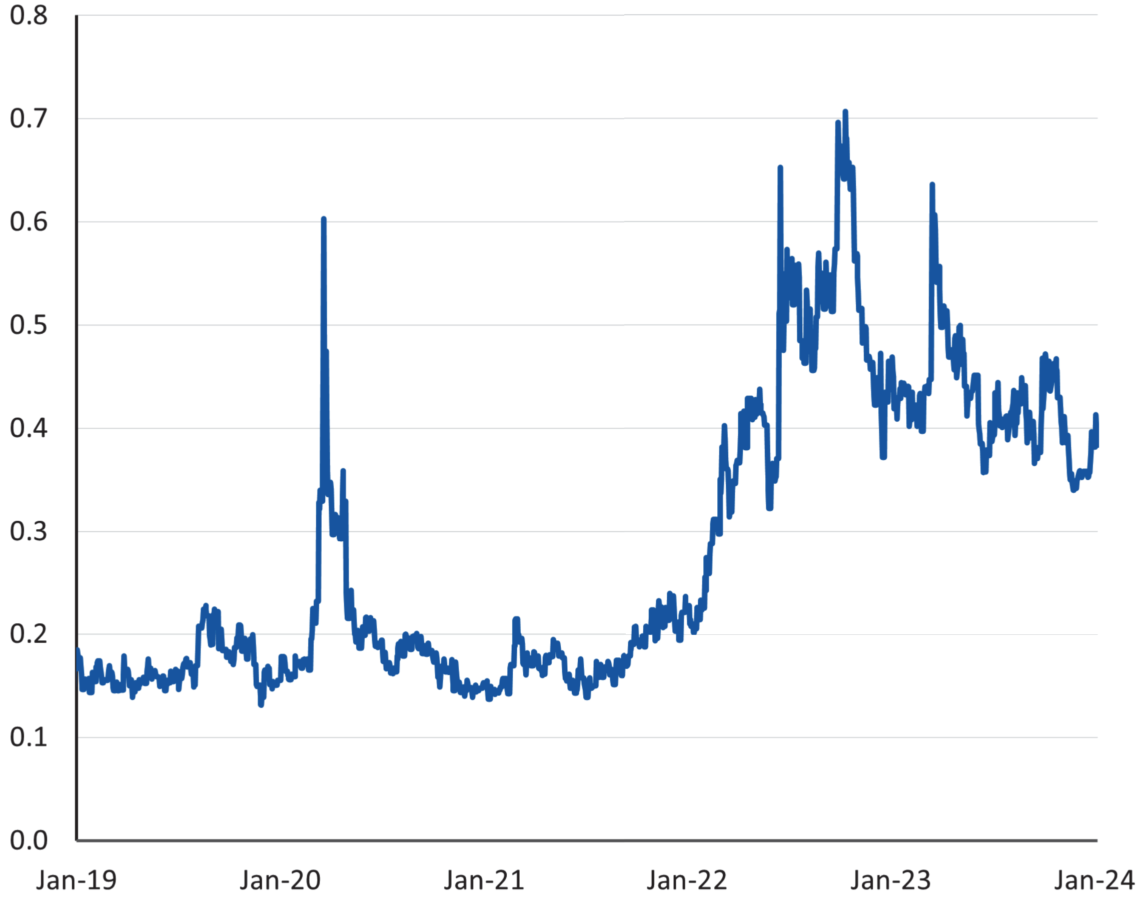

Note: Standard deviation of the distribution of 45-day-ahead option-implied bond yields for Germany.

Source: ESM calculation based on Bloomberg LP data.

Figure 4: OBYR and yield volatility (GARCH)

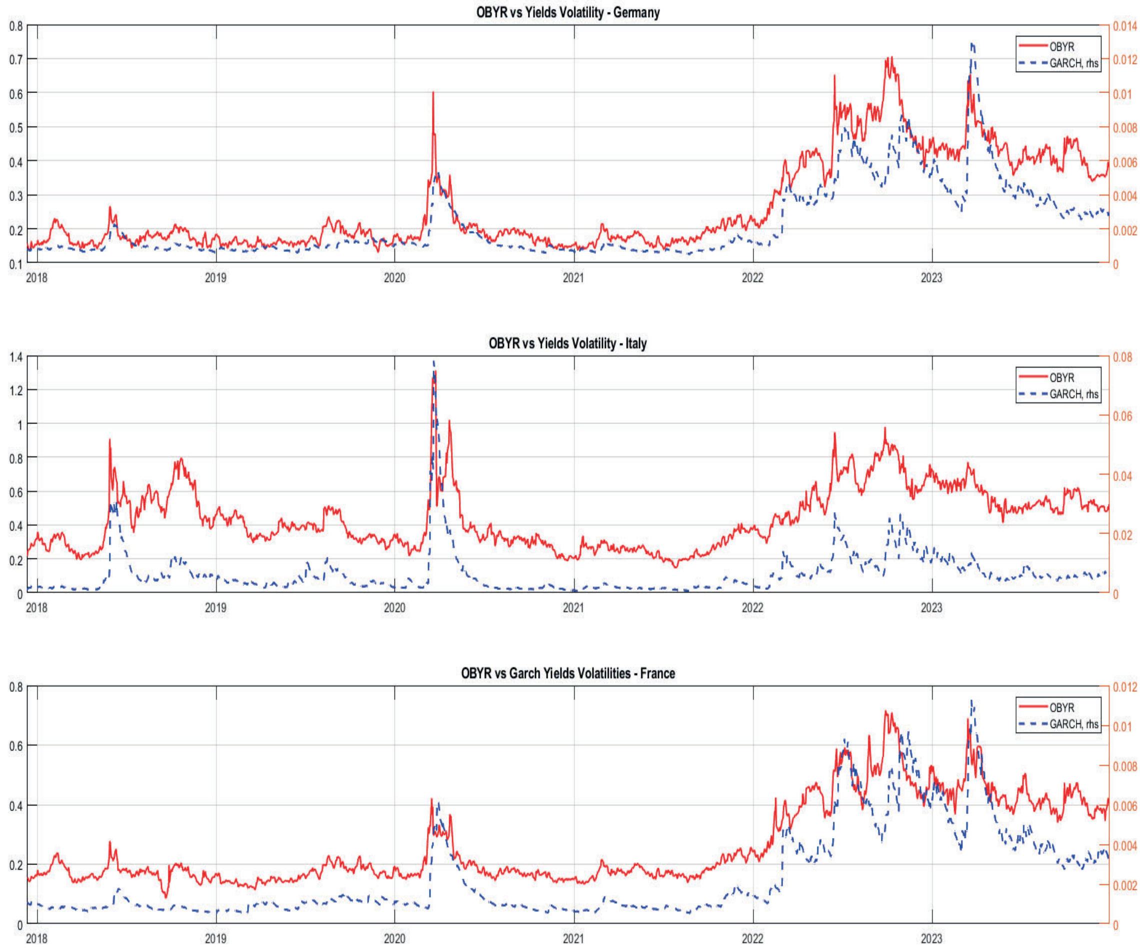

Note: OBYR (lhs axis) refers to the volatility of the 45-day-ahead option-implied bond yield distribution. GARCH (rhs axis) refers to the 45-day-ahead implied volatility of a GARCH(1,1) model.

Source: ESM calculation based on Bloomberg LP data.

Figure 5: OBYR for Germany and OIS volatility (GARCH)

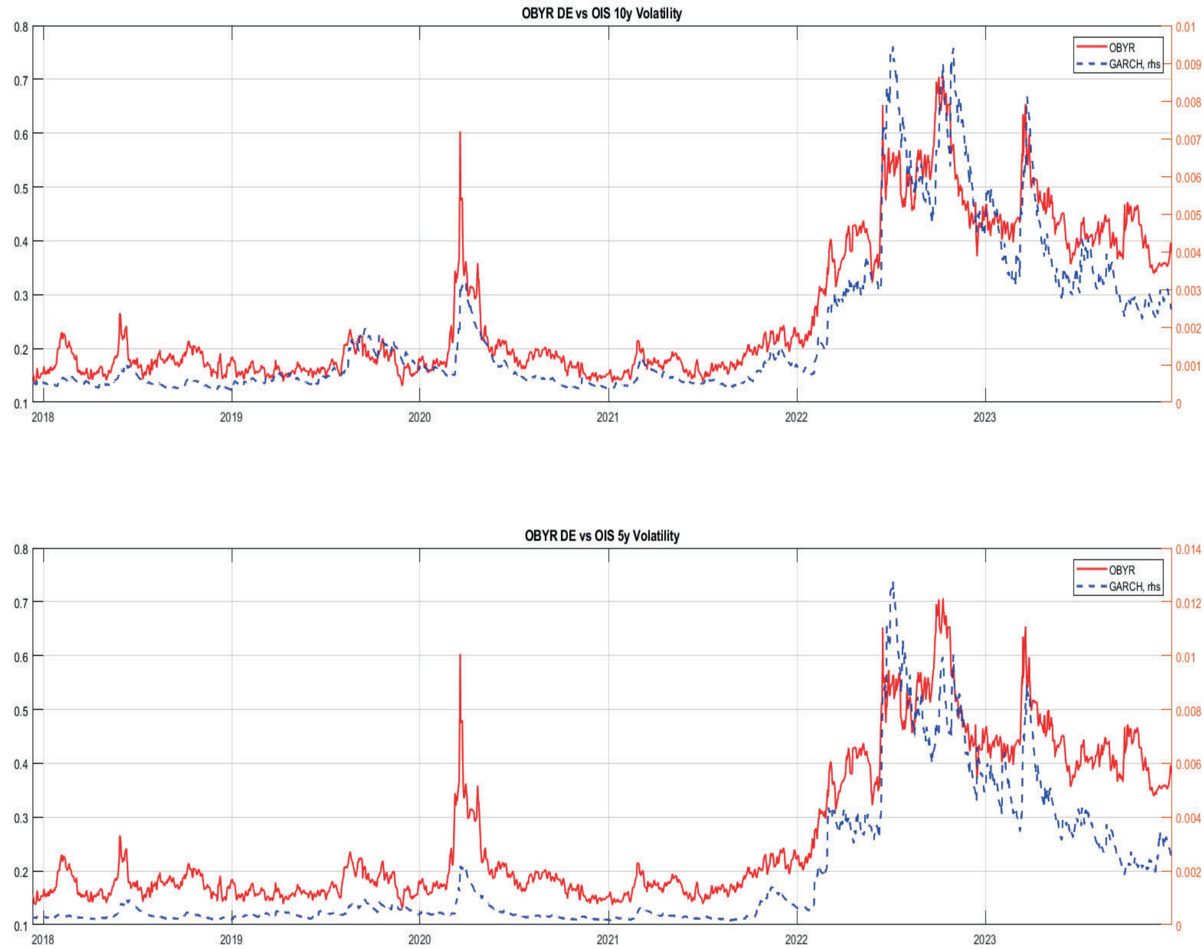

Note: This figure compares the standard deviation of the 45-day-ahead option-implied bond yield distribution of Germany and the GARCH implied volatility of 10-year and 5-year Overnight Index Swaps.

Source: ESM calculation based on Bloomberg LP data.

Figure 6: Bivariate distribution - probability of spread increase

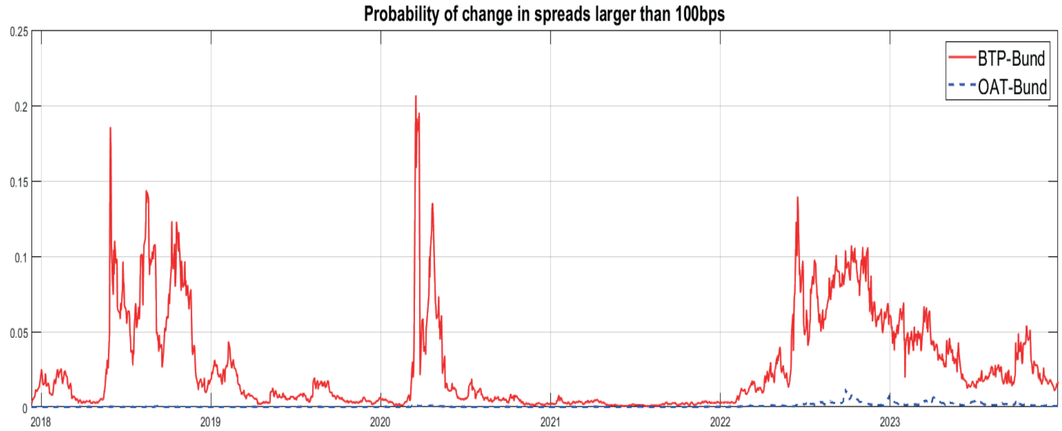

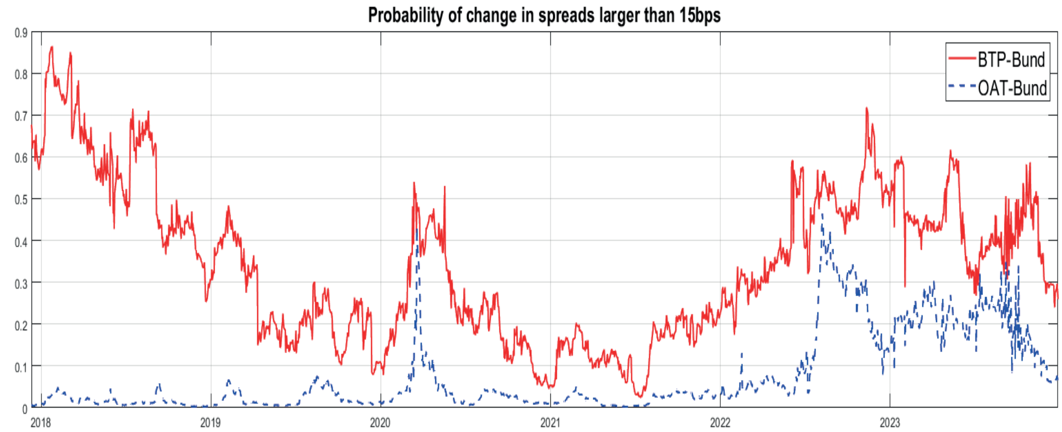

Note: The bivariate distribution relating changes in Italian and German yields (BTP-Bund, red line) and the one relating changes in French and German yields (OAT-Bund, blue line) are derived from univariate 45-day-ahead option-implied bond yield distributions using a copula approach. We extract the probability of an increase in spreads being larger than 100bps and 15bps, by sampling from the bivariate copula distributions.

Source: ESM calculation based on Bloomberg LP data.

Figure 7: Tails vs realised Italian and French spreads

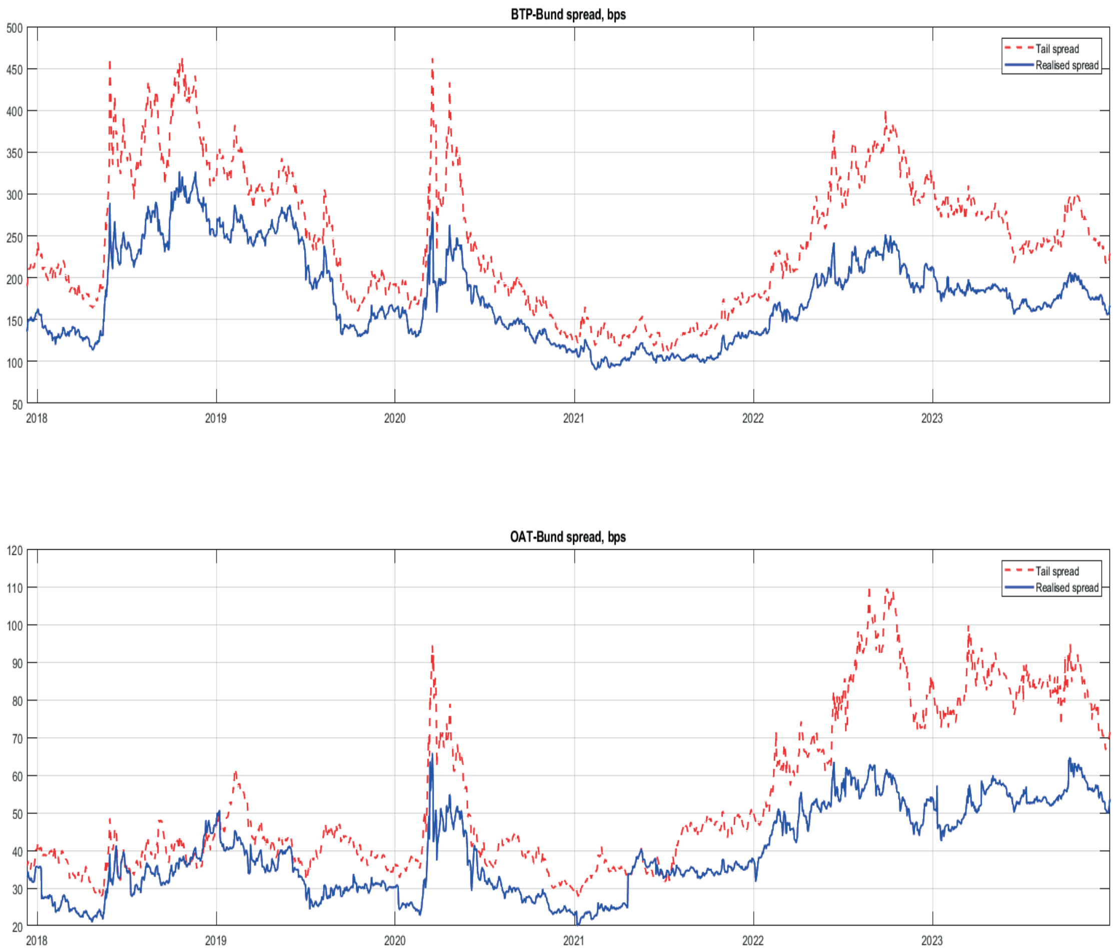

Note: We plot the realised spread at the time and the 95th percentile of the predicted spread in 45-days, both for Italy-Germany and France-Germany.

Source: ESM calculation based on Bloomberg LP data.

Figure 8: Average predictions' accuracy : log-likelihood comparison

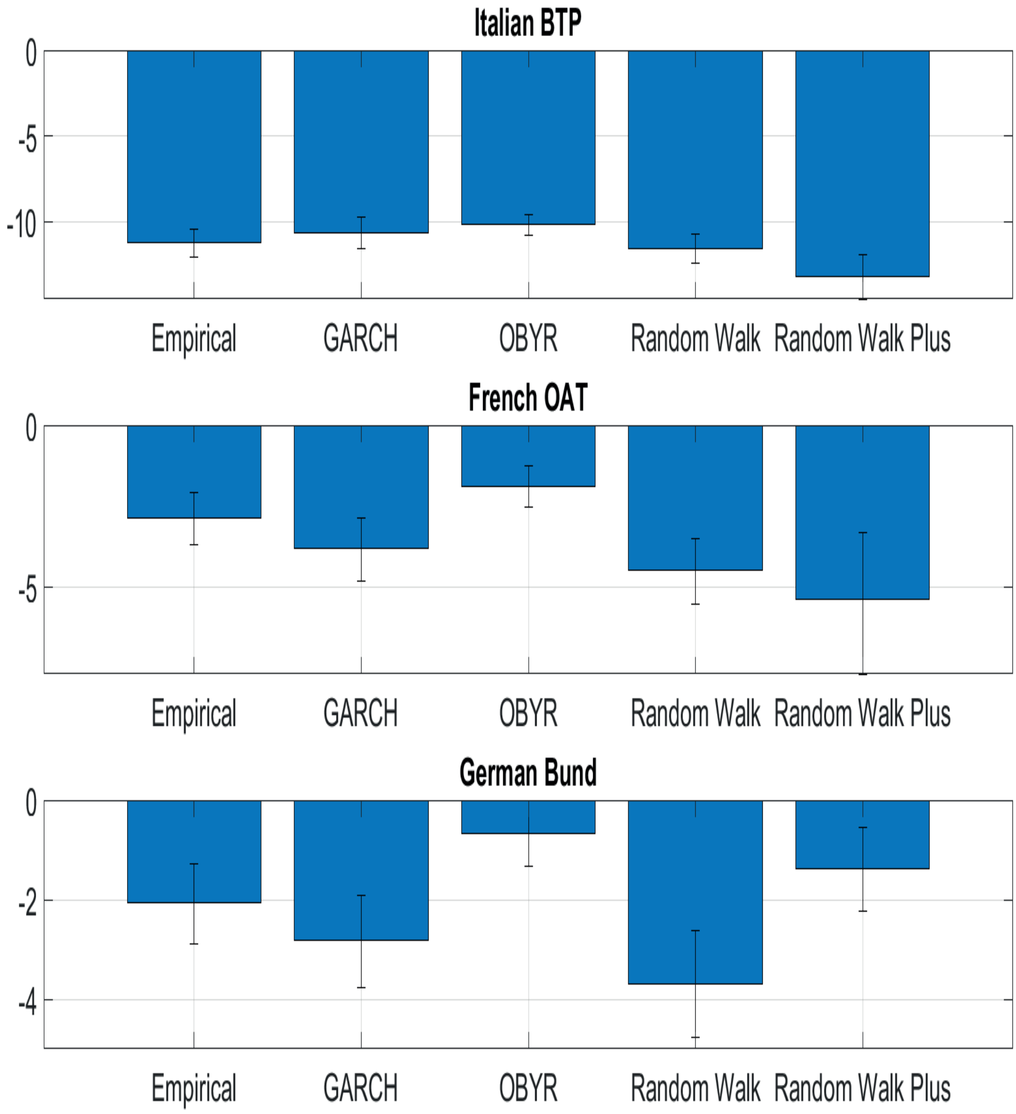

Note: The y axis refers to sum of log-likelihood divided by 100 of different univariate models in explaining 45-day-ahead yield changes. Confidence bands are derived using bootstrapping. Empirical model refers to the non-parametric distribution of observed rate changes since 2000. GARCH refers to a Generalised Autoregressive Conditional Heteroskedasticity model with one lag in both the autoregressive term and the moving average term. OBYR refers to a model based on option-implied bond yields. Random Walk is a normal-distributed model with mean equal to the forward yield and the standard deviation estimated from the standard deviation of rate changes since 2000. For Random Walk Plus, we replace the historical standard deviation with the option-implied one.

Source: ESM calculation based on Bloomberg LP data.

Figure 9: Tail predictions' accuracy: tick-loss comparison

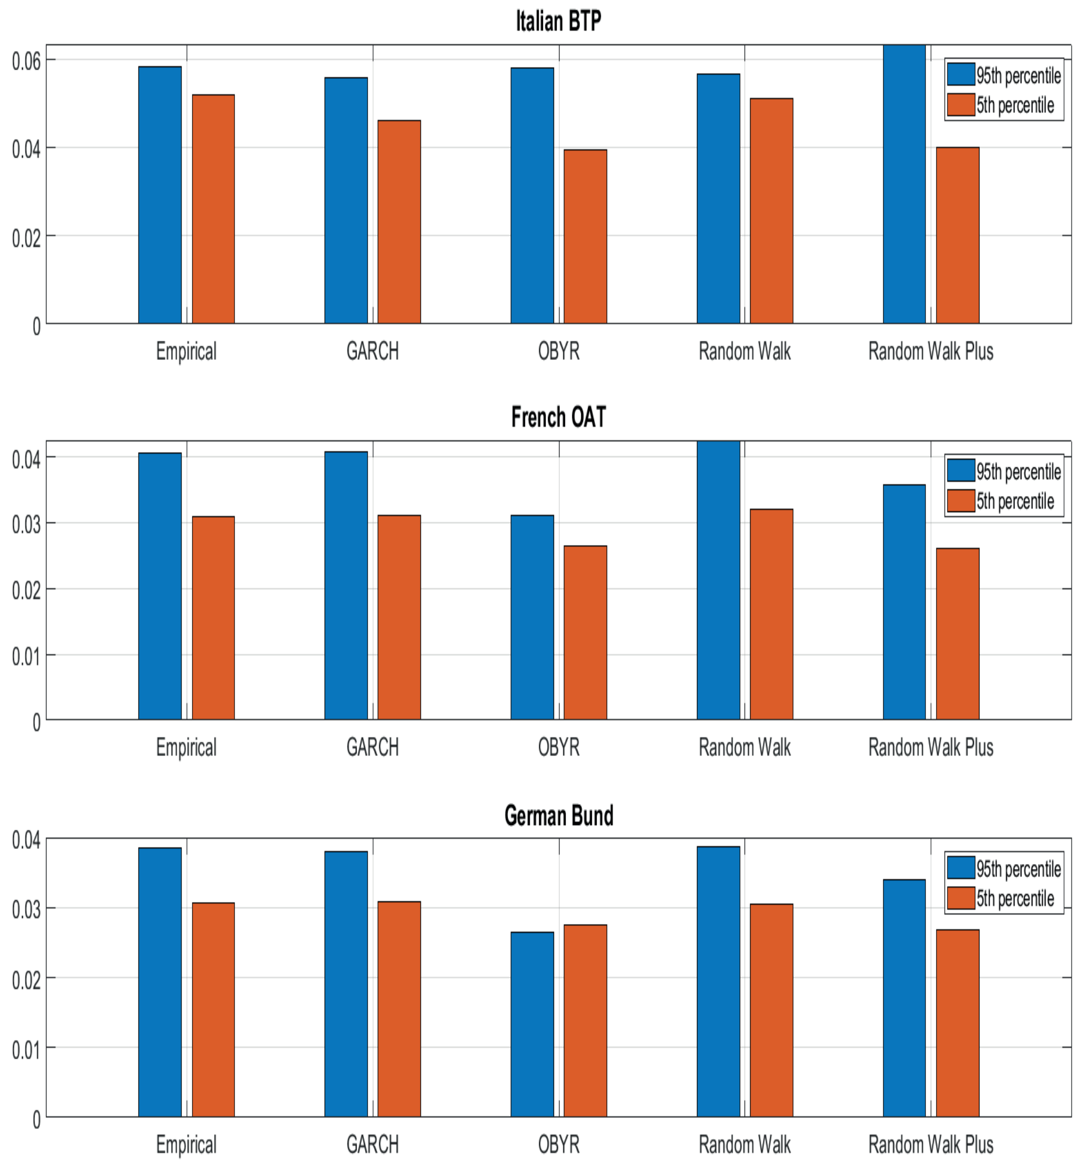

Note: See the note to Figure 8 for a description of models. Contrary to Figure 8 log-likelihood, which measures the model performance across the full distribution, Figure 9 tick-loss measures the models' performance in predicting 45-day-ahead tail events (large yield increases or decreases).

Source: ESM calculation based on Bloomberg LP data.

Figure 10: Overall log-likelihood comparison - Italian and French spreads

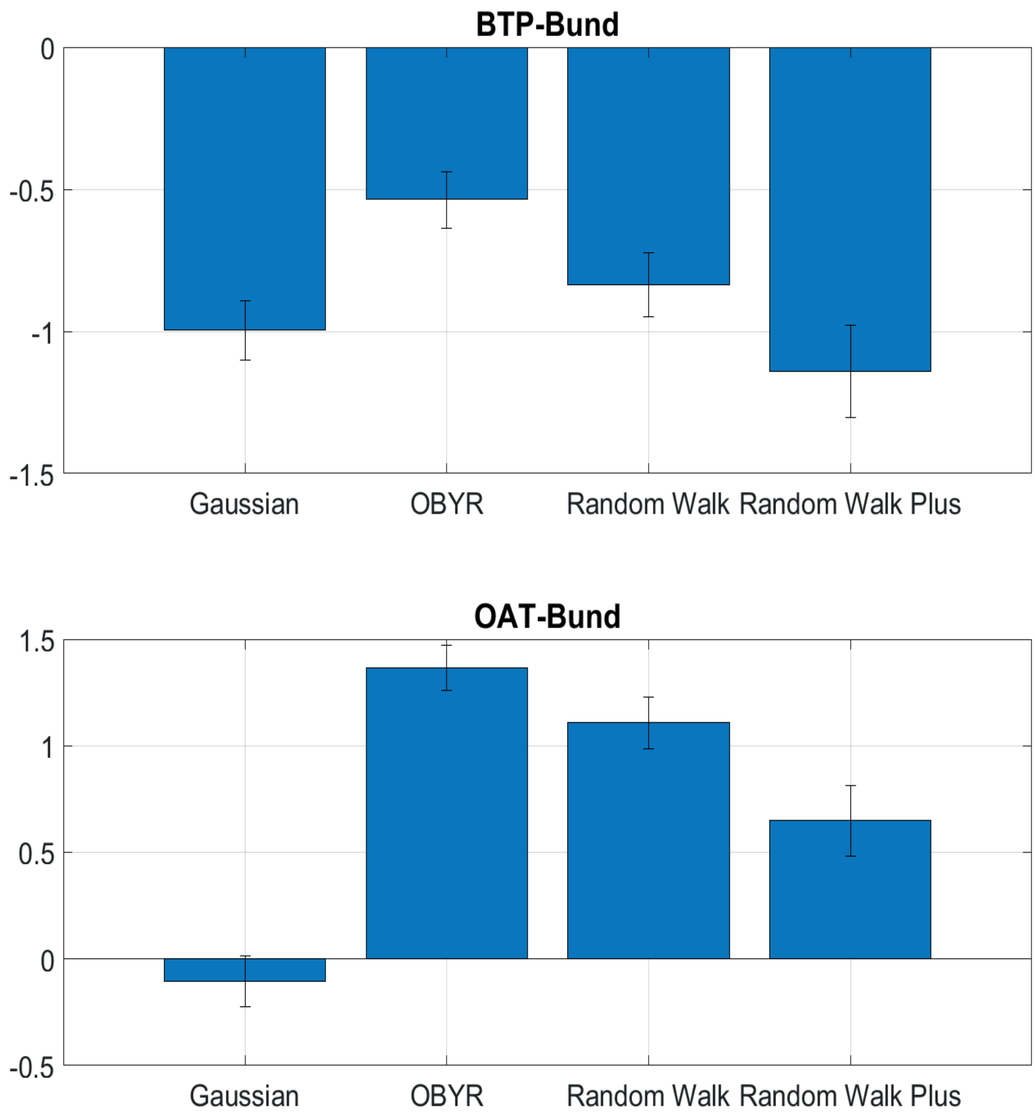

Note: The y axis refers to sum of log-likelihood divided by 1000 of different bivariate models in explaining 45-day-ahead spread changes. Confidence bands are derived using bootstrapping. Gaussian refers to the option-implied bond yield model, where yields of bonds from two countries are related using a Gaussian copula. OBYR also refers to the option-implied bond yield model, but here yields of bonds from two countries are related using a t-distribution copula. Random Walk and Random Walk Plus refer to the univariate models described in the note to Figure 8, related using a t-distribution copula.

Source: ESM calculation based on Bloomberg LP data.

Figure 11: Log-likelihood comparison over time - BTP-Bund spread

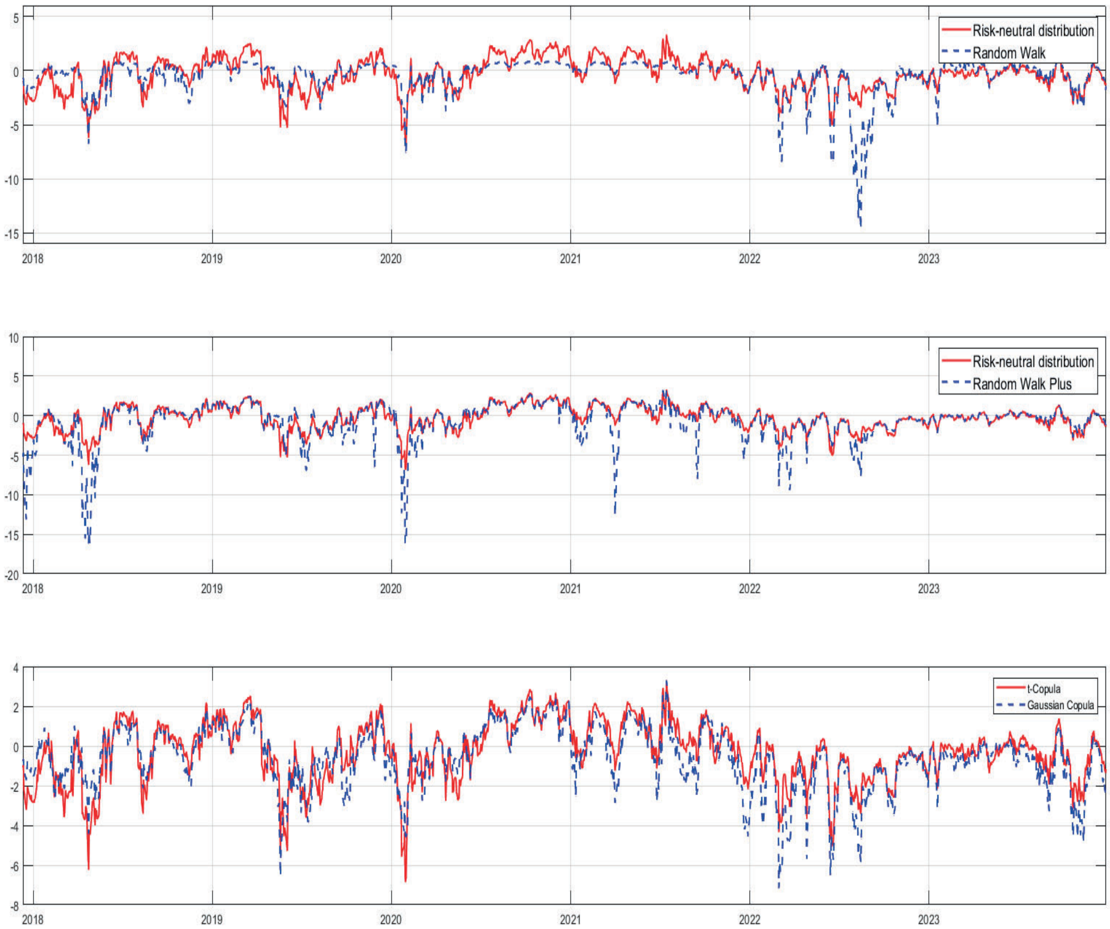

Note: The y axis refers to the log of the probability for different models of observing a realised bond-yield spread based on data from 45 days ago. A higher probability suggests a better predictive performance of the model. Risk-neutral distribution and t-Copula, both refers to the option-implied bond yield model, where yields of bonds from two countries are related using a t-distribution copula. Gaussian copula also refers to the option-implied bond yield model, but here yields are related using a Gaussian copula. Random Walk and Random Walk Plus refer to the univariate models described in the note to Figure 8, related using a t-distribution copula.
Source: ESM calculation based on Bloomberg LP data.

Figure 12: Log-likelihood comparison over time - OAT-Bund spread

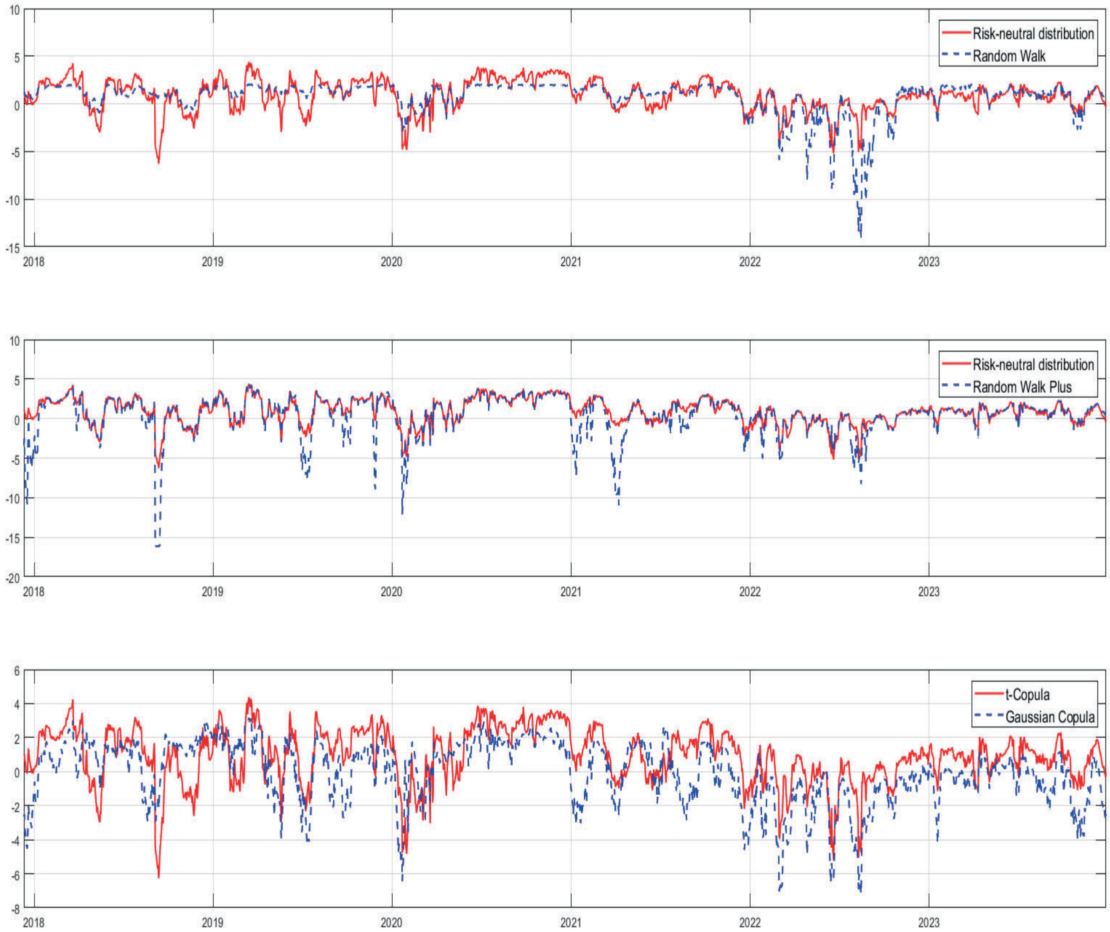

Note: The y axis refers to the log of the probability for different models of observing a realised bond-yield spread based on data from 45 days ago. A higher probability suggests a better predictive performance of the model. Risk-neutral distribution and t-Copula both refer to the option-implied bond yield model, where yields of bonds from two countries are related using a t-distribution copula. Gaussian copula also refers to the option-implied bond yield model, but here yields are related using a Gaussian copula. Random Walk and Random Walk Plus refer to the univariate models described in the note to Figure 8, related using a t-distribution copula.
Source: ESM calculation based on Bloomberg LP data.

A option-implied empirical PDF and CDF

The probability of options to be exercised depends on how prices and the realised volatility of their underlying asset's prices evolve over time. Therefore, option prices contain useful information about future bond yields and prices.

Specifically, European options allow to buy (call option, C) or sell (put option) an underlying asset at a predetermined future (maturity) date, T, and for a predetermined strike price, K. The opportunity of exercising the buy option (hence to sell the underlying asset at T and earn the difference) depends on the difference between strike and market prices.

Defining market prices at the maturity date as S_T, the above definition implies that the payoff is either $S_T - K$ or 0[11]. Furthermore, call prices are equal to the expected pay-off[12], which implies:

$$C = E[max(S_T - K, 0)]. \tag{A.1}$$

Equation A.1 suggests that call prices represent the probability-weighted average of all possible outcomes, or:

$$C = E[max(S_T - K, 0)] = \sum_{S_{T,i} > K} (S_{T,i} - K)\pi(S_{T,i}). \tag{A.2}$$

Denoting the probability distribution of the underlying price at T as π_{S_T} and expressing A.2 in continuous we obtain:

$$C = E[max(S_T - K, 0)] = \int_{S_T > K} (S_T - K)\pi_{S_T}(S_T)dS = \int_{S_T = K}^{\infty} (S_T - K)\pi_{S_T}(S_T)dS_T. \tag{A.3}$$

[11]Notice that when the market price is lower than the strike price (i.e., $(S_T - K) < 0$), the contract expires.
[12]Notice that we refer to call only in order to simplify the notation, but the same argument holds for put options.

In A.3, our object of interest is $\pi_{S_T}(S_T)dS_T$. In order to eliminate the other term, $(S_T - K)$ from the equation, we can differentiate (A.3) with respect to K, obtaining:

$$\frac{\partial C}{\partial K} = \int_{S_T=K}^{\infty} \frac{\partial}{\partial K}(S_T - K)\pi_{S_T}(S_T)dS_T = -\int_{S_T=K}^{\infty} \pi_{S_T}(S_T)dS_T = -\pi(S_T > K) \qquad \text{(A.4)}$$

What A.4 implies is that the sensitivity of option prices to variation in strikes depends on the probability of the underlying price being higher than the strike at T. Since the call option will be exercised only when this happens, if this probability is zero, the call price is insensitive to changes in the strike. On the other extreme, the price changes one to one with the strike when the probability is one. Moreover, from A.4 we can easily recover the complementary probability (i.e., the cumulative density function, or CDF) $\Pi_{S_T}(K) = \pi(S_T \leq K) = 1 - \pi(S_T > K)$. By substituting this definition in A.4 we obtain: $\frac{\partial C}{\partial K} = \Pi_{S_T}(K) - 1$

Finally, by taking the derivative of the CDF, we obtain the PDF (i.e., $\pi_{S_T}(K)$), that we are looking for:

$$\frac{\partial^2 C}{\partial^2 K} = \pi_{S_T}(K) \qquad \text{(A.5)}$$

Notice that A.5 is the curvature of the call price with respect to the strike K (i.e., it is the second derivative). By looking at this indicator, we can infer the probability of the underlying asset price being equal to that strike at the maturity.

B Practical implementation and definitions

In this section we describe how the (univariate and bivariate) risk-neutral distributions have been derived, and we provide definitions of the various metrics used in the paper. The software used is Matlab.

The univariate distributions for German, Italian, and French sovereign futures are constructed following these steps.

- Get option prices from a data provider. Specifically, for each option, we use information on the strike, the last tradeable date, the bid, the ask and the last (end of day) prices. Furthermore, for the underlying future, we also observe its expiry date, the forward rate of the cheapest to deliver (CTD) bond and the conventional CTD conversion factor.[13]

- Convert the option prices into implied volatility using the Matlab function "*impvbyblk*". This results in a typical convex series known as the "volatility smile," where the minimum is near the strike price where put and call options meet (i.e., at-the-money).

- Use the SABR (Stochastic Alpha, Beta, Rho) model to interpolate the volatility smile. The SABR model is calibrated by minimising the sum of squared differences between the implied volatility and the observed volatility smile. The calibration process involves optimising the alpha, rho, and nu parameters of the SABR model using Matlab's "*fmincon*" function. Note that SABR function is undefined at the point where strike and underlying price meet. This issue is mitigated by computing the function at a fine granularity (we apply a step-size of 0.01 strikes) and applying a cubic-spline interpolation at the discontinuity. Furthermore the interpolation mitigates an issue where prices in the tails are erratic, mainly because option prices on the market are fixed in discrete steps. Figure B.1 depicts the implied volatility extracted from market data only (see previous step) and the one obtained after the SABR interpolation. As the figure shows, the interpolated volatility series is smoother, much more granular, and it covers strikes which are well beyond the quoted sample.

- After obtaining the interpolated volatility, the Matlab function "*blkprice*" is used to convert it back into option prices. These reconstructed prices provide a smoother and more granular price series compared to the original market data.

- Derive the distribution from option prices by taking the first derivative of the reconstructed prices for the CDF and the second derivative for the PDF. This is

[13] Both the CTD conversion factor and the CTD forward rate are Bloomberg fields which are used to convert the distribution from strikes to yields. If using Bloomberg LP, active tickers can be retrieved using OPT_CHAIN and FUT_CHAIN.

achieved using Matlab's *diff* function. The PDF is re-scaled by the the step-size of strikes.

- After obtaining option prices and implied volatilities across different strikes, the next step is to convert these strike prices into yield space. The strike-based distribution is initially computed in terms of the option strike prices, but for bond futures, it is more intuitive to express the distribution in terms of yields. To do this, a transformation is applied using two key parameters: the forward rate of the cheapest-to-deliver (CTD) bond and the CTD conversion factor. The forward rate is the projected yield of the CTD bond at a future date, and the conversion factor adjusts for differences between the specific bond and the CTD bond. The transformation from strike prices to yields is expressed as:

$$\text{Yield} = \left(\frac{\text{Futures Price} - \text{Strike Price}}{\text{CTD Conversion Factor}} \right) + \text{CTD Forward Rate}$$

- When constructing the risk-neutral distribution, irregularities can occur at the extremes, where the probability density function (PDF) may turn negative or exhibit erratic behavior. To address this, we smooth these regions by first filtering for positive PDF values and then by approximating the tails with an exponential distribution. The process begins by identifying the left tail region where the PDF becomes negative. Specifically, the left tail cutoff is located by finding the largest strike where the PDF becomes negative, for strikes located below the mode. If no such point exists, the leftmost strike is assumed to be the tail's start. Once the left tail cutoff is determined, we set the PDF as 0 for all values to the left of the tail cutoff. The right tail is treated symmetrically. Furthermore we rescale the distribution to sum to 1. Next, we model the tails using an exponential distribution. The main objective is to ensure the distribution is unbounded and yields strictly positive probabilities for any strike. To that end, the normalised PDF is converted into a logistic form suitable for exponential fitting:

$$\text{logform} = -\log\left(\frac{1}{\text{CDF}} - 1\right)$$

This transformation essentially maps the smoothed CDF values (which range between 0 and 1) onto a real-number space, where the tails can be more easily modeled by an exponential distribution. The transformation ensures that values near 0 and 1 map to negative or positive infinity. After this transformation, the 'logform' is used to define the shape of the tail regions. Next, a small threshold $\epsilon = 0.005$ is introduced. This threshold is used to determine when the exponential fitting should start in the tail. The left tail cutoff K_{left} is identified by finding the largest strike K where the following conditions are jointly met: 1. The strike is located below the mode. 2. The second difference of the 'logform' is negative. 3. The normalized CDF value, is smaller than the threshold ϵ, signaling that the PDF has become sufficiently small to warrant approximation with an exponential distribution. If a valid cutoff point is found, the exponential tail is fitted to ensure continuity at this point. The exponential fit for the left tail is extrapolated with a constant slope, such that it matches the slope and value of the transformed 'logform' at K_{left}. Finally the 'logform' computation is inverted to back out the underlying CDF. As a result, due to the linear extrapolation, situations are avoided where the 'logform' reaches negative infinity, and consequently the underlying PDF is unbounded. In a symmetric fashion the same logic also applies to the right tail.

- Risk-neutral distributions are typically computed at different maturities, corresponding to the expiration dates of options on the underlying futures. Figure B.2 offers an example of the distribution for Bund yields as of the 12/10/2022 across two horizons. As it shows, the two distributions have a different shape. Not surprisingly, the one with the closer horizon shows a substantially lower uncertainty (i.e., the density based on options expiring on 21/10/2022 is "taller" and "narrower" than the other one). However, to compare distributions across different time horizons, it is useful to interpolate these distributions to a common horizon (e.g., 1 month, 3 months, etc.). This step involves creating a linear interpolation of probabilities across percentiles between the distributions computed at different maturities, allowing for the extraction of constant-horizon distributions. By doing so, one can consistently compare the market's expectations for bond yields over various time frames, even when the available option data has different expiry dates. The interpolation ensures that the resulting distributions are aligned on a single time axis, making it easier to analyse the evolution of risk-neutral expectations over time. This is particularly important for understanding how market sentiment and expectations shift as bonds approach maturity or as macroeconomic conditions change.

Next, **bivariate risk-neutral distributions** of spread changes can be constructed based on (univariate) risk-neutral marginals. Since these densities are not independent, constructing the *joint* risk-neutral distribution requires accounting for the correlation structure characterising them. To this end, we exploit the theory of copulas (Sklar, 1959), and follow these steps.

- To compute copulas, we first approximate the original distributions of realised bond yield changes for both the German and French or Italian markets. We do this by estimating the empirical cumulative distribution functions (CDF) for each dataset using kernel density smoothing. Specifically, the differences in yields are calculated by subtracting the 45-days lagged yield values from the contemporaneous yields. The empirical CDFs are then calculated using the Matlab function "*ksdensity*". Using these CDFs, we fit a t-copula model to the data, estimating both the correlation and the degrees of freedom through approximate maximum likelihood estimation applying the Matlab function "*copulafit*".
- After fitting the copula, for some functions like assessing the likelihood of observing the

realisation of a specific event, we can derive the result analytically in a closed form. For other functions such as deriving the likelihood of spread increases we simulate random draws.

- For deriving the likelihood of observing the realisation of an event in the bivariate distribution, we first need to derive the percentile of each event occurring on the respective univariate distribution. Next we apply the Matlab function "*copulapdf*" to back out the probability. Finally we need to scale the result appropriately for the further analysis of spreads. To achieve this, we multiply the unscaled probability with a scaling factor.

$$\text{scaling\_factor}_c = \frac{(\Delta\text{PCT}_{\text{DE}} \cdot \Delta\text{PCT}_c)}{(\Delta\text{YLD}_{\text{DE}} \cdot \Delta\text{YLD}_c)}$$

  Here the variables refer to the univariate distributions that are sampled at a discrete scale. $\Delta\text{PCT}_{\text{DE}}$ refers to the German univariate distribution and is the difference in the two percentiles closest to the percentile of the event in question. $\Delta\text{YLD}_{\text{DE}}$ is the difference in yields corresponding to the two percentiles. The c subscript refers to the French or Italian univariate distribution. In other words we are selecting points in the CDF for both univariate distributions that correspond to the specific quantiles we are interested in. These points are used to calculate the "area" under the copula curve, which represents the joint probability mass for a given region of the copula distribution. Similarly, we compute the area under the curve for the yields themselves. By dividing the copula area by the yield area, we obtain a scaling factor that adjusts the copula distribution to match the observed yield data.

- As described above we also use the fitted t-copula to simulate random samples, generating synthetic realisations of the joint distribution of yield changes. We apply the Matlab function "*copularnd*" to generate a random sample of 1 million observations. These realisations allow us numerically to capture the dependence structure between the markets.

- Simulating random samples from the *joint* distribution allows to derive the risk-neutral probability of a spread increase being larger than a given threshold. We start with the random samples generated from the fitted t-copula described above. For each sample we compute the change in spread as the difference in yields at the 45-day horizon minus the difference in spread at t_0. Our measure of interest finally refers to the share of samples that has a spread change larger than a specific threshold. We compute the measure for thresholds of 15 and 100 basis points.

- Another measure obtained as a result of random sampling from the fitted t-copula are the "spread tails". Specifically, we simulate distributions of (Italian and French) spread changes, and define this measure as the 95th percentile of the sample distribution density.

- To compute the confidence bands for likelihood estimates, we apply the Matlab function "*bootstrp*" for 10,000 draws, and taking in each draw the mean of the drawn log-likelihood series. From the resulting distribution we report the median as well as the 5th and 95th percentile.

Figure B.1: Example of SABR interpolation

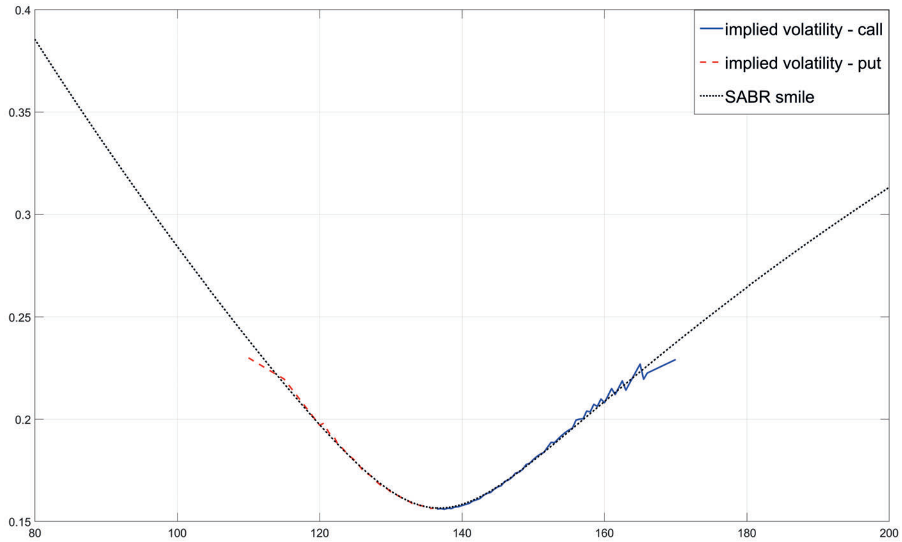

Source: ESM based on Bloomberg LP

Figure B.2: Example of option-implied densities for two horizons

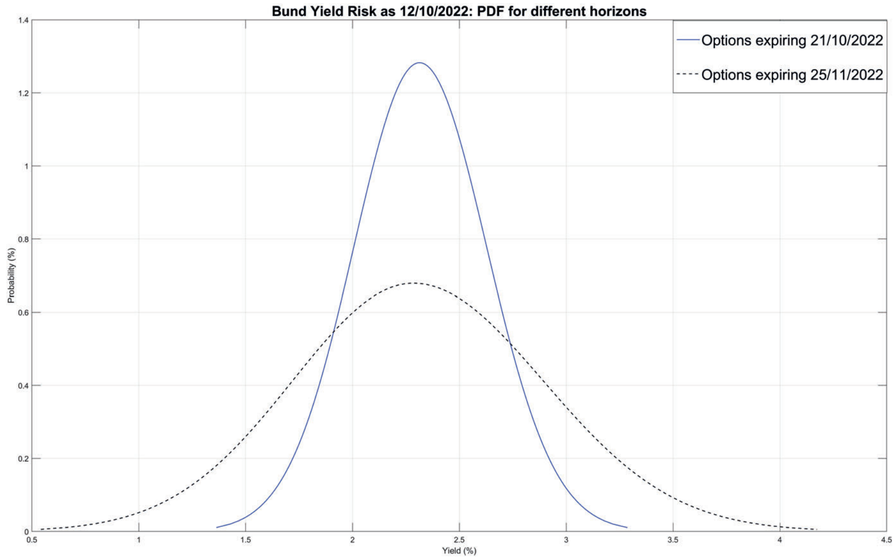

Source: ESM based on Bloomberg LP

European Stability Mechanism

6a Circuit de la Foire Internationale
L-1347 Luxembourg

Tel: +352 260 292 0

www.esm.europa.eu

info@esm.europa.eu

Publications Office
of the European Union