Working Paper Series

No 65 / January 2018

Revised June 2019

Sovereign bond-backed securities: a VAR-for-VaR and Marginal Expected Shortfall assessment

by
Maite De Sola Perea
Peter G. Dunne
Martin Puhl
Thomas Reininger

ESRB
European Systemic Risk Board
European System of Financial Supervision

Abstract

The risk reducing benefits of the sovereign bond-backed security (SBBS) proposal of Brunnermeier et al. (2016) have been assessed in terms of the likely losses that different kinds of holders would suffer under simulated default scenarios. However, the effects of mark-to-market losses that may occur when there is rising uncertainty about defaults, or when self-fulfilling destablising dynamics are prevalent, have not yet been examined. We apply the "VAR-for-VaR" method of Manganelli et al. (2015) and the Marginal Expected Shortfall (MES) approach of Brownlees and Engle (2012, 2017) to estimated yields of SBBS to assess how *ex ante* exposures and marginal contributions to systemic risk are likely to play-out for different SBBS tranches under various securitisation structures. We compare these with exposures/MES of single sovereigns and a diversified portfolio of sovereigns. We find that the senior SBBS has extremely low *ex ante* tail risk and that, like the low-risk sovereigns, it acts as a hedge against extreme market-wide yield movements. The mezzanine SBBS has tail risk exposure similar to that of Italian and Spanish bonds. Yields on SBBS appear to be adequate compensation for their risks when compared with single sovereigns or a diversified portfolio.

JEL classification: E43, E44, E52, E53, G12, G14.

Keywords: Safe Assets; Sovereign Bonds; Value-at-Risk; Spillover; CAViaR; Co-Dependence.

1 Introduction

This paper examines tail risk exposure and safe-haven properties of estimated yields on sovereign bond-backed securities (SBBS) proposed by Brunnermeier et al. (2016). The proposed securitisation involves the replacement of a proportion of the supplies of individual European sovereign bonds with senior, mezzanine and junior tranches backed by the purchased sovereigns. The junior and mezzanine components would, in turn, be exposed to agreed levels of losses from defaults on the underlying securities (perhaps 10% and 20% respectively) before senior bond holders representing the remaining 70% of the securitisation would become liable. Simulations by Brunnermeier et al. (2016, 2017) have shown that the senior bond-holders in a securitisation with such proportions are *ex post* as unlikely to experience losses as any existing individual low-risk sovereign while the mezzanine and junior tranches (representing 10% or more) should expect to experience losses comparable to those of euro area sovereigns with intermediate and higher credit risks respectively.

However, the existing risk assessments are based on simulated default outcomes - based on a variety of assumed probabilities of defaults, default correlations and expected losses given defaults - rather than on the losses that arise from fluctuations in the market valuation of the securities. This leaves a gap in our understanding of the effects of 'flights-to-safety' and any other type of market panic that could arise from perceptions and doubts about risk exposures of SBBS in very extreme circumstances (and the real feedback effects that these produce). There remains a concern that, when expected losses on mezzanine or junior bonds are at very high levels, the senior SBBS may suffer much larger mark-to-market valuation risks than existing safe assets. Such fears can be assessed if SBBS yields can be estimated. In this case the risks and rewards that would have been faced when investing in the senior tranche could be compared with those faced when investing in a well established safe haven.

A related concern is that there may be insufficient interest in holding the junior bond given that its yield (when compared with single risky sovereigns) may be too low to compensate for its high level of embedded leverage. If yields on SBBS tranches can be estimated, then a supplementary analysis can be designed to address whether rewards appear high for the measurable risk exposures. This aids in categorising the junior tranche according to its reward for risk and is merely a first step in assessing likely demand. We leave a more exhaustive assessment for future work. Lastly, it is of interest to compare the risk attributes of the proposed securities with those of a diversified portfolio of sovereign bonds to assess the *ex ante* benefits that arise from tranching beyond pure diversification effects.[1] All of these issues can be addressed if realistic estimates of SBBS yields can be derived.

To derive probable yields on the SBBS components, for a variety of securitisation structures over roughly a 17-year history, we implement a pricing tool that employs a correlated multivariate Monte Carlo approach to the reallocation of historical yield spreads to SBBS components based on a static copula as described in Schönbucher (2003).[2] The present analysis uses the estimated yields to ascertain whether SBBS have Value-at-Risk (VaR) and Marginal Expected Shortfall (MES) characteristics that would make them relatively more attractive to investors when compared with individual sovereigns or a weighted pool of euro area sovereigns.

The MES and VAR-for-VaR metrics are measures of *ex ante* investment outcomes and the tail-threshold for such outcomes conditional on extreme market-wide/systemic declines or VaR movements in another asset.[3] The VAR-for-VaR reveals the expected threshold that, with some chosen likelihood, will be breached by future outcomes as a function of the lagged expected quantile and lagged (absolute) outcomes of each asset included in the VAR (see, Engle and Manganelli (2004) and Manganelli et al. (2015)). Hence the VaR is a measure of an investor's exposure to time-varying extreme tail-risks when holding an asset and we can examine how the VaR of one asset responds to shocks in the returns of another. The MES, in contrast, reveals how one asset performs in terms of expected return when the entire market is likely to be experiencing a tail event (we follow the MES measurement procedure outlined in Brownlees and Engle (2012) where MES is an important ingredient of the SRISK measure they propose). MES measurement complements VAR-for-VaR analysis since it allows the expected response to systemic shocks (i.e. MES) to be compared with contemporaneous responses assumed in a VAR-for-VaR impulse-response analysis.[4]

[1] The effects of banks diversifying into the senior tranche of a sovereign bond-backed securitisation, rather than merely diversifying across existing sovereigns, has recently been analysed by Alogoskoufis and Langfield (2018).

[2] This Monte Carlo estimation method is explored in more detail in the report of the ESRB High-Level Task Force on Safe Assets (2018).

[3] Related risk measures include Systemic Expected Shortfall (SES) by Acharya et al. (2012) and ΔCo-VaR as proposed by Adrian and Brunnermeier (2016).

[4] In assessing VAR-for-VaR impulse-responses we consider the approach outlined in Manganelli et al. (2015) as likely to produce an upper bound on how systemic shocks affect the conditional VaR of safe haven

MES has some advantages over VaR measures when allocation of the contributions to systemic risk are required. MES is, by definition, the marginal contribution of a change in the weight of an individual asset to 'systemic' Expected Shortfall (i.e. the Expected Shortfall of the whole market). The MES of market components can be summed to produce the market Expected Shortfall and this is not the case for VaR measures.[5] In a more general sense, MES is useful in identifying "safe assets" as those that contribute little or nothing to systemic risk. From this definition it becomes clear that the overall supply of safe assets is an important determinant of the size of market-wide/systemic expected shortfalls and their spillovers. VAR-for-VaR analysis also has some advantages over the MES approach. These are related to the fact that VAR-for-VaR quantile estimates are less reliant on distributional assumptions than is true for many other risk measures, including those relying on GARCH modelling.

We use MES and VaR (and conditional volatility measures) to ascertain whether investors are adequately rewarded for the actual and expected risk exposures they would face when holding the different tranches of a securitisation. A commonly used measure for comparisons of the rewards from holding different assets is the Sharpe ratio (Sharpe (1966, 1994)) which, in an *ex post* sense, is the average excess holding period return relative to historical standard deviation of such returns. However, it is also valid to examine expectations of excess returns relative to conditional expected standard deviation of returns or relative to conditional volatility, VaR, and MES-based risk exposure measures. We use GARCH-implied conditional standard deviation, VAR-for-VaR quantile projections and MES as risk measures to assess whether projections of holding period excess returns provide adequate rewards for risk exposures (i.e. to derive dynamic Sharpe and quasi-Sharpe ratios).[6] All

assets. This reflects the use of simple covariance in VAR-for-VaR analysis rather than covariance conditioned on the presence of systemic risks.

[5] VAR-for-VaR is a potential substitute for a MES analysis since it could be modified to relate median outcomes to market-wide tail events. But VAR-for-VaR will only have similar properties to MES if (i), movements in the median and MES are similar, (ii) the market-wide VaR is similar to market-wide Expected Shortfall and, (iii), if negative shocks to the market VaR are associated with an increased desire on the part of investors to hold assets perceived as safe (i.e., coincide with an increased willingness to pay a premium for such assets). The first condition is unlikely if tail behaviour is driven by non-linear feedback effects. Condition (ii) ignores the fact that market VaR cannot be interpreted as the sum of the VaRs of market components. Following from (ii) condition (iii) seems unlikely to hold.

[6] In addition to our analysis of the adequacy of holding period excess returns, we also examine yield-to-maturity relative to conditional volatility (and relative to absolute Value-at-Risk) of yield-to-maturity movements.

reward-for-risk measures are compared across SBBS tranches, a euro area GDP-weighted portfolio of sovereign bonds and individual sovereigns.

Our findings imply that the senior SBBS as proposed by Brunnermeier et al. (2017) has low levels of tail risk exposure (as low as, and often lower than, the tail risks of the lowest-risk euro area sovereign). We acknowledge that this result pertains mainly to the case of a senior tranche that remains approximately at 70% or lower as a proportion of the securitisation (increasing the size of the senior tranche eventually leads to higher credit risk and a less favourable comparison with existing low-risk sovereigns). MES analysis shows that senior bonds also retain similar hedging (or safe haven) properties to the lowest-risk sovereign - protecting against exposure to the spillover of losses from market-wide tail events.

The mezzanine tranche ranks as similarly exposed to tail risk as Italian and Spanish bonds. However, junior assets are very similar in terms of tail risk exposure to Irish and Portuguese bonds but (at 10% of the securitisation) they are less exposed to tail risks than Greek bonds. Reducing the size of the junior tranche proportion would eventually produce a credit risk premium in the junior tranche that is at least as high as that of the riskiest sovereign. This is simply because, with a very small junior tranche, expected default losses on **any** of the high-risk sovereigns would generate significant losses for junior tranche holders. With a junior tranche of around 10% there is some diversification of sovereign default risk. Only a high probability of multiple defaults would cause expected losses on the junior tranche to be as high as those for any individual high-risk sovereign. In our analysis we do not consider the case of a securitisation with a junior tranche proportion below 10%.

The reward for risk analysis implies that the junior tranche performs well relative to ex ante risks (and especially so in the post-sovereign debt crisis period). Some isolated large ex post losses that occurred at the height of the crisis would need to be considered by investors as indicative of potential future extreme tail risks. Notwithstanding the fact that such risks are to some extent diversified for junior SBBS investors, the typical investor in junior SBBS must be capable of withstanding extreme losses with non-negligible probability. The existence of such investors, and what they require to incentivise their demand, deserves further consideration going beyond the focus of the current analysis.

The note is arranged as follows. In the next section we briefly outline the data and methodologies employed. In section 3 we discuss the results from application of the VAR-for-VaR and MES applications. This is followed by concluding comments.

2 Data & Methodology

2.1 Data

The following senior:mezzanine:junior tranche proportions were used in the VAR-for-VaR and MES analysis; 70:30, 80:20, 90:10, 70:20:10, 80:10:10. We only discuss a subset of the results (i.e those for 70:20:10 and 70:30 tranche structures).[7] The VAR-for-VaR and MES for the estimated yields is compared with that of individual bonds of 11 euro area countries and with a euro area GDP-weighted portfolio. The results presented apply to assets with 10 years to maturity. Our sample runs from the beginning of January 2003 to the end of October 2016. Daily data for individual sovereign bond yields were sourced from Datastream. Yields for the securities backed by these sovereign bonds were estimated using the methodology discussed below and these were then treated in the same way as other individual sovereign bonds in the VAR-for-VaR and MES analysis. The negative of the daily yield change in basis points is used as the model variable.[8]

2.2 Simulating Yields

We estimate the yields of SBBS tranches with a multivariate Monte Carlo simulation that is based on a static copula approach as described in Schönbucher (2003). We create a joint distribution function of country-specific random variables to derive scenarios in which individual countries default. The joint distribution function is created in the first instance with a Gaussian copula and thereafter transformed into country-specific uniform variables between 0 and 1 that are correlated. Depending on the actual historical default probability (PD) of the respective country, a default is assumed to have occurred if the value of the random scenario drawn (i.e. one scenario within the simulation run per day and country)

[7] Results for other cases, including other maturity categories and other default correlations, are available from the authors.

[8] A related analysis of the price returns is available from the authors but this produces very similar results. In this case the bond price is approximated as (100 - 10 (yield - coupon)).

exceeds the threshold of (1-PD). For example random values above 0.96 lead to a default for a country with a PD of 4% or above.

In this way, the simulated scenarios define which of the sovereign bonds in the SBBS structure default and allows the calculation of the associated loss. The losses are assigned to the different tranches according to the predefined tranche structure and allows the construction of tranche-specific expected loss (EL) distributions. The overall risk premium (yield exceeding the risk free rate) of the bond portfolio is then allocated to the tranches according to their EL proportions. Consistency checks ensure that the weighted average yield of the tranches is identical to the yield of the underlying bond portfolio.

The estimation is conducted with historical market data. To represent sovereign-specific default probabilities we would normally use the bond yield minus a risk free rate (assuming the yield premium to be mainly explained by credit risk). We actually use the country-specific bond yield minus the lowest euro area sovereign yield each day (this is a close approximation of the yield minus EONIA but the latter produces some negative premiums which is not suitable for our purposes). In order not to understate the yield and resulting expected loss of the senior bond (in comparison to the German bund) and not to overstate the attractiveness of junior bonds, we decided to use a LGD of 100%.[9]

Datastream benchmark government indices are used for the yield time series of the respective governments. For the country correlation in the default scenario generator, we set a constant value of 0.60. This means that the random variables determining default events are noticeably interdependent.[10] 30,000 scenarios per day and country are used for the simulation. For the generation of (uniform) correlated default scenarios, we used the R package MASS. The simulation is based on yield data for 2, 5 and 10 year government bonds of Austria, Belgium, Germany, Spain, Finland, France, Greece, Ireland, Italy, the Netherlands and Portugal. Following a weighting scheme based on GDP (average between

[9]When the default simulator is set to deliver correlated uniform random variables (specifically, with this correlation set equal to 0.6), the estimated yield premium is not fully independent of the subdivision of EL into PD and LGD. For LGD values above 60% the yields start to diverge slightly from where they would be with a low LGD. For instance, on the basis of a 70/30 tranche structure and a default correlation assumption of 0.6, the junior bond yield for LGD values of 30%, 60% and 100% are 2.63%, 2.61% and 2.52%, respectively, as of 31/10/2016, while the senior bond yields are 0.08%, 0.09% and 0.13%. Using a 100% LGD gives the most conservative outcome.

[10]Note that the different PDs of the respective countries also influences the occurrence of default events. Even in the case of 100% correlation, the random variable would have to surpass a different PD threshold for e.g. Germany than for Portugal.

2006-2015), those countries cover 97.5% of the volume of the SBBS.

2.3 VAR-for-VaR

The VAR-for-VaR methodology of Manganelli et al. (2015) is essentially a vector autoregression applied to quantile relations in which autoregressive cross-effects are permitted. This extends the Conditional Autoregressive Value at Risk (CAViaR) model of Engle and Manganelli (2004) to a multivariate context. The most obvious benefit from estimating a CAViaR or VAR-for-VaR model is that it produces credible time varying VaR estimates. These can differ significantly from VaR estimates derived indirectly from GARCH estimates of variance. GARCH-implied VaRs rely on an assumed distribution for the underlying observations (usually Gaussian) and this constrains quantiles behaviour. The CAViaR and VAR-for-VaR approaches use probability-based weighted deviations of observations around the VaR estimates to achieve the best fitting process and this is not constrained by distributional assumptions.

The VAR-for-VaR approach fits a parameterised smoothed step-function of predetermined variables to choose a quantile value and a categorisation of returns that produces the appropriate number of exceedences of the quantile using an optimisation strategy based on Koenker and Basset (1978). For a pair of asset returns, r_{mt}, r_{at}, and associated quantiles, q_{mt}, q_{at}, the period 't' quantiles are related to the most recently estimated quantiles in 't-1' and lagged absolute returns as follows;

$$
\begin{aligned}
q_{mt} &= c_1 + a_{11}|r_{mt-1}| + a_{12}|r_{at-1}| + b_{11}q_{mt-1} + b_{12}q_{at-1} \\
q_{at} &= c_2 + a_{21}|r_{mt-1}| + a_{22}|r_{at-1}| + b_{21}q_{mt-1} + b_{22}q_{at-1}
\end{aligned}
\tag{1}
$$

Although different quantile levels can be jointly modelled, we estimate the case where a common quantile (e.g. 1%) is chosen for each distribution. Significance of parameters on cross-terms is used as evidence of cross-effects.

Manganelli et al. (2015) suggest a framework for assessing the response of conditional VaR to shocks in the underlying returns. In this framework two uncorrelated structural shocks to the returns outcomes underpin each bivariate VAR-for-VaR. Only the first structural shock contemporaneously impacts on the first variable (say the market or the junior SBBS) while both structural shocks may affect the second variable. This implies a Cholesky decomposition of the covariance matrix of returns such that for L lower triangular, LL' gives the covariance matrix.[11] Since one of the off-diagonal elements of L is zero the observed outcome in one of the assets is related to only one of the structural shocks while the other is a response to a combination of shocks. If L is standardised so that it produces returns with the observed covariance from unit variance uncorrelated shocks $\{\epsilon_1, \epsilon_2\}$, then setting $\epsilon_1 = 1$ and $\epsilon_2 = 0$ and following the VaR process in Equation 1, the contemporaneous response to a one standard deviation shock to the first factor is obtained as (assuming that all terms in L are positive);

$$\begin{bmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{bmatrix} \begin{bmatrix} l_{1,1} & 0 \\ l_{2,1} & l_{2,2} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \tag{2}$$

where $l_{1,1} = \sigma_m$ is the standard deviation of first factor/variable and $l_{1,2}$ describes how a unit shock to the first factor contemporaneously affects the second variable (i.e., $l_{1,2} = Cov(r_m r_a)/\sigma_m$).

The Cholesky decomposition has some appeal in the case of interactions between the tranches of a securitisation because there is a natural causal ordering running from junior to more senior tranches. The shock that contemporaneously affects the more senior tranche and not the junior would then represent the effects of defaults that are expected to occur after the junior tranche has been fully devalued. It may also be valid to motivate a causal ordering for the case of interactions between the conditional VaRs of individual sovereigns. Here the sovereign with weakest fiscal circumstance, or with the weaker banking sector, would be considered the causal source. In this case it is more difficult to motivate a complete lack of contemporaneous feedback from the stronger sovereign to the weaker one but it is likely to be a reasonable approximation.

A slight weakness of the impulse-response set-up proposed by Manganelli et al. (2015) is that it is applied to the unconditional covariance of asset returns. The covariance that matters for VaR cross-effects is arguably the return conditional on a systemic event. This implies that the VAR-for-VaR impulse-response analysis could overstate how shocks to returns transmit to declines in conditional VaRs (particularly for assets considered as safe havens). In most cases we find that the VAR-for-VaR impulse-response analysis produces

[11] See, Cholesky (1910), which is the 2005 publication of the 1910 manuscript.

a subdued (and usually insignificant) response in the low-risk sovereigns (or more senior SBBS tranches) but, it should be noted that, this probably understates the degree of protection that is afforded to the safe assert holder. In these cases Marginal Expected Shortfall provides a better indication of the (contemporaneous) benefits from holding the safe asset.

2.4 Marginal Expected Shortfall

Marginal Expected Shortfall is the expected return on a specific asset where the expectation takes account of (i) the joint distribution of the standardised returns of the specific asset and that of the market portfolio, and (ii), the expectation of idiosyncratic and market returns conditional on a market-wide tail event. The correlation between returns in the joint density of the observations therefore plays a vital role in determining the marginal expected shortfall but tail correlation is also important. For example, when the individual asset is a safe haven, it is possible that the marginal expected shortfall would be a highly positive outcome.

More formally, let the expected value of the standardised market return in the tail be $E_{t-1}(\epsilon_{m,t}|\epsilon_{m,t} < C/\sigma_{m,t})$ and the conditional expected asset-specific return be denoted $E_{t-1}(\xi_{a,t}|\epsilon_{m,t} < C/\sigma_{m,t})$ where C is set to the market non-parametric 1% VaR and we assume that $\epsilon_{m,t}$ and $\xi_{a,t}$ are uncorrelated. Crucially, although $\epsilon_{m,t}$ and $\xi_{a,t}$ are uncorrelated the expectation of $\xi_{a,t}$ conditional on $\epsilon_{m,t}$ being an extreme negative value is not necessarily zero. MES for asset a is then measured as follows (where ρ is the correlation between asset-specific and market returns);

$$\begin{aligned} MES_{a,t-1}(C) &= E_{t-1}(r_{a,t}|r_{m,t} < C) \\ &= \sigma_{a,t}E_{t-1}(\rho_t\epsilon_{m,t} + \sqrt{1-\rho_t^2}\xi_{a,t}|\epsilon_{m,t} < C/\sigma_{m,t}) \\ &= \sigma_{a,t}\rho_t E_{t-1}(\epsilon_{m,t}|\epsilon_{m,t} < C/\sigma_{m,t}) + \\ &\sigma_{a,t}\sqrt{1-\rho_t^2}E_{t-1}(\xi_{a,t}|\epsilon_{m,t} < C/\sigma_{m,t}). \end{aligned} \tag{3}$$

Following Scaillet (2005) we measure the conditional tail expectations of the components using a kernel estimation method as follows:

$$E_{t-1}(\epsilon_{m,t}|\epsilon_{m,t} < C/\sigma_{m,t}) = \frac{\sum\limits_{t=1}^{T} \epsilon_{m,t}\Phi(\frac{c-\epsilon_{m,t}}{h})}{\sum\limits_{t=1}^{T} \Phi(\frac{c-\epsilon_{m,t}}{h})};$$

$$E_{t-1}(\xi_{a,t}|\epsilon_{m,t} < C/\sigma_{m,t}) = \frac{\sum\limits_{t=1}^{T} \xi_{a,t}\Phi(\frac{c-\epsilon_{m,t}}{h})}{\sum\limits_{t=1}^{T} \Phi(\frac{c-\epsilon_{m,t}}{h})}.$$

where Φ denotes application of the cumulative normal density function that produces probability-based weightings on observations where the weights are greatest for the most extreme standardised market return observations, $c = VaR(\epsilon_{m,t})$ is the constant empirical 1% VaR of market returns standardised using the GJR-GARCH volatility estimates and Silverman's "rule of thumb" method is used to determine the bandwidth h for the kernel (see, Silverman (1986)).

Practically, the estimation of MES requires conditional volatilities, a time varying correlation and conditional tail expectations. The conditional volatilities and correlation are estimated using an asymmetric DCC-GJR-GARCH process (see, Glosten et al. (1993)). This specification allows for both a leptokurtic distribution in the returns and common volatility characteristics such as volatility clustering. The asymmetry captures a leverage effect commonly observed in asset markets where negative returns increase volatility more than positive returns. As discussed by Engle (2002) an important feature of this specification is that the autoregressive parameters are restricted so that the expected long-run correlation is equal to unconditional correlation.

We test the adequacy of the GJR-GARCH specification by testing for autocorrelation in squared standardised returns. In the majority of cases the remaining ARCH effects are insignificant. In section 3.2 below we assess whether the DCC-GJR-GARCH model provides a reasonable MES estimate compared with alternatives.

2.4.1 VAR-for-VaR and MES Complementarity

Some further discussion on the differences, relative merits and complementarity of the VAR-for-VaR and MES methodologies is warranted. The VAR-for-VaR approach allows quantiles of related distributions to be modelled jointly and this reveals tail-risk exposure relations as opposed to conditional expected outcomes obtained from the MES approach. Estimation in the VAR-for-VaR case is not based on maximisation of a joint probability distribution and it doesn't require a distributional assumption. VaR estimates are simply projected from a set of right hand side (rhs) variables with parameters chosen to minimise the sum of weighted deviations of the observations from the quantile estimates (where the weights are chosen to make the expected number of exceedances equal to the chosen quantile probability). In this way, the fitting process is more localised to the tail of the distribution than would be true for GARCH-based VaR estimates (the latter are used in the MES analysis).

Our main interest in the VAR-for-VaR estimation is in obtaining the best possible indication of the time varying VaR (tail-risk exposures) across a variety of historical circumstances. This facilitates an accurate comparison of the tail risks being encountered by investors in the various SBBS tranches, individual sovereigns and a market-wide portfolio. The VAR-for-VaR impulse-response analysis also provides valuable insights regarding how exposures to extreme tail risks transmit from one SBBS tranche (or individual sovereign) to another. An insignificant VaR response is more likely for the case of safe-haven assets when lower-tail VaRs of other assets receive negative shocks. This provides complementary evidence of the risk protection benefits that stem from holding the safe-haven asset. The MES, on the other hand, identifies the conditional expected return of the safe asset conditional on a negative systemic event. The two measures (one for tail-risk exposure and the other for expected returns) are complementary methods for identifying safe haven assets.

An important complementarity that arises in the context of MES and VAR-for-VaR analysis is the ability to compare VaR responses and MES for the case of assets considered as safe havens. As discussed above, MES will be more adept in identifying a safe asset because it allows for the fact that the correlation between returns on a safe asset and the market (or a high risk asset) is more likely to be negative when the market is experiencing a tail event. Since the VAR-for-VaR impulse-response analysis does not condition the correlation of impulses on a scenario with elevated systemic risk, it is less likely to identify the opposing moves of safe assets relative to market returns during a crisis. Despite this weakness, we find that the VAR-for-VaR responses are muted in the case of low-risk sovereigns and senior tranches of SBBS indicating a high degree of protection from a more general worsening of tail risk exposures elsewhere.

2.5 Measuring Reward-for-Risk

To ascertain whether investors would be adequately rewarded for the expected risks in their holdings we calculate quasi-Sharpe ratios for SBBS tranches, a euro area GDP-weighted portfolio of sovereign bonds and individual sovereigns. The quasi-Sharpe ratio differs from the standard Sharpe in that we asses risks using conditional expectations rather than historical values. In this sense we are assuming that investors are taking a forward looking approach and a relatively long investment horizon so that temporary blips in the actual return relative to risk are avoidable. We consider two alternative approaches to the calculation of expected returns (in excess of the risk free rate) for the quasi-Sharpe numerator as well as two ways to represent conditional risk for the denominator. The expected returns variable will be either an annualised yield-to-maturity or an expected holding period return. The risk measures will be either a conditional standard deviation produced by a GJR-GARCH model or a VaR projection from a VAR-for-VaR model.

The annualised yield to maturity and expected returns (based on yield changes) are not always highly correlated so they give slightly different perspectives on reward for risk (the measures based on yield-to-maturity would generally be considered a longer term perspective). We do not observe the holding period returns directly so we infer these from yields subject to coupon/duration assumptions. The calculation of the excess return involves adding capital gains (or losses) to accrued coupons and subtracting a risk free rate. The coupon can sometimes be the dominant driver of holding period returns.[12] The standard deviation of the holding period return (i.e. standard deviation of the bond price changes) can be approximated as duration times the standard deviation of yield changes.

In our analysis the monthly holding period return is calculated as follows (where the euribor is the 1 month swap of fixed for floating interbank rate recorded at the beginning

[12] Unfortunately, we do not have an obvious way to determine what coupon would apply to the senior, mezzanine and junior tranches of SBBS. We address this by assessing the Sharpe calculations for a range of coupon assumptions based on the historical relation between coupons and yield spreads across issuers.

of period t);

$$r_t = \frac{\Delta P_t}{P_{t-1}} + (coupon(i) - euribor)_t/12 \tag{4}$$

We calculate the price change using an approximation based on yield changes suggested by Shirvani and Wilbratte (2005) as follows (where D is the duration);

$$\frac{\Delta P_t}{P_{t-1}} \approx (1 + \Delta y_t/(1 + y_{t-1}))^{-D} - 1 \tag{5}$$

This is also the basis for the conversion of yield volatility to price volatility as follows;

$$StDev\left(\frac{\Delta P_t}{P_{t-1}}\right) \approx (1 + StDev\left(\Delta y_t\right)/(1 + y_{t-1}))^{-D} - 1 \tag{6}$$

We approximate a high and low coupon based on the fitted relation between yield spreads and coupon spreads (where the yield spread is the spread relative to the bond with the lowest yield and the coupon spread is the spread relative to the risk free return). Using euro area sovereign bond coupon data from Bloomberg, we find that the approximate relation between the coupon spread and yield spread from January 2012 to end-October 2016 is as follows (standard errors of coefficient estimates are shown below the coefficient estimates);[13]

$$\begin{aligned} (coupon(i) - euribor)_t \quad &= \quad \underset{(0.35)}{3.08} + \underset{(0.21)}{0.42} \quad (yield(i) - min(yield))_t \\ Adj\bar{R}^2 = 0.23 & \end{aligned} \tag{7}$$

We use a two standard deviation shift of the intercept (in both directions) to derive lower and upper bounds for the coupon assumption for each yield and we restrict the coupon to a maximum of 10.5%. We also round the coupon to the nearest whole percentage and this leads to a variable that mimics more closely what we observe in the individual sovereign historical coupon series. In all related calculations we make sure that the duration is appropriately adjusted for the different levels of the coupons. The fitted coupon rises with yield. This may bias our results towards the finding of a higher Sharpe ratio for the relatively higher-risk sovereigns (and for the higher-risk SBBS). To counter this bias we also consider

[13]We excluded Greek bonds from this exercise due to the excessively large yield spread during the crisis since this dominates the regression when included.

a case where we imposed a static equal coupon across all the bonds/portfolios/SBBS in our analysis. This reveals a remarkably similar pattern and relative position for the Sharpe ratios.

3 Results

3.1 Tail Risk Exposure and Marginal Expected Shortfall

We present the majority of our estimation results in chart form while VAR-for-VaR diagnostics, assessment of the robustness of the MES approach and impulse-response analyses are each considered separately below. Figures will usually include the estimated 1% VaR (Value-at-Risk) and MES over time for the distribution of sign-reversed yield changes (i.e. minus the daily yield change measured in basis points). Comparisons are made between SBBS components, individual low-risk and high-risk sovereigns, and a GDP weighted portfolio of euro area sovereigns. For the two-tier 70:30 tranche design the VAR-for-VaR analysis was conducted with the junior bond placed first in the ordering of the system estimator - implying that the tail risk of the junior tranche is causal. For the three-tier tranche structure the VaR analysis was first conducted for the mezzanine as a function of the junior component and then for all other variables as a function of the mezzanine (i.e. for the senior SBBS, individual sovereigns and the EA portfolio). The MES for various tranches under the two- and three-tier tranche structures are compared with each other and with MES of individual sovereigns.

Figure 1 shows the case of a 70:20:10 tranche structure. This figure displays dot plots of the yield changes of mezzanine (light-blue dots) and junior (light-green dots) tranches of SBBS. We see that the mezzanine observations are distributed with considerably smaller variance than the junior bond observations. Other variables shown in this figure include the 1% VaR for the mezzanine and junior tranches (blue and purple lines respectively) and the MES of the mezzanine tranche (the dark-red line). The depiction of the VaRs through time give a first indication of how volatility is concentrated within the junior tranche. During the most volatile episodes the VaR of the junior tranche is much more negative than that of the mezzanine tranche (by a factor of three during the most volatile conditions).

Continuing with Figure 1 we observe that the mezzanine MES measure, which is conditioned on the probability of the entire market having a yield change more negative than its 1% VaR, is often significantly above the mezzanine 1% VaR and this indicates that market-wide tail events tend to coincide with non-tail yield movements in the mezzanine tranche (i.e. holders of the mezzanine benefit from a degree of insurance against both general and extreme systemic declines due to initial losses being accepted by junior bond holders). This is a theme that is repeated in the case of the senior tranche discussed below.

Since the junior SBBS generally has a very negative 1%VaR (and this is the most extreme SBBS VaR that we examine) it is immediately of interest to consider how this compares with VaRs of some of the high-risk individual sovereigns. Figure 2 shows the 70:20:10 junior SBBS 1%VaR compared with the 1%VaRs of each of three individual sovereigns that experienced the highest volatility during the Sovereign Debt Crisis (Portugal, Ireland and Greece). The Portuguese VaR lies mostly between the junior and mezzanine SBBS VaRs. It is interesting however that even the junior SBBS VaR is not as volatile as the highest-risk single sovereign VaR and quite often the PT-VaR falls well below the junior SBBS VaR. This reflects some diversification of the exposure to defaults among high-risk sovereigns available by holding the junior tranche so long as this tranche remains a large enough proportion of the securitisation. The Greek sovereign bond VaR is exceptional and often plunges several orders of magnitude below the junior VaR.[14] For the case of the Irish sovereign VaR, there is clearly a period between the Irish sovereign debt crisis (when guarantees given to depositors and senior bank bond holders imposed losses and re-capitalisation costs on the State) and the beginning of 2012 when risks were as great as, and occasionally in excess of, that of the junior SBBS. Otherwise the IE-VaR is practically inseparable from the mezzanine SBBS VaR.

Our next concern is whether the senior tranche is as safe as the lowest-risk sovereigns and whether it has safe-haven characteristics. Figure 3 shows the time profiles of a number of comparisons of risk measures between the German 10-year sovereign bond and the two most senior tranches of the SBBS under a 70:20:10 structure. The mezzanine 1% VaR is shown as the bright blue line. The senior SBBS and German 1% VaRs are shown as bright-green and purple lines respectively. These VaRs are almost indistinguishable from each other implying that the senior bond is just as low-risk as the German bond on the basis of a VaR comparison. The senior SBBS and German bond MES measures are shown

[14]Here it is worth noting that the scale of the vertical axis is roughly twice that of the other two panels.

as dark-blue and dark-red lines respectively. The MES profiles are very similar, but when volatility is at it highest (and when VaRs in general go sharply negative) we notice that both the German and senior SBBS MES estimates move in a positive direction (with the German MES rising marginally above the senior SBBS MES).

So the senior SBBS is almost (but not quite) as good a hedge against extreme market declines as is the German bond. This is to be expected since German bonds historically benefited from a flight to safety effect that does not apply so readily to the senior SBBS - the latter suffers losses in the most extreme circumstances by definition (and according to how the yield estimation procedure allocates yield premiums) whereas the German bond does not necessarily become exposed. Still, there are clearly only slight MES differences between the senior SBBS and the lowest-risk sovereign.

It is worthwhile considering how various single sovereign bonds and a diversified portfolio compare with the risk characteristics of the mezzanine and junior SBBS. Figure 4 shows VaR and MES risk measures for 6 of the relatively low-risk single European sovereigns compared with the VaR for the mezzanine SBBS under the 70:20:10 tranche structure. During volatile periods, these low-risk individual sovereign bonds have 1% VaRs that stay relatively stable and, in many cases, the whole distributions of returns on the low-risk sovereigns shift upward. Moreover, the MES of these low-risk sovereigns (i.e. the shortfall conditional on negative market outcomes below the market 1% VaR) tends to rise significantly, above their own 1%VaR, as the crisis intensifies and decline as it passes. This implies that there is a tendency towards positive returns for these lower-risk single-name sovereign bonds when there are extreme market-wide losses (there are several instances where expected returns on these lower-risk sovereigns are positive when there are extreme systemic declines). This is particularly apparent in the case of the German, Finnish and Dutch government bonds.

Figure 5 shows VaR and MES risk measures for the Italian and Spanish sovereign bonds along with the VaR and Expected Shortfall for the euro area GDP-weighted portfolio (which is our market/systemic indicator). In each case the 1%VaR for mezzanine and senior SBBS in the 70:20:10 tranche structure are included. The Italian and Spanish bond VaRs coincide almost exactly with the mezzanine SBBS VaR. The mezzanine tranche is clearly highly comparable with the Italian and Spanish cases. Somewhat surprisingly, the euro area portfolio has a VaR more in line with the senior SBBS. This demonstrates the sizeable benefits of diversification but such an investment is also guaranteed to experience systemic losses when they occur.

Likewise, the country-specific MES comparisons in Figure 5 suggest that shortfall experienced by investors in these single-named sovereigns conditional on market tail events, is generally of a similar magnitude to the country specific VaRs and (unlike for the low-risk sovereigns) this indicates that they do not benefit from virtually any of the rise in MES that is characteristic of sovereigns with safe-haven status. The cases of Italy and Spain are very clear (i.e. MES and VaR movements are almost exactly the same as each other in these two cases for the majority of the sample period). In the case of the euro area portfolio, despite being quite similar to other low-risk assets in terms of VaR, the diversified portfolio has Marginal Expected Shortfall which is always less severely negative than its own VaR (by definition) and is quite similar to that of the investment in Spanish or Italian bonds. This implies that the portfolio provides risk reduction through diversification but also exposes investors to all systemically risky events. This contrasts with an investment in the senior SBBS which has a Marginal Expected Shortfall that occasionally becomes positive during crisis periods.

Turning now to the case where there is a two-tier tranche structure (specifically a 70:30 structure). Figure 6 facilitates a comparison of the 70:20:10 and 70:30 SBBS structures. VaR comparisons are mainly of interest here, but the MES for the senior tranche is also shown as the dark blue line. Note that the 1% VaR for the senior SBBS (shown as the orange line) is the same for any SBBS structure that has a 70% weight for the senior tranche. The grey-shaded area identifies the 70:30 case and this grey area is bordered by the MES of the senior tranche and the VaR of the 70:30 junior tranche. In systemically risky periods, the mezzanine SBBS from the 70:20:10 structure has a VaR which is significantly above the junior VaR of the 70:30 structure. However, the junior VaR associated with the 70:20:10 structure is much more prone to extremely negative outcomes than that of the 70:30 SBBS.

It is noticeable that the senior and junior VaRs show some negative correlation during the onset of the Great Financial Crisis but not beyond that. More interestingly, from the beginning of the crisis, the senior SBBS MES rises towards the middle of the distribution of (sign-reversed) yield-changes while the junior SBBS VaR becomes extremely negative (in the pre-crisis period the senior SBBS MES stays at the lower fringes of the yield change distribution). This suggests that the senior SBBS is a hedge against the junior SBBS tail risk when tail risk is extremely high.

The expected value of the return on each entity conditional on the return for the junior yield change being in the extreme tail depends critically on the correlation between the returns on individual sovereigns with those of the junior SBBS. When the correlation is close to 1 the MES tends to be close to or below the VaR. If the correlation moves into negative territory the asset pairing involves a hedge relationship and in this case the returns will tend to be observed in opposite sides of their respective distributions. Figure 7 shows this relationship between correlation and MES (where the correlation is for the DE and junior SBBS yield changes while the DE MES is conditioned on the junior SBBS yield change being more negative than the junior SBBS 1%VaR).

It is possible to calculate the contribution of SBBS tranches to systemic risk when the latter is defined in terms of Expected Shortfall. The SBBS securitisation is simply a rearrangement of the entire market. Hence, it is the case that the weighted-sum of the Marginal Expected Shortfalls of the tranches equals the market Expected Shortfall (with the weight in each case being the proportion of the market portfolio that acts as collateral for the tranche, i.e. the tranche proportions). The weighted contributions to Expected Shortfall are displayed in Figure 8. The horizontal lines show the tranche proportions. The contributions to systemic risk of the SBBS senior and sub-senior tranches move away from their tranche proportions at the height of the sovereign debt crisis.[15] For the case of the senior tranche we observe a sharp fall in its proportional contribution to Expected Shortfall from early 2010 and at this time the sub-senior tranches start to account for up-to 70% of Expected Shortfall despite representing only 30% of the sovereign bond market.

3.2 Robustness of MES measurement

Diagnostic tests for the effectiveness of the GJR-GARCH analysis are provided in Table 1, Appendix A. P-values are shown for Langrange-Multiplier test statistics (that are robust to heteroscedasticity) where the null is the joint insignificance of autocorrelation (at lags 1 to 3) in squared standardised residuals where the demeaned returns are standardised using volatility estimates from the GARCH models used in the MES estimation. The tests are

[15] Note that the mezzanine and junior tranche contributions have been combined for clarity.

applied for the cases of each country and each of the SBBS tranches in the 70:20:10 SBBS structure.

The test results indicate that almost all of the cases are adequately described by the GJR-GARCH specification used in the MES analysis. The most obvious case where there is some doubt about the adequacy of the GARCH is for the junior tranche of the SBBS. In this case there is evidence that the standardised residuals still contain some significant ARCH effects at lag 3. Other cases where there is some evidence of remaining ARCH effects include Germany, Portugal and the Mezzanine tranche.

In these cases we examined whether an alternative volatility modelling procedure would deliver a very different MES measurement. Specifically, we examined the outcome obtained by replacing the GJR-GARCH estimates with conditional volatility obtained from a stochastic volatility (SV) model proposed by Chan and Grant (2016). The SV model easily passed L-M tests for insignificance of remaining ARCH (for the junior SBBS the relevant p-values are 0.2044, 0.2618 and 0.4165 for tests that include the first, first and second and first three lags respectively).[16] Figure 9 provides a comparison of the original MES for the junior tranche with that obtained using the Stochastic Volatility modelling over a subsample of the data from Sept 2010 to Oct 2016. We can see that the original MES measure tends to be lower than the alternative SV estimate and the difference between these becomes larger for the very extreme negative values.[17] Given that the original measure (the Brownlees and Engle methodology) tends to provide a more negative outcome, we base our analysis and conclusions on this approach.

3.3 VAR-for-VaR Diagnostics & Impulse-Response Analysis

Diagnostic tests for the effectiveness of the VAR-for-VaR analysis are provided in Appendix B. The first column of the table indicates the Figure displaying the VaR time series to which the test results refer. The second column indicates the country or SBBS that is the subject of the test (the causal variable is the mezzanine SBBS). Column 3 shows the p-values for the DQ tests and column 4 shows the percentage of exceedances of the estimated VaR in each

[16]Specifically, we use the SVMA version which allows the returns process to have an MA(1) process. We retained the original DCC correlation parameter for the MES estimation. We found that replacing the DCC with a moving sample correlation did not change the conclusion drawn from the comparison.

[17]We only show the comparison of MES for the case of the junior SBBS but the other borderline cases also had similar MES profiles and these comparisons are available from the authors on request.

case (for the 1% VaR this should be close to 1%). Overall, these test results indicate that the VaR-for-VaR estimates have good properties in terms of the number and randomness of exceedances (with all p-values for the DQ test greater than or equal to 0.10 and the number of exceedances very close to 1% in all cases).

The last two columns of Table 2 also provide test results for the joint significance of cross-effects in each of the of VAR-for-VaR models displayed in Figures 1 to 5. The most obvious conclusion to draw from these cross-effects tests is that (with the exception of Finnish and Austrian yields) cross-effects to core countries and to the senior tranche of the SBBS are absent. The German and Senior tranche have insignificant spillovers from the Mezzanine SBBS. All of the higher risk sovereigns show evidence of cross-effects involving interaction with the VaR of the mezzanine SBBS. Since there are significant effects between the junior and mezzanine SBBS (and since increased risks in the junior are a prerequisite for losses in the mezzanine) cross-effects also exist between the junior and higher-risk sovereigns.[18]

The cross-effect tests can only be regarded as indicative of interactions for two reasons. Firstly, the contemporaneous correlations between the estimated VaRs are not addressed in these tests. Secondly, as with all multivariate VAR applications, the right-hand-side variables tend to be highly correlated with each other and this multicollinearity increases the variance of all individual parameter estimates. We therefore turn to an impulse-response analysis to capture the interactions more comprehensively. Our main interest in this analysis remains to provide evidence that cross-effects are absent for the case of safe haven assets (particularly for the case of the senior SBBS).

For VAR pairings of assets $\{i, j\}$ we consider the response in the VaR of variable j to a one unit shock in the absolute return of variable i. We therefore scale the L matrix in Equation 2 above as follows;[19]

$$\begin{bmatrix} l_{1,1} & 0 \\ l_{2,1} & l_{2,2} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \frac{\sigma_{i,j}}{\sigma_i^2} & \frac{\sqrt{\sigma_i^2 \sigma_i^2 - \sigma_{i,j}^2}}{\sigma_i^2} \end{bmatrix} \tag{8}$$

For 6 VAR-for-VaR pairings, Figure 10 shows 1% VaR responses to a one unit absolute

[18]Test results for these additional effects are available from the authors on request.

[19]In this case the normal Cholskey decomposition $LL' = Cov(r_i r_j)$ is modified by pre and post-multiplication by D with diagonal elements σ_i^-1 and zero elsewhere.

return shock in the junior or mezzanine SBBS of a 70:20:10 tranche structure. Blue lines indicate two standard deviations above and below the mean response. The following pairings are considered and the responses shown are for the second variable 1% VaR in response to a one unit shock to the absolute returns of the first; {junior, mezzanine}, {junior, market}, {junior, senior}, {junior, DE}, {mezzanine, senior} and {mezzanine, DE}.

The first two cases indicate that the mezzanine SBBS and the market portfolio both have significant negative responses to the shock to returns of the junior SBBS. The contemporaneous responses in both cases are much smaller than the unit shock to the junior SBBS (the mezzanine reacts with a size of 0.15 while the market portfolio responds with a 0.035 reaction). These small responses are not surprising given that the junior protects the mezzanine from first losses and the market involves significant dilution through the presence of safer assets in the portfolio as well as through diversification benefits.

The second and third cases concern the 1% VaR response in two very safe assets to a return shock in the junior SBBS returns. This reveals that the senior SBBS is just as insensitive to return shocks in the junior bond as is the German sovereign. The final two cases in the bottom row of Figure 10 depict the response of 1% VaR of the the senior SBBS and German sovereign to a one unit shock in the mezzanine returns. In this case there is still an insignificant pass through to the conditional 1% VaR of two low risk assets despite the fact that such a shock to the mezzanine would imply large expected losses for the junior bond).

It is also interesting to note that, although statistically insignificant, the point estimates for the contemporaneous responses of the VaRs for the senior SBBS and the German sovereign are both slightly positive (this is more noticeable for the DE case). As mentioned above, the correlation of returns conditional on a systemic tail event (the MES) will generally be larger than this response for the case of safe haven assets.

3.4 Return for Risk

The above analysis considers only comparisons of risk measures. We now consider ex ante reward for risk. We begin by describing the results of a Sharpe ratio analysis for the holding period returns. As mentioned in the methodology section, this ratio is a little unusual for the case of bonds because coupons sometimes represent a large and changeable component of holding period returns. The results now discussed are based on the case where an equal (and static) coupon and duration is used for the calculations of all holding period returns. The results for the estimated coupons (where coupons are allowed to rise with the yield spread) tend to show relatively larger Sharpe ratios for the riskier bonds/SBBS.

Figure 11 shows the case of the monthly holding period Sharpe ratios for the senior, mezzanine and junior tranches of the SBBS combined with the Sharpe ratios for the monthly holdings of German sovereign bonds (these are all for the 10 year maturity and for the 70:20:10 tranche structure). In each case we allow the coupon to have an upper and lower bound of one standard deviation around the chosen coupon rate based on the standard error of the intercept coefficient in the regression of the coupon spread on yield spread as discussed in the methodology section. This reveals that Sharpe ratios are generally close together. All Sharpe ratios are low and declining during the Great Financial Crisis until 2010. The financial crisis began to affect peripheral sovereigns during 2010. This seems to coincide with an increase in the Sharpe ratios for the mezzanine, senior and German sovereigns (perhaps reflecting safe haven flows). The junior SBBS turns around later and it is plausible that the ECB's non-standard monetary policy measures were responsible for reducing the risks associated with holding the higher-risk sovereigns and therefore the junior SBBS. Sharpe ratios tend to rise towards early-2012 (with the junior SBBS peaking far above the others at a value near 9) and then all but the junior SBBS tend to stay around a value of 4 for the remainder of the sample while the Sharpe ratio of the junior SBBS declines to zero or below for the end of the sample.

Figure 12 is the same as the previous figure except that the German Sharpe ratio is replaced by the Sharpe for the euro area portfolio (EA Sharpe). Firstly, the EA Sharpe is very similar to the German case. The EA portfolio gives slightly more reward for risk than the senior and mezzanine SBBS during the crisis periods but otherwise it is very similar to these. Figure 13 shows the SBBS Sharpe ratios against the background of the spread between the lowest and highest individual sovereign Sharpe ratios over time (the purple shaded area). This indicates that the SBBS Sharpe ratios (excluding the junior) are generally contained within the range covered by individual sovereigns. On this basis it appears that the SBBS provide as good (if not better) holding period reward for risk to that of the individual sovereigns.

We now examine the compensation for risk where risk is measured either as GJR-GARCH conditional volatility or as VAR-for-VaR (these are therefore variations of the concept underlying the simple Sharpe ratio).[20] Figure 14 shows the dynamic relative reward for GJR-GARCH conditional volatility of individual sovereigns, the euro area portfolio and senior, mezzanine and junior SBBS. The top panel considers the comparison for the group of low-risk sovereigns while the bottom panel pertains to the relatively more risky sovereigns. The sample average comparison is shown in Figure 15.

It is clear from the time-varying Sharpe ratios in the top panel of Figure 14 that the senior SBBS is rewarded in a very similar manner to the German bond. Both seem under-rewarded when compared with other low-risk sovereigns but this reflects the fact that they have substantial flight-to-safety price premiums. The black dotted line represents the Sharpe ratio for the senior SBBS in the top panel. This almost always lies directly over the German reading (shown in purple). There is some evidence of a difference between the German and senior SBBS during the pre-crisis period (this was a period when German bonds attracted a slightly higher return for risk - perhaps due to the fiscal situation in Germany driving bond prices down slightly with little change in volatility). Overall however, there is very little difference between the senior SBBS and the German bond in terms of their yield-to-maturity relative to GJR-GARCH volatility. This confirms the ranking of the senior bond in our earlier analysis based purely on risk. The Sharpe ratio on average in Figure 15 confirms the approximate equality between German and senior SBBS. We also see that the mezzanine has a similar reward ratio to that of the French bond. It is notable that the GDP-weighted euro area portfolio frequently has the highest yield relative to volatility. This reflects the strong diversification effects available (with many low and sometimes negative pairwise correlations).

The time varying Sharpe ratios in the bottom panel of Figures 14 and 15 reveal that the junior SBBS is generally not well compensated for its volatility. This under-performance is more prominent in the pre-crisis period. During the crisis and post-crisis period there

[20]This analysis is conducted without considering the fact that coupons could be distributed differently than is presumed under the estimation of SBBS historical yields. Since the senior SBBS is as low-risk as the German bund it is likely that it would be rewarded with a coupon no greater than the bund. This is less than what is paid on the pool of low-risk sovereigns and the prospect of redistributing such extra coupons to the junior SBBS arises. Such a redistribution would change the warranted yield-to-maturity of the junior tranche relative to what we have assumed in our SBBS yield estimations.

is regularly a relatively high ranking of the junior SBBS in terms of reward for risk. As mentioned above, there is a possibility that the junior bonds may be subject to receipt of extra coupons relative to what is assumed in the Monte Carlo analysis. On this basis it seems that junior tranches could be made more attractive to investors. While the junior SBBS does not appear to be well compensated it is clear that the mezzanine is very much in line with the relatively high-risk sovereigns such as Spain, Italy, Ireland and Portugal.

In Figures 16 and 17 the dynamic and average yield for absolute Value-at-Risk earned by individual sovereigns, the euro area portfolio and the senior, mezzanine and junior SBBS are shown (the average in Figure 17 is for the whole sample from January 2002 to October 2016). Once again, the top panel considers the comparison between a euro area portfolio, senior and mezzanine SBBS and the group of lower-risk individual sovereigns. The comparison in the bottom panel pertains to the riskier sovereigns (the junior SBBS is included). It turns out that there are very few differences between a VaR based risk-return assessment and one based on conditional volatility - except for the case of the junior SBBS where the reward for risk is no longer such an outlier. The euro area portfolio performs best overall (but recall that this portfolio has little of the hedging properties possessed by the German or senior bonds). The senior bond has a reward for risk which is very similar to that of the German bond. The bottom panel of Figure 17 reveals that the junior SBBS provides a reward for risk (using Value-at-Risk) than borders the lowest performers (albeit not as outlying). Overall, the SBBS are not excessively out-of-line with other similarly risky single sovereigns.

It is important to note that the dynamic quasi-Sharpe ratios considered in this analysis are based on forward-looking ex ante holding period returns in excess of the risk free rate or yields-to-maturity relative to conditional volatility (or VaR) of such returns/yields. During the crisis these Sharpe measures are driven high for the junior SBBS by the fact that coupons (i.e., accrued interest) stay relatively static while the cost of a junior SBBS investment declines markedly following unexpected falls in the prices of higher risk bonds. Hence, the forward-looking returns appear very high relative to the cost of the investment and even relative to the standard deviation or VaR of the returns. It is important to note that junior SBBS would in fact have been exposed to very large unexpected negative returns during the crisis. But these were rare and unpredictable events that are not well represented using ex ante or average returns. Such tail outcomes are best assessed by reference to time varying VaR metrics as discussed in section 3.1 above.

4 Conclusion

The analysis above examines the *ex ante* tail-risk characteristics of sovereign bond-backed securities as proposed by Brunnermeier et al. (2011) and assessed in Brunnermeier et al. (2016, 2017). The results of this analysis largely confirm the simulation-based results of Brunnermeier et al using correlated probabilities of defaults and expected loss given default calibrated on historical data. The original simulations found that senior SBBS would be exposed to smaller amounts of risk than even the lowest-risk existing euro area sovereign bond and that junior SBBS would surpass most euro area periphery sovereigns in their pricing. Using Value-at-Risk and Marginal Expected Shortfall measures we have found that senior SBBS are almost as low-risk as the lowest-risk euro area sovereign (including as a hedge against the extreme risk of many defaults). The mezzanine SBBS is very similar to Spanish and Italian bonds using all risk measures. The junior SBBS under a 70:20:10 design is not as high-risk as the highest-risk single sovereign but is usually more exposed to market-based losses than sovereigns with intermediate levels of risk exposures. On this basis, the junior bonds may attract investor interest when one considers their likely higher liquidity than existing single high-risk sovereigns.

We found that the senior SBBS has a yield-to-maturity relative to Value-at-Risk that is similar to that of the German bund. The junior bond frequently outperforms in terms of the dynamic Sharpe ratio. In the pre-crisis period and at the end of the sample it under-performs. However, since junior bonds may benefit from higher coupons (as a net benefit from the securitisation), and are likely to be more liquid than comparable individual sovereigns, it should be possible to enhance the junior components of SBBS to make them attractive to investors.

While the VAR-for-VaR approach is flexible, in that it allows for cross-effects, it does not allow for changing parameters in the VAR. Tail risk spillovers may primarily be a feature of crisis circumstances so allowing for parameters to switch in such circumstances may materially affect the findings above. The MES analysis reminds us that the correlation between low- and high-risk assets matters a lot for the perceived (and actual) risk exposure in crisis situations. Some individual sovereigns have very attractive hedging properties (i.e. high MES) and, while this feature is passed on to the senior SBBS, it may be counteracted by actual exposure to losses in the rare circumstance of a very large number of defaults. Our simulations have accounted for these circumstances by assuming a very high correlation for defaults in each period. While we have calibrated the yield estimation process to guard against a safe-haven bias we have not tried to represent the benefits that securitisation might bring in terms of reducing risks due the breaking of a bank-sovereign diabolic-loop. In this respect, our findings are more likely to be an understatement of the benefits of securitisation.

References

Acharya, V., R. Engle, and M. Richardson (2012). Capital Shortfall: A New Approach to Ranking and Regulating Systemic Risks. *The American Economic Review 102*(3), 59–64.

Adrian, T. and M. K. Brunnermeier (2016). CoVaR. *American Economic Review 106*(7), 1705–1741.

Alogoskoufis, S. and S. Langfield (2018). Regulating the doom loop. *ESRB Working Paper Series No 74*.

Brownlees, C. T. and R. F. Engle (2012). Volatility, correlation and tails for systemic risk measurement. *University of Chicago, Becker Friedman Institute, Working Paper*.

Brownlees, C. T. and R. F. Engle (2017). SRISK: A Conditional Capital Shortfall Measure of Systemic Risk. *Review of Financial Studies 30*(1), 48–79.

Brunnermeier, M. K., L. Garicano, P. Lane, M. Pagano, R. Reis, T. Santos, D. Thesmar, S. Van Nieuwerburgh, and D. Vayanos (2011). European safe bonds. *The Euronomics Group, Summary available at VoxEU.org*.

Brunnermeier, M. K., L. Garicano, P. R. Lane, M. Pagano, R. Reis, T. Santos, D. Thesmar, S. Van Nieuwerburgh, and D. Vayanos (2016). The Sovereign-Bank Diabolic Loop and ESBies. *American Economic Review Papers and Proceedings 106*(5), 508–512.

Brunnermeier, M. K., S. Langfield, M. Pagano, R. Reis, S. Van Nieuwerburgh, and D. Vayanos (2017). ESBies: Safety in the Tranches. *Economic Policy 32*(90), 175–219.

Chan, J. and A. Grant (2016). Modeling Energy Price Dynamics: GARCH versus Stochastic Volatility. *Energy Economics 54*, 182–189.

Cholesky, A. (1910). On the Numerical Resolution of Systems of Linear Equations. *Bulletin de la Société des Amis de la Bibliothèque de lÉcole Polytechnique 2005*(39).

Engle, R. F. (2002). Dynamic Conditional Correlation - A Simple Class of Multivariate GARCH Models. *Journal of Business and Economic Statistics 20*, 339–350.

Engle, R. F. and S. Manganelli (2004). CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles. *Journal of Business & Economic Statistics 22*(4), 367–381.

ESRB High-Level Task Force on Safe Assets (2018). Sovereign bond-backed securities: A feasibility study. European Systemic Risk Board.

Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993). On the Relation Between the Expected Value and the Volatility of the Nominal Excess Rreurn on Stocks. *The Journal of Finance 48*(5), 1779–1801.

Koenker, R. and G. Basset (1978). Regression Quantiles. *Econometrica 46*(1), 33–50.

Manganelli, S., H. White, and T. Kim (2015). VAR-for-VaR: Measuring Tail Dependence Using Multivariate Regression Quantiles. *Journal of Econometrics 187*, 169–188.

Scaillet, O. (2005). Non-Parametric Estimation of Conditional Expected Shortfall. *Insurance and Risk Management Journal 85*, 639–660.

Schönbucher, P. J. (2003). *Credit Derivatives Pricing Models: Models, Pricing and Implementation.* London: Wiley.

Sharpe, W. F. (1966). Mutual fund performance. *Journal of Business 39*, 119–138.

Sharpe, W. F. (1994). The sharpe ratio. *Journal of Portfolio Management 21*, 49–58.

Shirvani, H. and B. Wilbratte (2005). Duration and Bond Price Volatility: Some Further Results. *Journal of Economics & Finance Education 4*(1), 1–6.

Silverman, B. W. (1986). *Density Estimation for Statistics and Data Analysis.* London: Chapman and Hall.

5 Figures

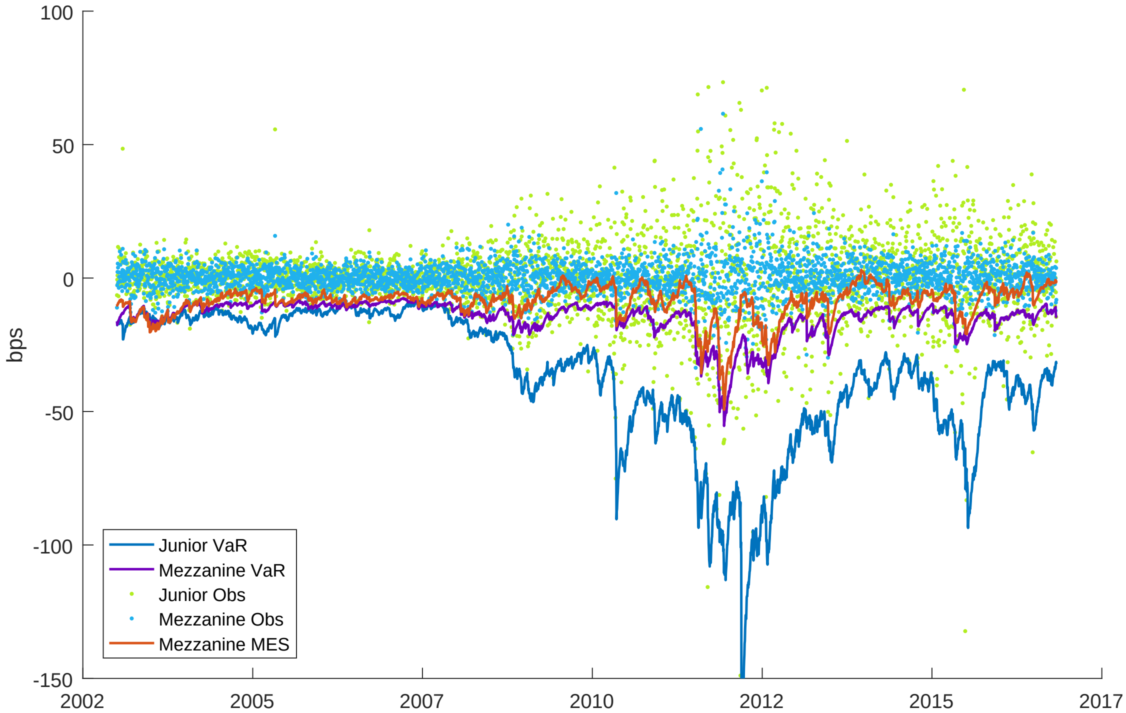

Figure 1: For the case of a 70:20:10 tranche structure this figure displays dot plots of the observed changes in yields in basis points on mezzanine (Light-Blue dots) and junior (Light-Green dots) tranches of SBBS along with 1% VaR and MES measures of risk exposures. The blue and dark-blue lines represent the 1% VaRs for the junior and mezzanine SBBS respectively while the MES for the case of the mezzanine tranche (i.e. expected return conditional on the entire market experiencing a yield change more negative than its 1% VaR) is displayed as a dark-red line.

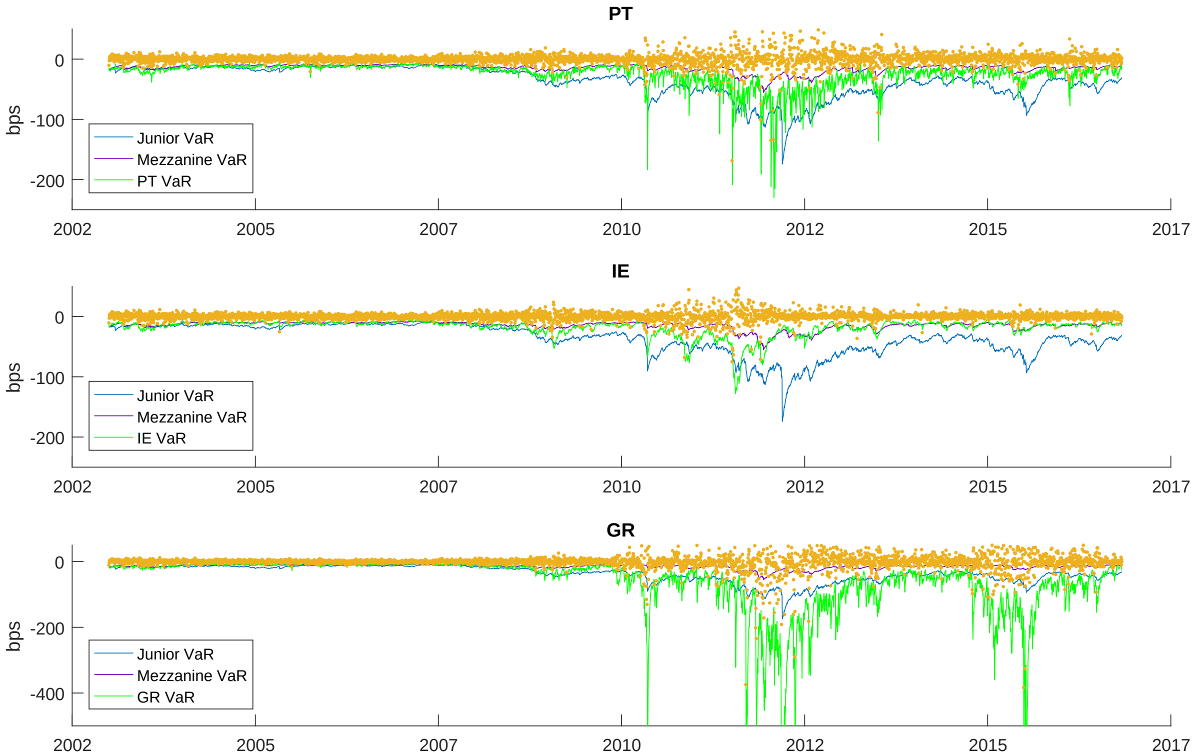

Figure 2: This figure shows the 70:20:10 junior SBBS 1%VaR compared with the 1%VaRs of each of three individual high-risk sovereigns (Portugal, Ireland and Greece).

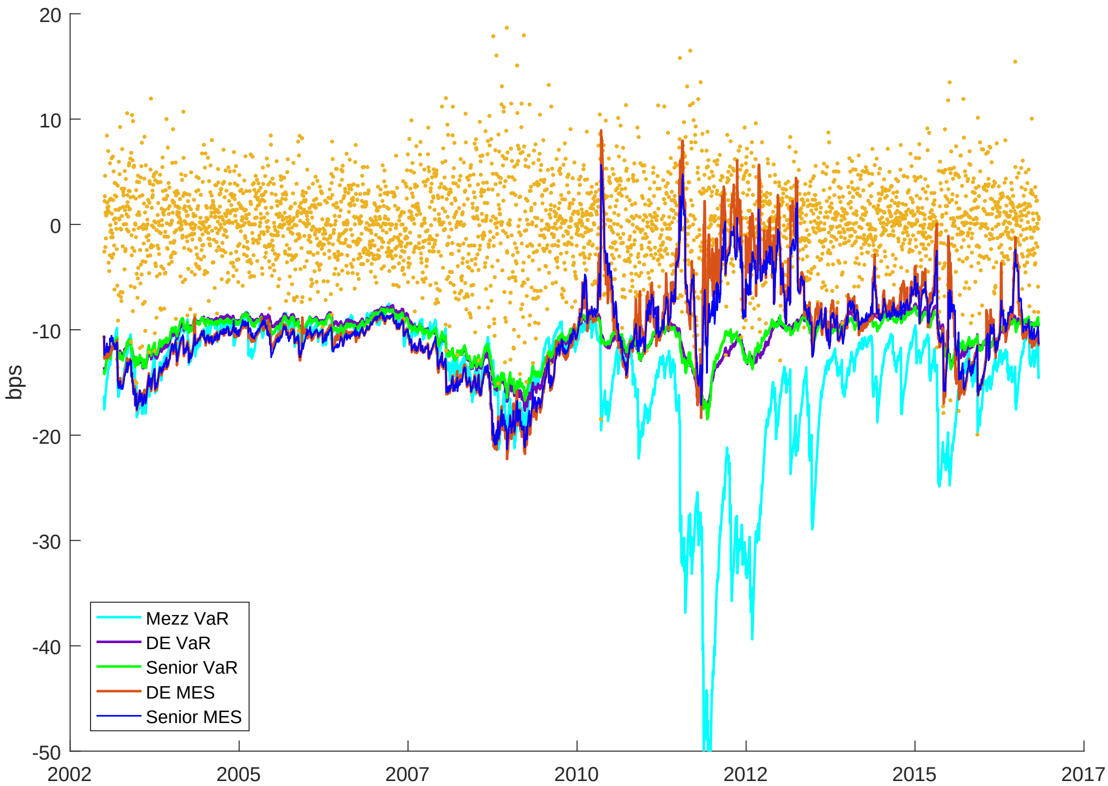

Figure 3: This figure shows the time profiles of comparisons of risk measures (in basis points) of the German 10-year sovereign bond and the two most senior tranches of the SBBS under a 70:20:10 structure. The mezzanine 1% VaR is shown as the bright-blue line. The senior SBBS and German 1% VaRs are shown as bright-green and purple lines respectively. The senior SBBS and German bond MES risk exposure measures are shown as dark-blue and dark-red lines respectively.

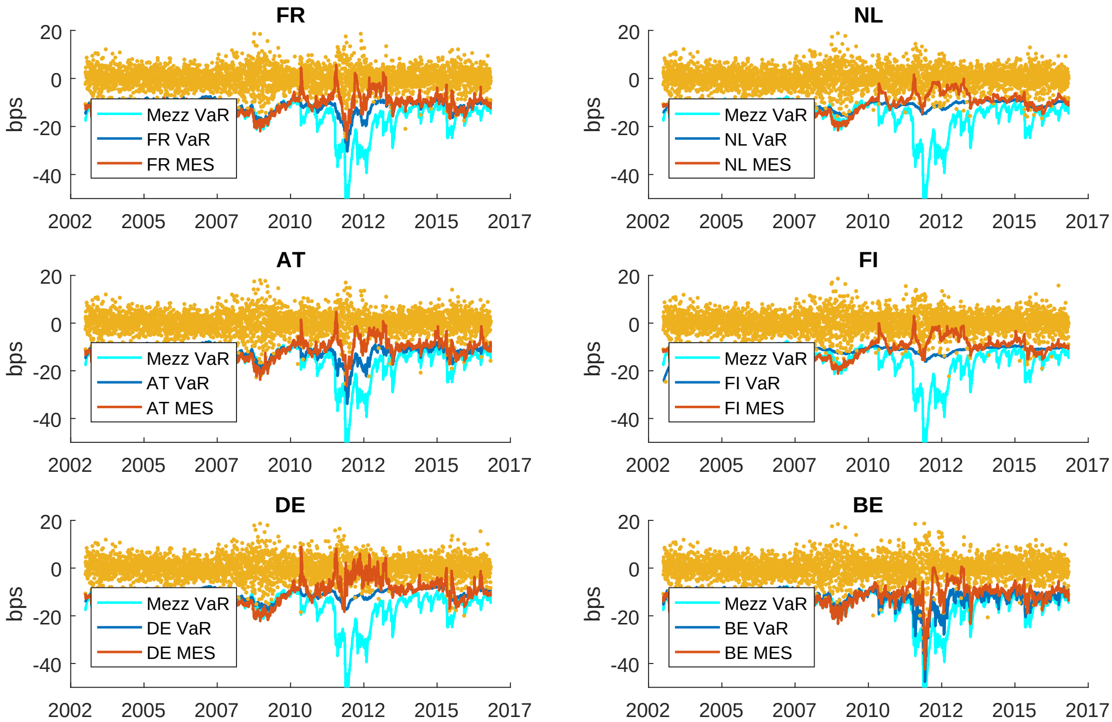

Figure 4: This figure shows VaR and MES risk measures for 6 of the relatively low-risk single European sovereigns compared with the VaR for the mezzanine SBBS under the 70:20:10 tranche structure.

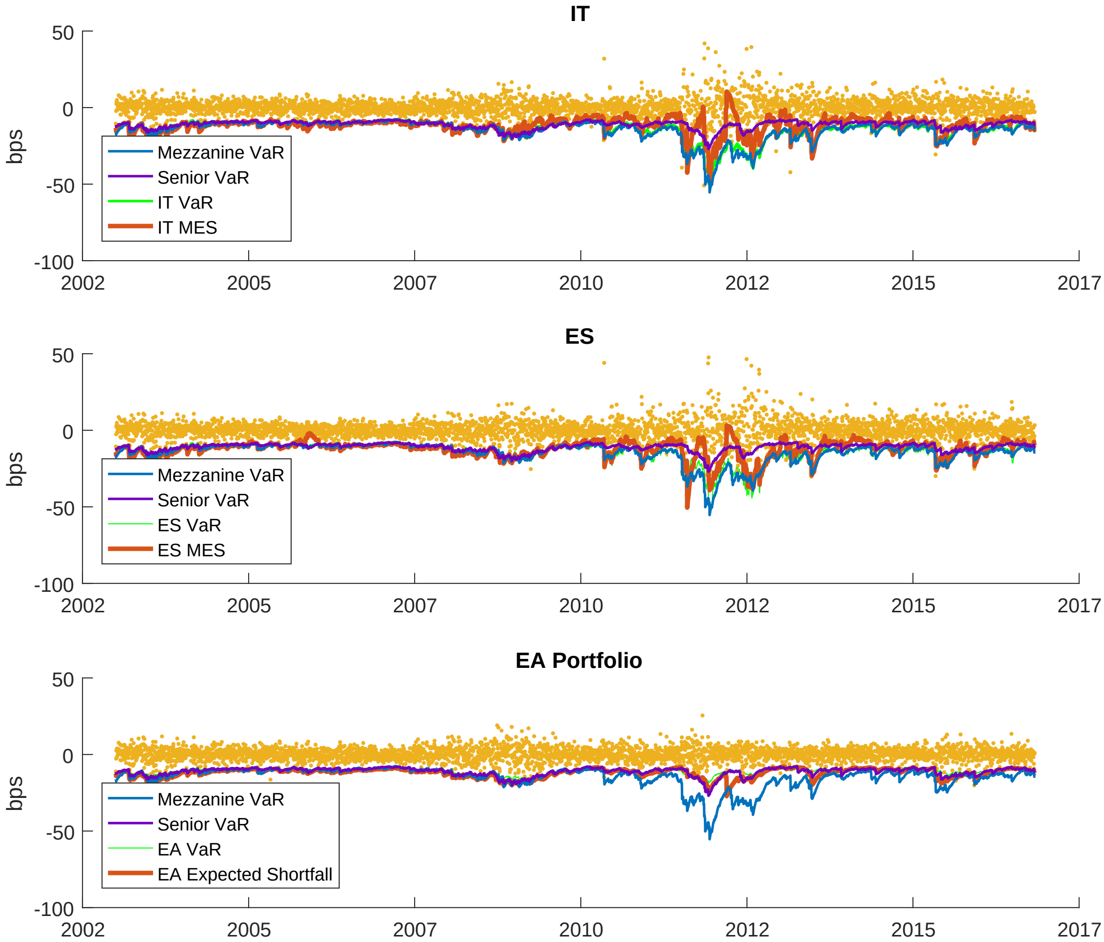

Figure 5: This figure shows VaR and MES risk measures for Italian and Spanish sovereign bonds along with the VaR and Expected Shortfall for a euro area GDP-weighted portfolio. These are compared with the 1% VaR of senior and mezzanine SBBS under the 70:20:10 tranche structure.

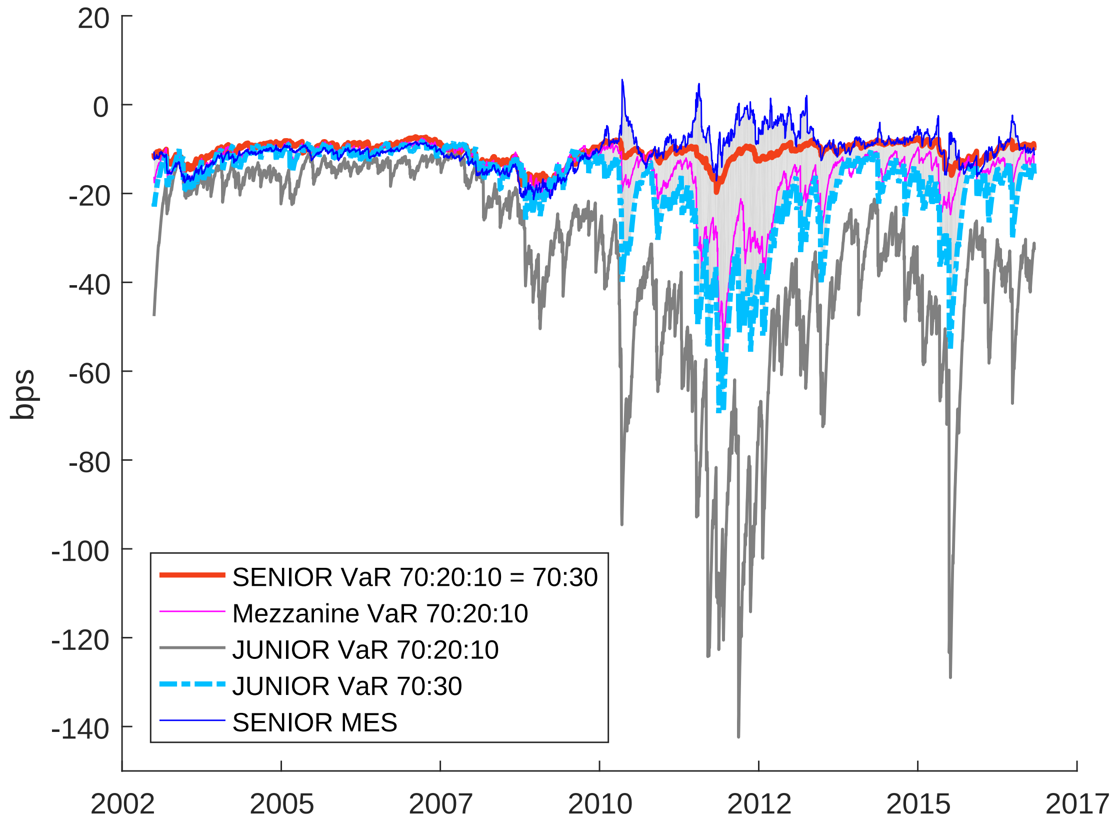

Figure 6: This figure facilitates a comparison of the 70:20:10 and 70:30 SBBS structures. VaR comparisons are mainly considered, but the MES for the senior tranche is also shown as the dark blue line. Note that the 1% VaR for the senior SBBS (shown as the orange line) is the same for the 70:20:10 and 70:30 SBBS structures. The grey-shaded area identifies the 70:30 case and this region is bordered by the senior MES and the VaR of the 70:30 junior tranche. In systemically risky periods, the mezzanine SBBS from the 70:20:10 structure has a VaR which is significantly above the junior VaR of the 70:30 structure. However, the junior VaR from the 70:20:10 structure is much more prone to extremely negative outcomes than that of the 70:30 SBBS.

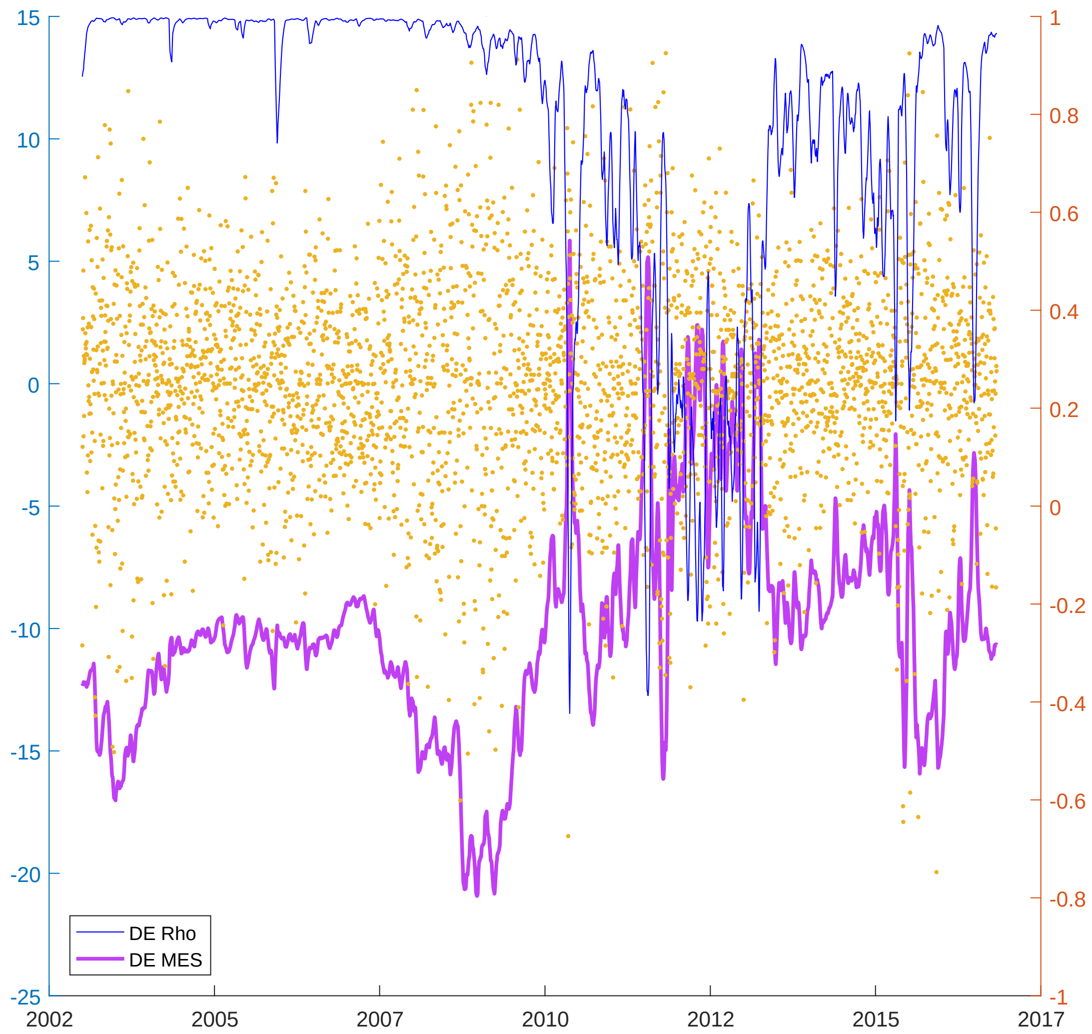

Figure 7: This figure displays the 1% VaR and MES for 10 Year German Sovereign Bond along with the coefficient of correlation between yield changes of the Bund and junior SBBS (derived from DCC Garch analysis). The correlation coefficient changes to negative during the crisis and this affects the estimate of the MES making it more positive.

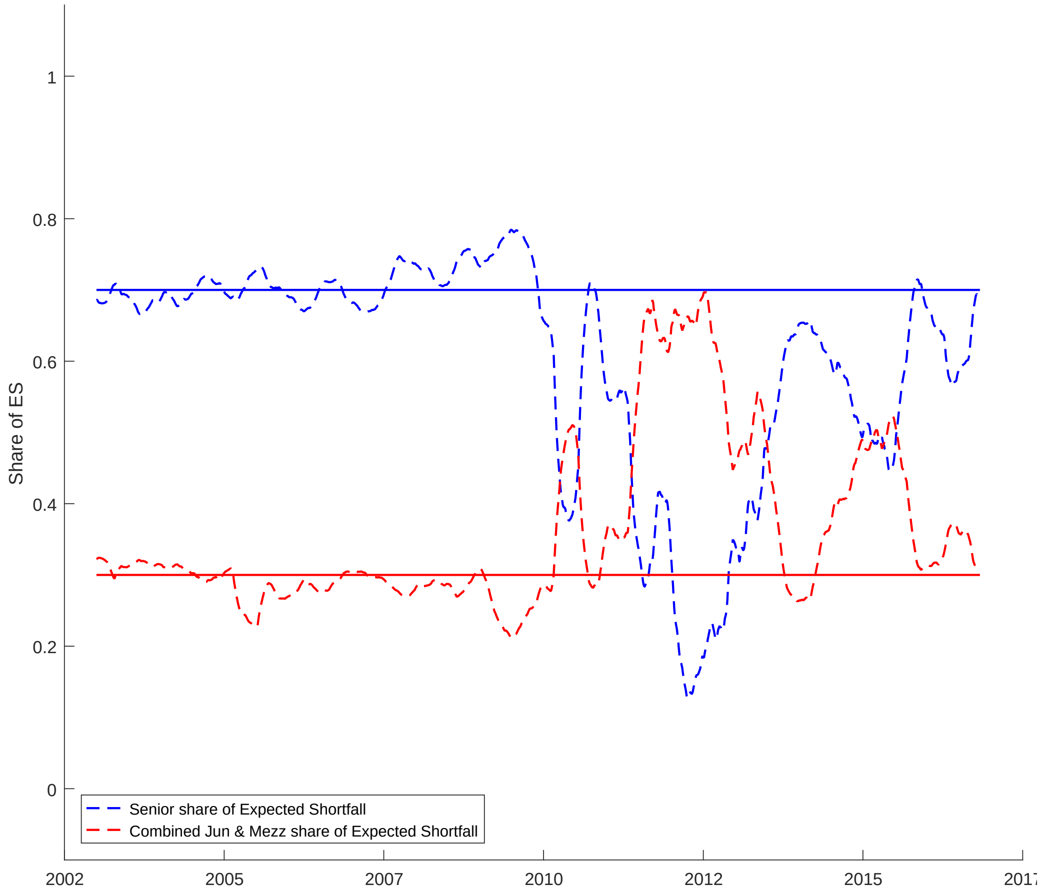

Figure 8: This figure displays the contributions to systemic risk of the SBBS senior and sub-senior tranches (where systemic risk is measured as the Expected Shortfall of the entire market). The horizontal lines give the shares of the entire market collateralising each of the tranches (expected contributions will be equal to these market shares if market risks are not specifically tail events).

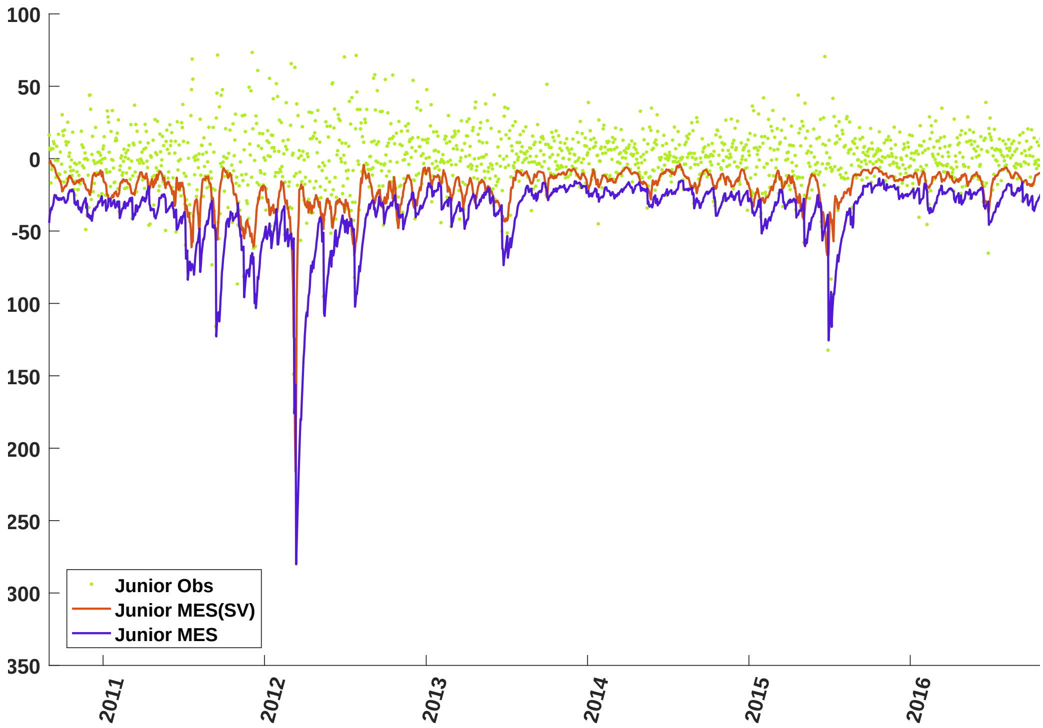

Figure 9: This figure displays MES calculated in two different ways. One is based on a strict application of the Brownlees and Engle (2012) approach. The other, MES(SV), replaces the GJR-GARCH volatility estimates used in the standard approach with alternative estimates from a stochastic volatility model suggested by Chan and Grant (2016).

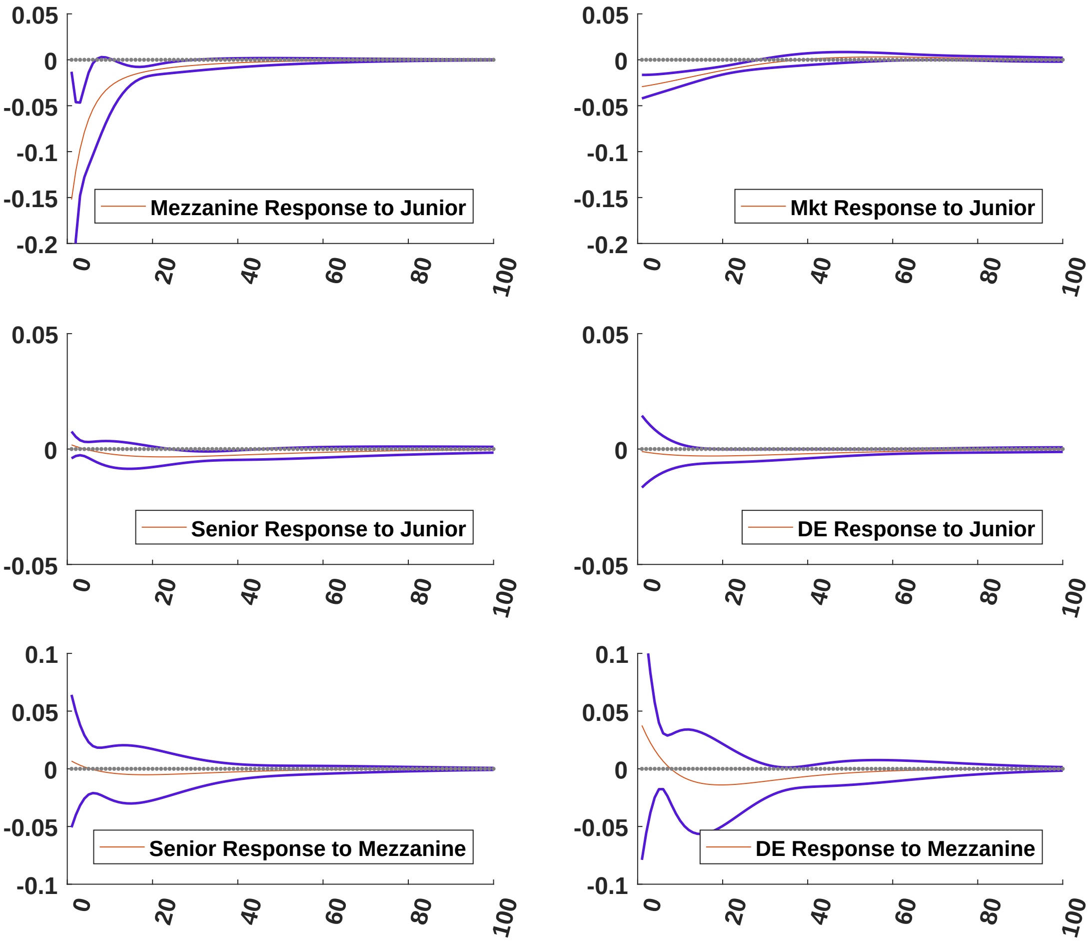

Figure 10: For 6 VAR-for-VaR pairings, this figure shows VaR Responses to a one unit absolute return shock in the junior or mezzanine SBBS of a 70:20:10 tranche structure. Blue lines represent two standard deviations above and below the mean response.

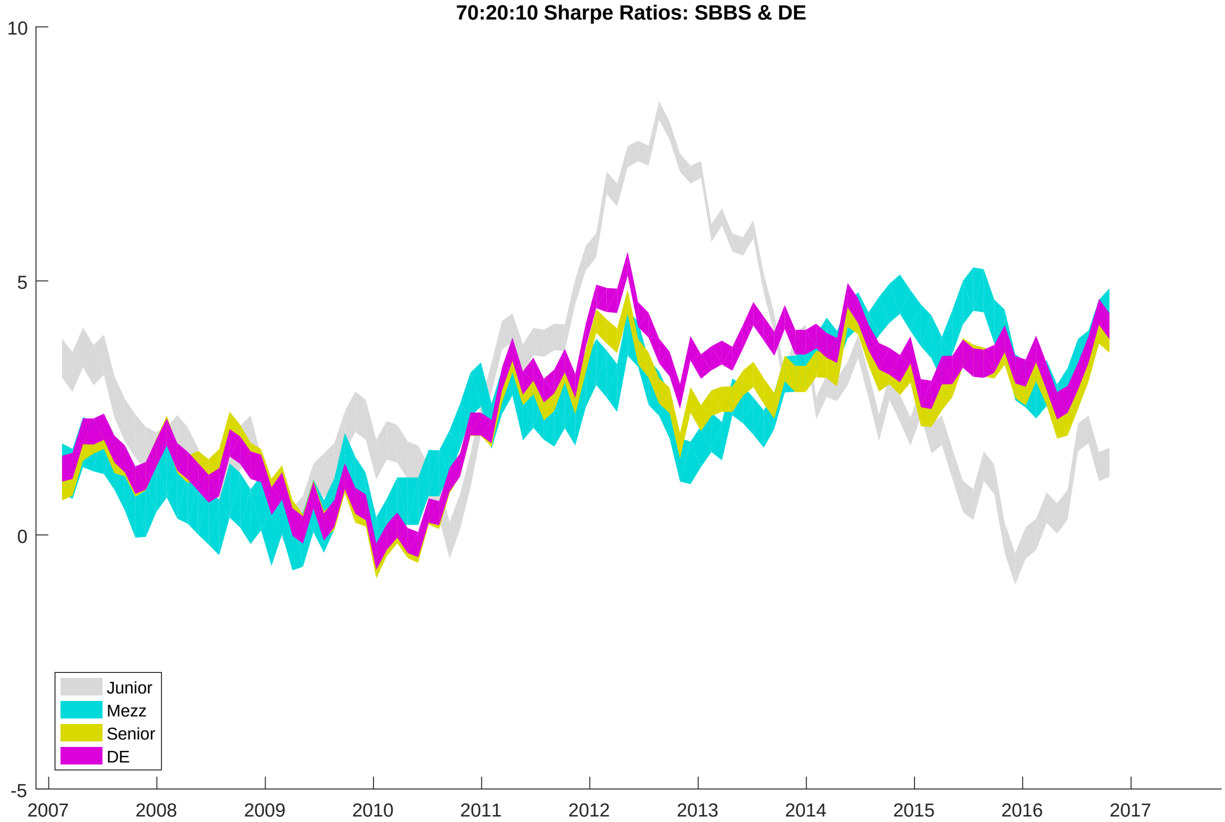

Figure 11: This figure displays the dynamic Sharpe ratio associated with monthly holding of senior, mezzanine and junior tranches of the SBBS and of German sovereign bonds (these are all for the 10 year maturity and for the 70:20:10 tranche structure). In each case we allow the coupon to have an upper and lower bound of one standard deviation around the chosen rate based on the standard error of the intercept coefficient in the regression of the coupon spread on yield spread as discussed in the methodology section.

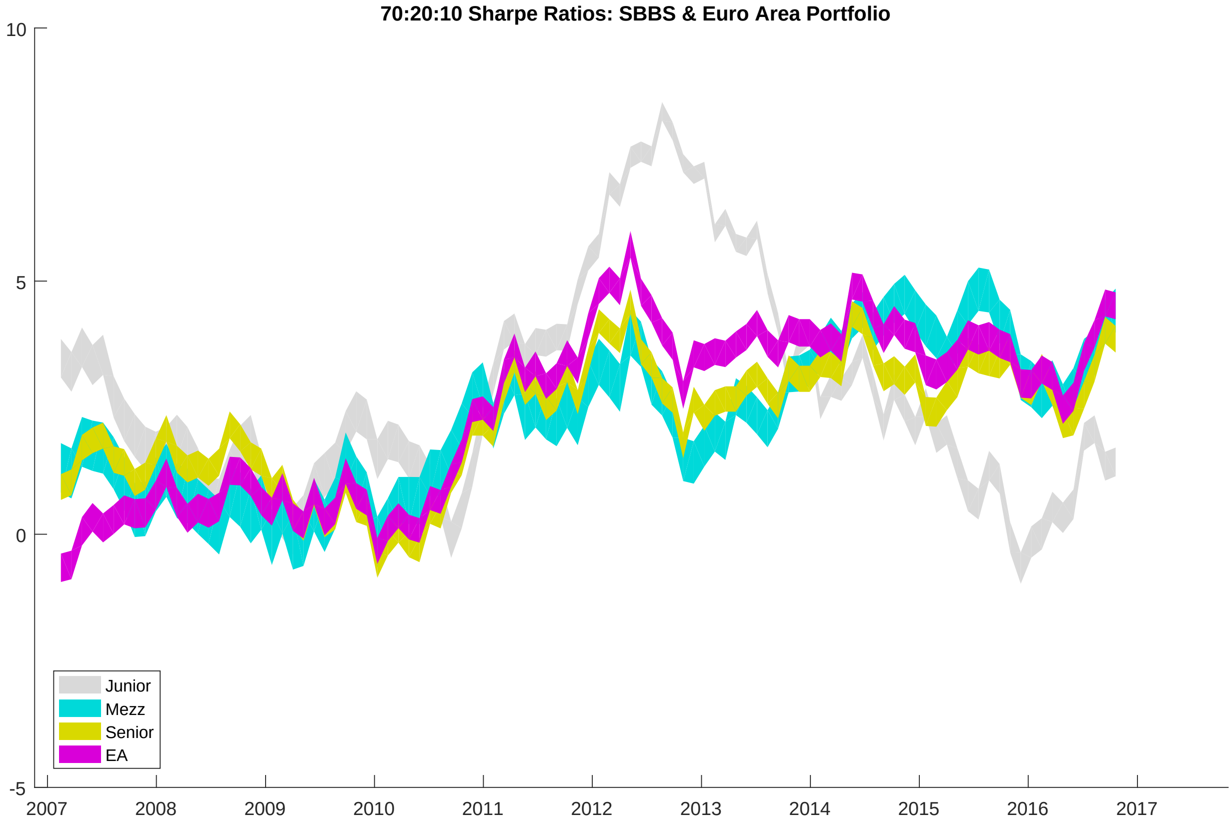

Figure 12: This figure displays the dynamic Sharpe ratio associated with monthly holding of senior, mezzanine and junior tranches of the SBBS and of the GDP weighted euro area portfolio of sovereign bonds (these are all for the 10 year maturity and for the 70:20:10 tranche structure). In each case we allow the coupon to have an upper and lower bound of one standard deviation around the chosen rate based on the standard error of the intercept coefficient in the regression of the coupon spread on yield spread as discussed in the methodology section.

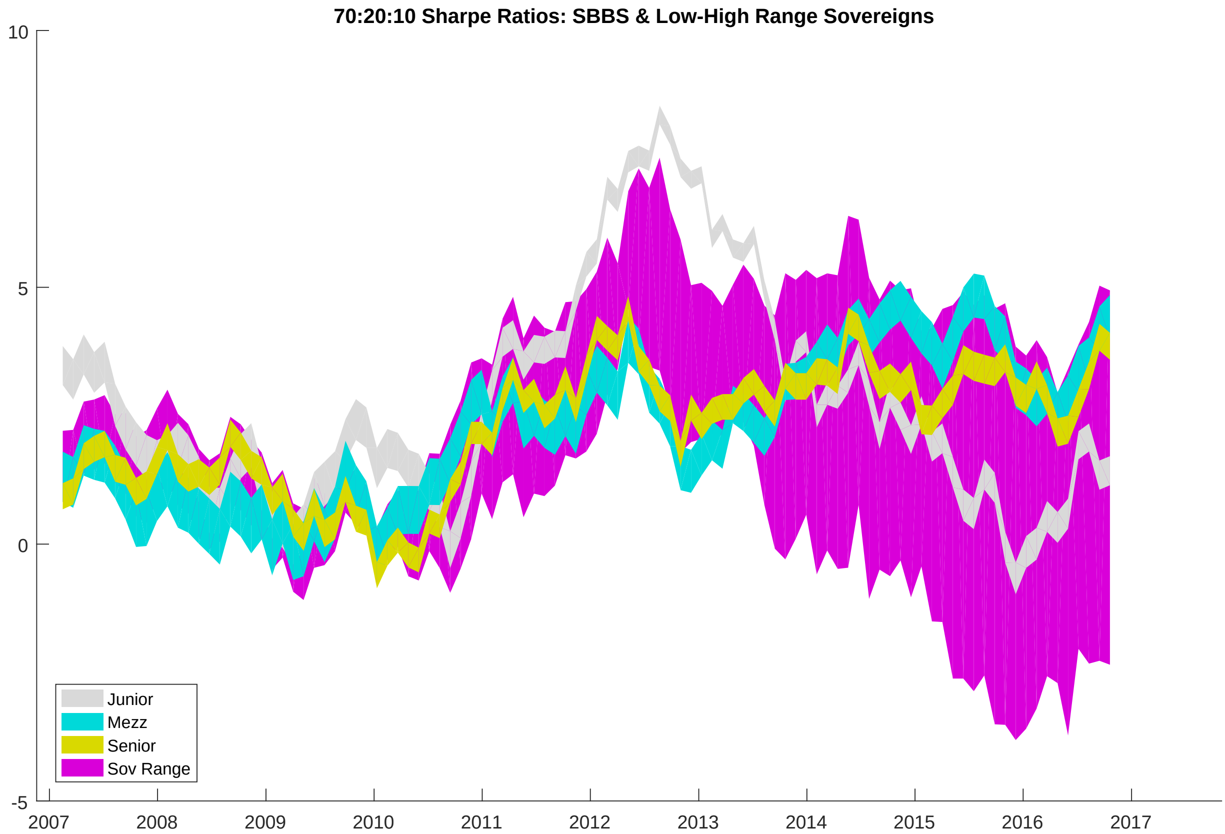

Figure 13: This figure displays the dynamic Sharpe ratio associated with monthly holding of senior, mezzanine and junior tranches of the SBBS (these are all for the 10 year maturity and for the 70:20:10 tranche structure). The purple shaded area shows the area between the lower and upper bounds of observed Sharpe ratios across the euro area individual sovereigns. In each case we allow the coupon to have an upper and lower bound of one standard deviation around the chosen rate based on the standard error of the intercept coefficient in the regression of the coupon spread on yield spread as discussed in the methodology section.

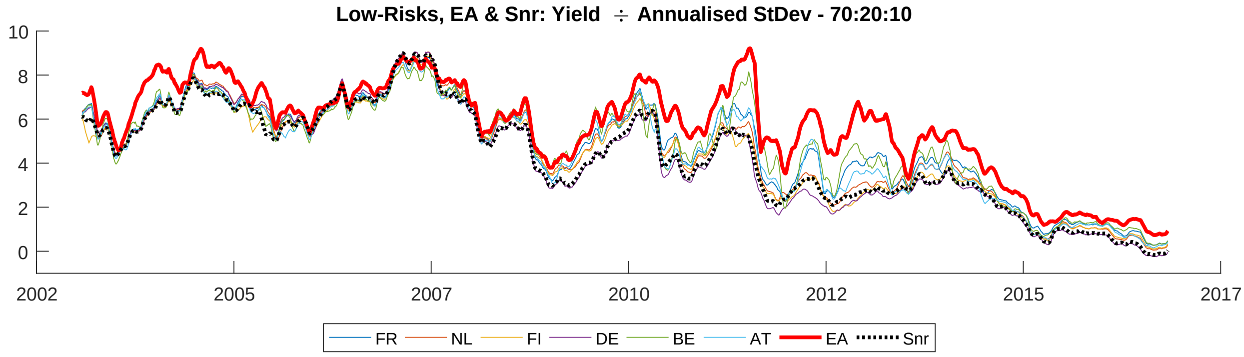

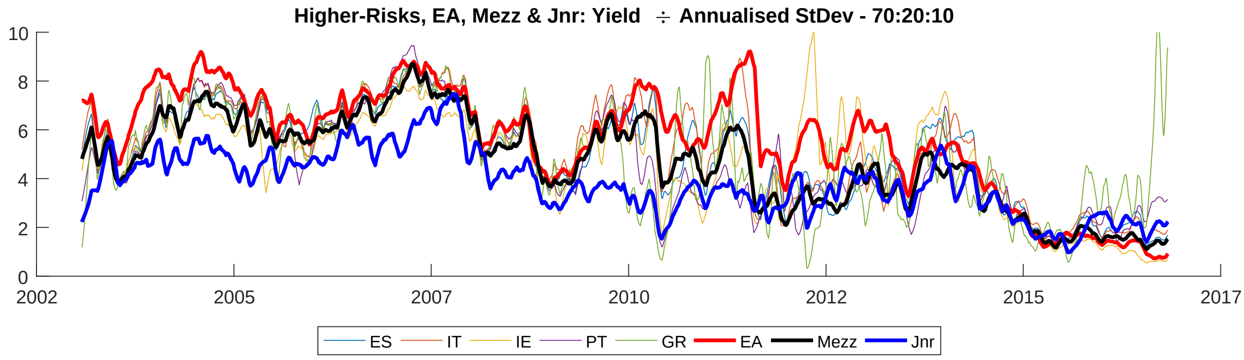

Figure 14: Dynamic Quasi-Sharpe Ratio: Yield for Annualised Risk. In this case risk is measured as GJR-GARCH Conditional Volatility.

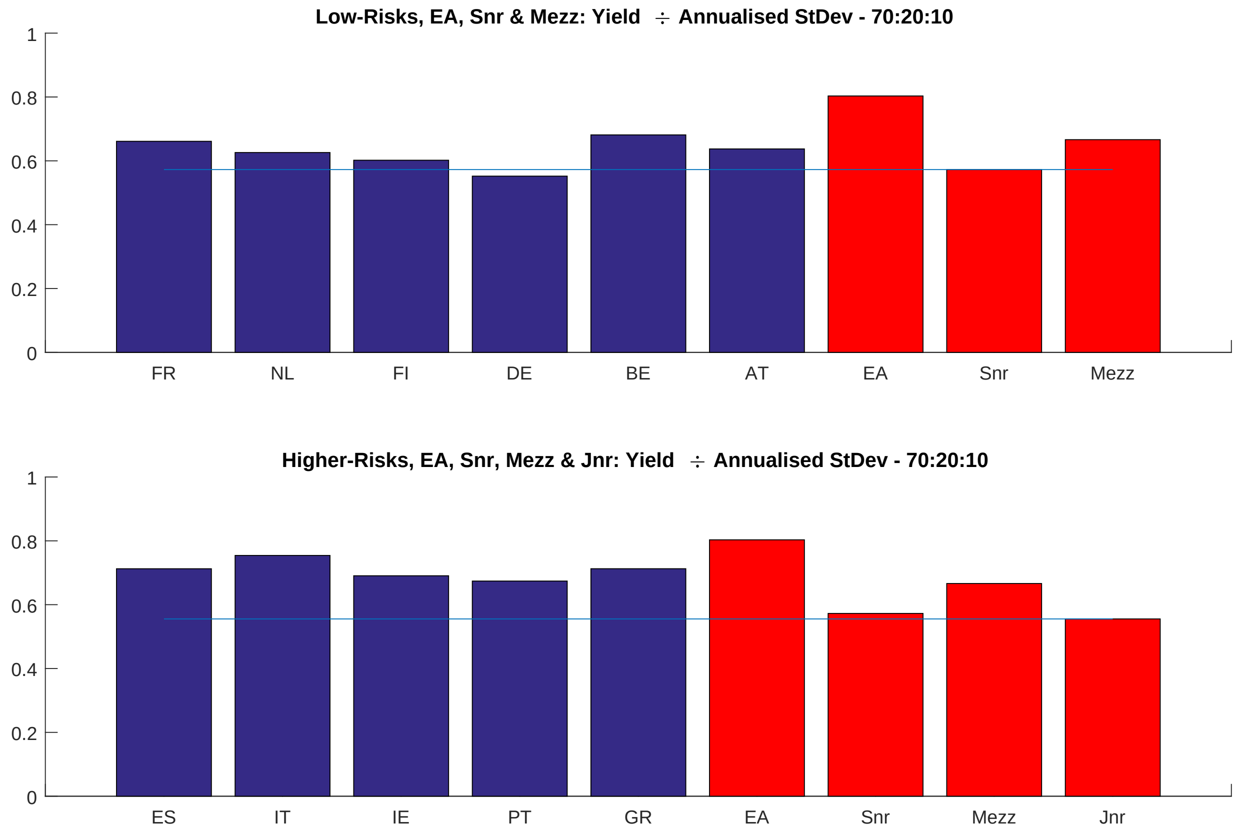

Figure 15: Average Quasi-Sharpe Ratio: Yield for Annualised Risk. In this case risk is measured as the average of the GJR-GARCH Conditional Volatility across the sample.

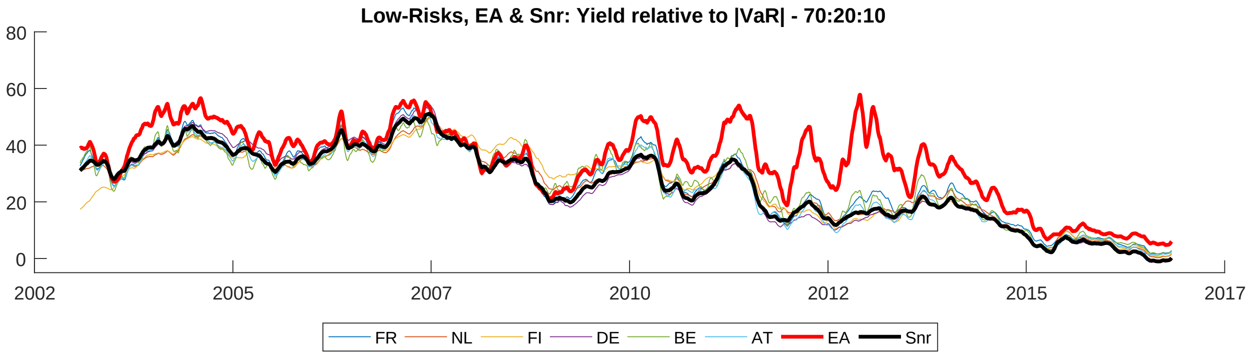

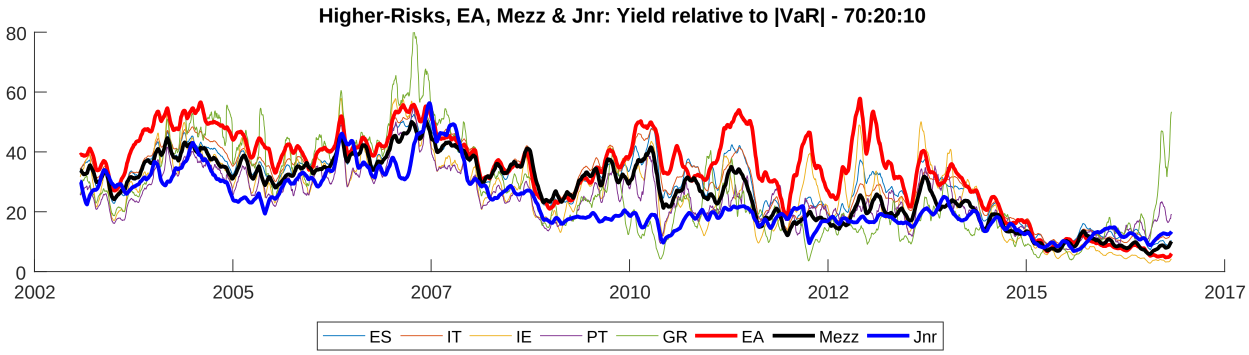

Figure 16: Dynamic Yield for Tail Risk: $\text{YTM}_t \div \text{Abs}(\text{VaR})_t$. The graphed observations have been smoothed by taking the centered moving average over a rolling 20 observations.

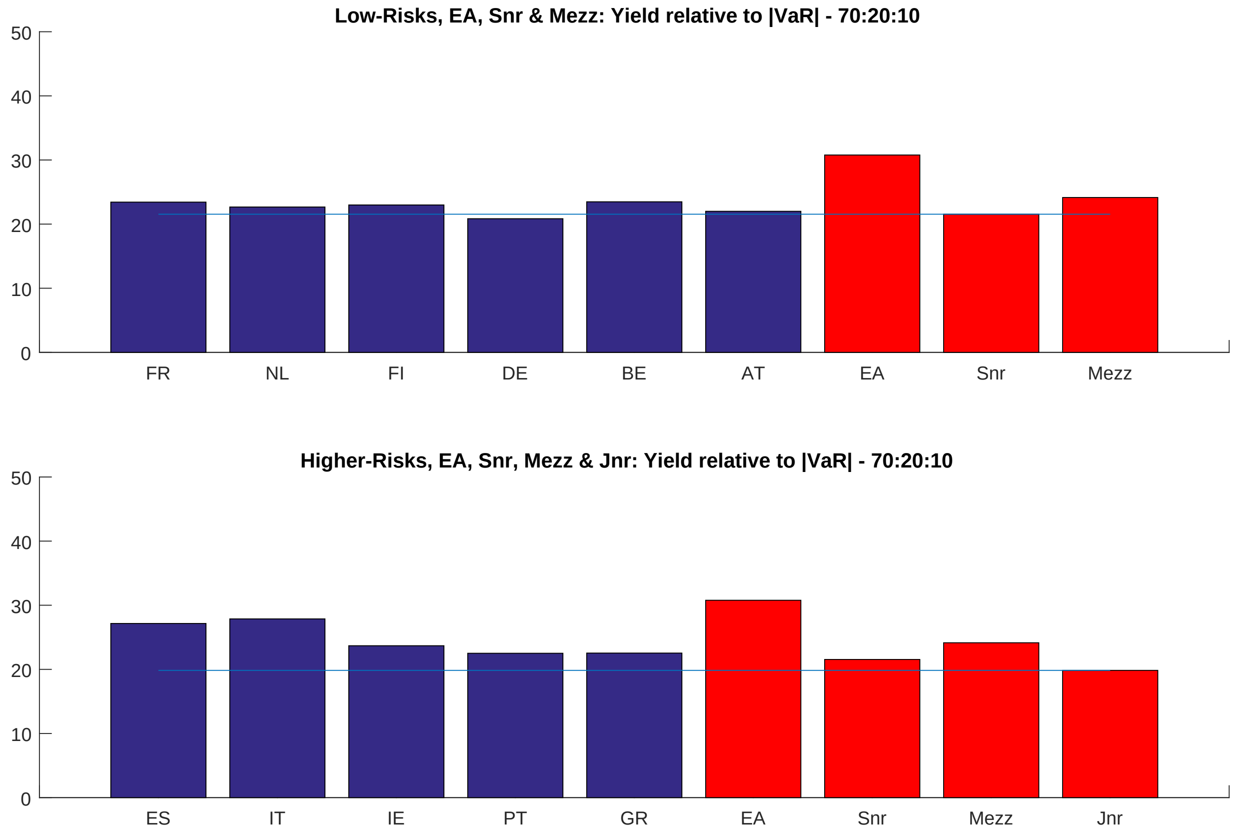

Figure 17: Average Yield for Tail Risk: Average ($\mathrm{YTM}_t \div \mathrm{Abs}(\mathrm{VaR})_t$).

A Appendix - GARCH Diagnostics

Table 1: **Joint insignificance of ARCH in standardised residuals.** **Backlink to page 18.**

	Lags				Lags		
Case	1	2	3	Case	1	2	3
AT	0.85	0.97	0.98	IE	0.28	0.52	0.68
BE	0.58	0.56	0.62	IT	0.58	0.68	0.81
DE	0.03	0.10	0.16	NL	0.32	0.31	0.41
ES	0.19	0.43	0.63	PT	0.01	0.02	0.05
FI	0.34	0.55	0.76	SEN	0.12	0.25	0.22
FR	0.24	0.50	0.70	MEZZ	0.05	0.14	0.08
GR	0.98	0.80	0.30	JUN	0.48	0.54	0.00
EU	0.95	0.33	0.46	MKT	0.83	0.97	0.10

Note: P-values are shown for Langrange-Multiplier test statistics (robust to heteroscedasticity) where the null is joint insignificance of autocorrelation in squared standardised residuals from GARCH models used in MES estimation. Test are presented for the inclusion of 1 lag, lags 1 and 2 and lags 1 to 3. The test is applied for the cases of each country, each of the SBBS tranches in the 70:20:10 SBBS structure and for an EU and EA/MKT portfolio of sovereigns (the latter is used as the market portfolio). Tranche names are abbreviated as follows: Sen(Senior), Mezz(Mezzanine), Jun(Junior). Country-specific results are abbreviated as: AT(Austria), BE(Belgium), DE(Germany), ES(Spain), FR(France), GR(Greece), IE(Ireland), IT(Italy), NL(Netherlands), PT(Portugal).

B Appendix - VAR-for-VaR Diagnostics

Table 2: **Diagnostics for Conditional Quantile Estimates** **Backlink to page 19.**

Figure	Case	DQ(p)	Exceed(%)	Cross(t)	Cross(p)
1	JUN	0.53	0.99	26.19	0.00
1	MEZZ	0.10	0.99	26.19	0.00
2	PT	0.52	0.99	52.56	0.00
2	IE	0.48	0.99	21.11	0.00
2	GR	0.47	0.93	13.95	0.01
3	DE	0.20	1.06	0.34	0.99
3	SEN	0.16	0.99	0.18	0.99
4	FR	0.12	0.99	1.19	0.88
4	NL	0.53	0.99	0.78	0.94
4	AT	0.99	0.99	74.25	0.00
4	FI	0.16	0.99	10.20	0.04
4	BE	0.46	0.93	2.01	0.73
5	IT	0.99	0.93	19.90	0.00
5	ES	0.46	0.93	95.49	0.00

Note: SBBS tranches are abbreviated as Sen(Senior), Mezz(Mezzanine) and Jun(Junior) and these are associated with the 70:20:10 SBBS design. Country-specific results are abbreviated as follows: AT(Austria), BE(Belgium), DE(Germany), ES(Spain), FI(Finland), FR(France), GR(Greece), IE(Ireland), IT(Italy), NL(Netherlands), PT(Portugal). The first column of the table indicates the figure where the VaR is displayed. The second column indicates the country or SBBS that is under analysis. Column 3 shows the p-value for the DQ test. Column 4 shows the percentage of exceedances of the estimated VaR (this should be close to 1%). A joint test-statistic for the joint significance of the cross-effect parameters and the associated p-values are provided in the final two columns.

Imprint and acknowledgements

We are grateful to Sam Langfield, Marco Pagano, Richard Portes, Philip Lane, Javier Suárez, Kitty Moloney, Patrick Haran, Angelo Ranaldo, Giovanna Bua, Joe McNeill, Spyros Alogoskoufis and members of the ESRB High-Level Task Force on Safe Assets for helpful comments. This work also benefited from comments following presentation at the Annual Irish Economic Association Conference, 2017.

Maite De Sola Perea
National Bank of Belgium, Brussels, Belgium; email: **maite.desolaperea@nbb.be**

Peter G. Dunne
Central Bank of Ireland, Dublin, Ireland; email: **peter.dunne@centralbank.ie**

Martin Puhl
Oesterreichishe Nationalbank, Vienna, Austria; email: **martin.puhl@aon.at**

Thomas Reininger
Oesterreichishe Nationalbank, Vienna, Austria; email: **thomas.reininger@oenb.at**

Postal address	60640 Frankfurt am Main, Germany
Telephone	+49 69 1344 0
Website	**www.esrb.europa.eu**

Note:
The views expressed in ESRB Working Papers are those of the authors and do not necessarily reflect the official stance of the ESRB, its member institutions, or the institutions to which the authors are affiliated.

ISSN	2467-0677 (pdf)
ISBN	978-92-9472-017-7 (pdf)
DOI	10.2849/651144 (pdf)
EU catalogue No	DT-AD-18-003-EN-N (pdf)