Risk endogeneity at the lender-/investor-of-last-resort*

Diego Caballero,[(a)] *André Lucas,*[(b)] *Bernd Schwaab,*[(a)] *Xin Zhang,*[(c)]

[(a)] European Central Bank [(b)] Vrije Universiteit Amsterdam & Tinbergen Institute

[(c)] Sverige Riksbank, Research Division

***** preliminary conference submission *****

January 30, 2018

Abstract

The riskiness of a central bank's balance sheet depends on the financial health of its counterparties, which in turn depends on the central bank's liquidity provision and asset purchases. This two-way dependence is particularly pronounced during a liquidity crisis. We propose a novel framework to study the time-variation in central bank portfolio credit risks associated with different conventional and unconventional monetary policy operations. The framework accommodates numerous bank and sovereign counterparties, fat tails, skewness, and time-varying dependence parameters. In an application to items from the European Central Bank's weekly balance sheet between 2009 and 2015, we find that unconventional monetary policy measures tended to generate beneficial risk spill-overs across operations. Some were 'self-financing' in risk terms.

Keywords: Credit risk; risk measurement; central bank; lender-of-last-resort; score-driven model.

JEL classification: G21, C33.

*Author information: Diego Caballero, European Central Bank, Sonnemannstrasse 22, 60314 Frankfurt, Germany, Email: diego.caballero@ecb.int. André Lucas, Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, Email: a.lucas@vu.nl. Bernd Schwaab, European Central Bank, Sonnemannstrasse 22, 60314 Frankfurt, Germany, Email: bernd.schwaab@ecb.int. Xin Zhang, Sverige Riksbank, Brunkebergstorg 11, 103 37 Stockholm, Sweden, Email: xin.zhang@riksbank.se. Parts of this paper were written while Schwaab was on secondment to the ECB's Risk Management Directorate (D-RM). The views expressed in this paper are those of the authors and they do not necessarily reflect the views or policies of the European Central Bank or the Sverige Riksbank.

1 Introduction

During the global financial crisis between 2008 and 2010 and the euro area sovereign debt crisis between 2010 and 2012, severe liquidity squeezes and market malfunctions forced the European Central Bank (ECB) to act as a lender-of-last-resort (LOLR) to the entire financial system. Large-scale central bank lending to banks ensured the proper functioning of the financial system and, with it, the transmission of monetary policy. Such lending occurred mainly via ECB main refinancing operations (MROs), multiple long-term refinancing operations (LTROs) with maturities of up to one year, and two very-long-term refinancing operations (VLTROs) with a three-year maturity, all backed by repeated expansions of the set of eligible collateral. In addition, the ECB also acted as an investor-of-last-resort (IOLR). For example, it purchased sovereign bonds in illiquid secondary markets within its Securities Markets Programme (SMP) between 2010 and 2012, and committed to doing so again under certain circumstances within its Outright Monetary Transactions (OMT) program as announced in August 2012.

The ECB's actions as a lender- and investor-of-last-resort during the financial crisis and euro area sovereign debt crisis had a first-order impact on the size, composition, and, ultimately, the credit risk of its balance sheet. At the time, its policies raised substantial concerns that the central bank may be taking excessive risks (and supporting moral hazard by helping troubled banks). Particular concern emerged that the materialization of credit risk could impair the central bank's reputation, credibility, independence, and ultimately its ability to steer inflation towards its target of close to but below 2% over the medium term. The credit risk concerns were so pronounced that at the time some media reports referred to the ECB unflatteringly as the ECBB: Europe's Central Bad Bank (see e.g. Brendel and Pauly, 2011).

It is uncontroversial that lending freely in line with Bagehot (1873)-inspired principles during a liquidity crisis and purchasing credit-risky bonds during a severe sovereign debt crisis can increase the credit risk of a central bank's balance sheet. How different lender- and investor-of-last-resort policies interact from a risk perspective, however, is currently less clear. Specifically, we ask: Can central bank liquidity provision or asset purchases during a financial crisis be 'self-financing' in net risk terms? This could happen if risk-taking in one part of the balance sheet (e.g., more asset purchases) de-risks the other balance sheet positions (e.g., the collateralized lending book) by a commensurate amount, or more. How economically important are such risk spillovers across policies? Were the Eurosystem's financial buffers at all times sufficiently high to match its portfolio risk? And, from a methodological point-of-view, given that a central bank's actions affect its counterparties' point-in-time risks as well as their risk interdependence, how can a central bank assess the riskiness of its balance sheet in real time at a high frequency?

This paper proposes a novel credit risk measurement framework that allows us to study these questions. Our framework is based on a tractable high-dimensional dependence function that can accommodate a large number of bank and sovereign counterparties. The dependence (t-copula) model allows us to accommodate extreme tail dependence (fat joint tails), time-varying volatility and correlation parameters, as well as a potential asymmetry in the correlation dynamics that implies skewness also in the unconditional dependence of risks. Our model combines elements of the models in Creal, Koopman, and Lucas (2011), Lucas, Schwaab, and Zhang (2014), and Lucas, Schwaab, and Zhang (2017), and is here extended to accommodate asymmetric correlation dynamics and a large number of counterparties.

Banks, by engaging in maturity transformation, are by their very nature exposed to liquidity shocks. Central banks are uniquely able to provide liquidity-support in a crisis owing to the fact that they are never liquidity-constrained in the currency they issue; see e.g. Bindseil and Laeven (2017). In a seminal work, Bagehot (1873) argued that such support should be against good collateral, only to solvent banks, and at penalty rates when compared to normal times. Illiquidity-related credit risks may then dissolve quickly. To track a central bank's risk positions in real time requires a risk measurement framework that quickly adapts to changes in priced risks as well as changes in risk dependence. This paper introduces such a modeling framework.

We apply our model to six monetary policy operations from the European Central Bank's weekly balance sheet between 2009 and 2015. Point-in-time risk measures are obtained from either Moody's Analytics (for banks) or are inferred from CDS spreads (for sovereigns). All risk model parameters are estimated by the method of maximum likelihood. Portfolio risk measures are subsequently obtained through simulation. We compare the model-implied portfolio credit risks shortly before and after key policy announcements during the crisis, and report the time differences for different monetary policy operations. The 'high-frequency' (weekly) assessment allows us to identify the effect of each announcement on the portfolio credit risks associated with the main conventional and unconventional monetary policy operations. To distinguish size from composition effects we study portfolio credit risks both in absolute terms as well as in percent of total assets.

We find that LOLR- and IOLR-implied credit risks were almost always negatively related in our sample. Taking risk in one part of the central bank's balance sheet (e.g., the announcement of asset purchases within the SMP) de-risked other positions (e.g., the collateralized lending book from previous LTROs). Vice versa, the allotment of each VLTRO reduced the expected losses from the SMP portfolio at that time. Risk spillovers between monetary policy operations are economically significant, and are remarkably similar in sign and magnitude around the time of the policy announcements.

Finally, some unconventional policy measures were (almost) self-financing in net risk terms. For example, the announcement of OMT de-risked *all* other portfolios by an economically significant amount. We find that the initial OMT announcement de-risked the Eurosystem's balance sheet by €-9.1 bn in 95%-Expected Shortfall (ES), and that the anoucement of OMT technical details in September 2012 were associated with a reduction in 95% ES of €-36.5 bn. As another example, the announcement of the SMP on 10 May 2010, and the purchases that followed immediately afterwards, raised the 95% ES of SMP asset holdings from zero to approximately €8.5 bn, but de-risked the collateralized lending book such that the overall increase in 95% ES was significantly smaller (€1.3 bn). Our findings can have important implications for the design of central banks' post-crisis operational frameworks, and inform a debate on how to balance the need for a lender-/investor-of-last-resort during liquidity crises with recent banking-sector regulations that seek to lower the frequency of such crises.

Our study informs at least four directions of current research. First, several studies investigate the central bank's role of LOLR and IOLR during a financial crisis. Important contributions include Bagehot (1873), Allen and Gale (2000), and Rochet and Vives (2004). Freixas et al. (2004) provide a survey; see also Bindseil (2014) for a textbook treatment.

Second, a nascent strand of literature applies stress-testing methods to central banks' assets and income. Carpenter et al. (2013) and Greenlaw et al. (2013) stress-test the Federal

Reserve's ability to send positive remittances to the U.S. Treasury given that a large-scale sovereign bond portfolio exposes the Fed (and thus indirectly the Treasury) to interest rate risk. Christensen, Lopez, and Rudebusch (2015) advocate the use of probability-based stress tests, and find that the risk of suspended Fed remittances to the Treasury is small but non-negligible (at about 8%).

Third, we effectively apply 'market risk' methods to solve a 'credit risk' problem. As a result, we connect a growing literature on non-Gaussian volatility and dependence models with another growing literature on portfolio credit risk and loan loss simulation. Time-varying parameter models for volatility and dependence have been considered, for example, by Engle (2002), Demarta and McNeil (2005), Creal et al. (2011), Zhang et al. (2011), and Engle and Kelly (2012). At the same time, credit risk models and portfolio tail risk measures have been studied, for example, by Vasicek (1987), Lucas et al. (2001), Lucas et al. (2003), Gordy (2000), Gordy (2003), Koopman et al. (2011), Koopman et al. (2012), and Giesecke et al. (2014). We argue that our combined framework yields the best of these two worlds: portfolio credit risk measures (say, one year ahead) that are available at a market risk frequency (such as daily or weekly) for real time portfolio credit risk monitoring and impact assessment.

Finally, to introduce time-variation into our empirical model specification we endow our model with observation-driven dynamics based on the score of the conditional predictive log-density. Score-driven time-varying parameter models are an active area of research, see for example Creal, Koopman, and Lucas (2011), Creal, Koopman, and Lucas (2013), Harvey (2013), Oh and Patton (2014), Creal, Schwaab, Koopman, and Lucas (2014), Harvey and Luati (2014), Andres (2014), and more.[1] For an information theoretical motivation for the use of score driven models, see Blasques, Koopman, and Lucas (2015).

The remainder of the paper is set up as follows. Section 2 presents our data and reviews a subset of ECB unconventional monetary policy operations between 2008 and 2015. Section 3 introduces our high-dimensional credit risk measurement framework. Section 4 applies the framework to a subset of the ECB's weekly balance sheet. Section 5 concludes. A Web Appendix presents additional results and technical details.

[1] We refer to `http://www.gasmodel.com` for an extensive enumeration of recent work in this area.

2 Data

We are interested in studying the time variation in central bank portfolio credit risks, with a particular focus on such risks just before and after six key policy announcements. The announcements are related to three ECB unconventional monetary policy operations: The SMP, two VLTROs, and the OMT. This section discusses these programs and the relevant announcement dates. We then discuss our risk data.

2.1 ECB conventional and unconventional monetary policies

The ECB adjusts the money supply in the euro area mainly via so-called refinancing operations. ECB refinancing operations between 2009 and 2015 included main refinancing operations (MROs), long-term refinancing operations (LTROs), very-long-term refinancing operations (VLTROs), and targeted long-term refinancing operations (TLTROs).

Before the onset of the global financial crisis in 2007, the MROs and three-month LTROs were sufficient to steer short-term interest rates, to manage aggregate liquidity, and to signal the monetary policy stance in the euro area. Following the onset of the global financial crisis, however, the ECB was forced to significantly extend the scale and maturity of its operations (LTROs and VLTROs). TLTROs were set up in June 2014 mainly to subsidize bank lending to the non-financial sector.

During the sovereign debt crisis between 2010 and 2012 the ECB also conducted asset purchases within its SMP program. These purchases were not intended to affect the money supply. For this reason they were sterilized at the time.

Figure 2.1 plots selected items of the ECB's weekly balance sheet between 2009 and 2015. We distinguish five different liquidity operations: MRO, LTRO<1y, LTRO1y, VLTRO3y, and TLRTO. The figure also plots the total asset held in the SMP portfolio. Clearly, the ECB's balance sheet varied in size, composition, and thus credit riskiness during the course of the global financial crisis and euro area sovereign debt crisis. A peak in total assets was reached at the hight of the debt crisis in mid-2012, at approximately €1.5 trn, following two VLTROs and SMP government bond purchases.

The remainder of this section discusses three key non-standard monetary policy operations in chronological order: the SMP, the VLTROs, and the OMT. Each of these had a

Figure 1: ECB collateralized lending and SMP exposures

Total collateralized lending exposures associated with different liquidity operations (MRO, LTRO<1y, LTRO1y, VLTRO3y, TLTRO), as well as sovereign bond holdings from purchases within the Securities Markets Programme (SMP). Vertical axes are in trillion euro. Data is weekly between 2009 and 2015. The SMP1 horizontal line refers to the initial announcement of the SMP on 10 May 2010. SMP2 marks the cross-sectional extension of the program on 08 August 2011. VLTRO1 marks the allotment of the first three-year VLTRO on 20 December 2011. VLTRO2 marks the allotment of the second three-year VLTRO on 28 February 2012. OMT1 marks the initial announcement of the OMT on 02 August 2012. OMT2 marks the announcement of the OMT's technical details on 06 September 2012.

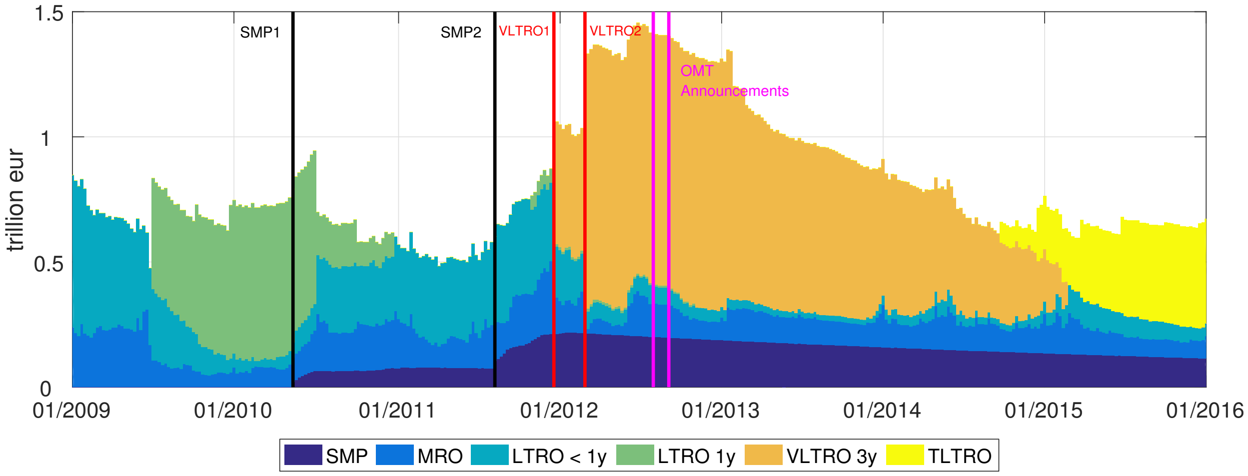

substantial impact on asset prices, point-in-time credit risks, and time-varying risk correlations; see e.g. ECB (2014) for a survey.

The SMP

The SMP was announced on 10 May 2010, with the objective to help restoring the monetary policy transmission mechanism by addressing the mal-functioning of certain government bond markets. The SMP consisted of interventions in the form of outright purchases which were aimed at improving the functioning of these bond markets by providing "depth and liquidity"; see González-Páramo (2011). Implicit in the notion of market mal-functioning is the notion that government bond yields can be unjustifiably high and volatile. For example, market-malfunctioning can reflect the over-pricing of risk due to illiquidity as well as contagion across countries, see Constâncio (2011).

SMP interventions occurred in government debt securities markets between 2010 and 2012 and initially focused on three euro area countries: Greece, Ireland, and Portugal. The SMP was extended to include Spain and Italy on 08 August 2011. Approximately €214 billion (bn) of bonds were acquired between 2010 and early 2012; see ECB (2013). At the end of 2012, the ECB held approximately €99.0bn in Italian sovereign bonds, €30.8bn in Greek debt, €43.7bn in Spanish debt, €21.6bn in Portuguese debt, and €13.6bn in Irish bonds; see the ECB (2013) Annual Report. For impact assessments of SMP purchases on bond yields, CDS spreads, and liquidity risk premia we refer to e.g. Eser and Schwaab (), Ghysels, Idier, Manganelli, and Vergote (2017), and De Pooter, Martin, and Pruitt (2018).

The VLTROs

Two large-scale VLTROs were announced on 08 December 2011, and subsequently allotted to banks on 21 December 2011 and 29 February 2012. The first installment provided more than 500 banks with €489 bn to at a low (1%) interest rate for an exceptionally long period of three years. The second installment in 2012 was even larger, and provided more than 800 euro area banks with €530 billion in three-year low-interest loans. By loading up on VLTRO funds stressed banks could make sure they had enough cash to pay off their own maturing debts, and at the same time keep operating and lending to the non-financial sector.

Figure 2: ECB collateralized lending across countries

Exposures across different euro area countries from five liquidity-providing operations; see Figure . The vertical axis is in trillion euro. Vertical lines indicate the events described in Figure 2.1. Data is weekly between 2009 and 2015.

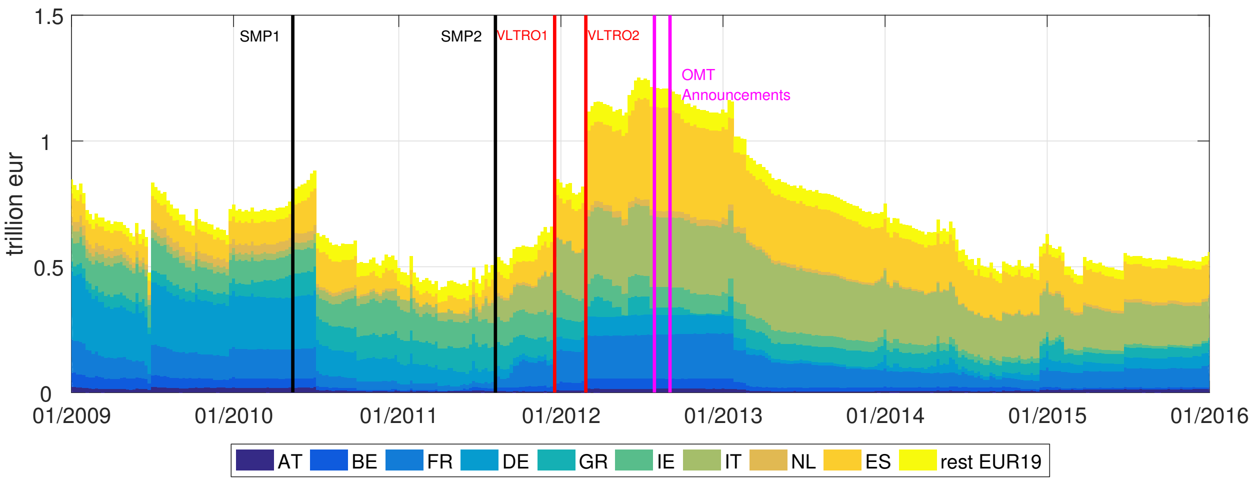

Incidentally, banks used some of the money to also load up on domestic government bonds, temporarily bringing down sovereign yields. This eased the debt crisis, but also worsened banks-sovereign risk dependence; see Acharya and Steffen (2015).

Figure 2 plots the ECB's country-level collateralized lending exposures, aggregating over five liquidity-providing operations; see Figure . The largest share of VLTRO funds was tapped by banks in Italy and Spain, and also Greece, Ireland, and Portugal. These sovereigns (and their banks) were perceived by markets to be particularly affected by the euro area sovereign debt crisis. Banks from non-stressed countries such as Germany and France relied less heavily on VLTRO and other ECB funding during the crisis.

The OMT

On 26 July 2012, the president of the ECB pledged to do "whatever it takes" to preserve the euro, and that "it will be enough". Communication announcing Outright Monetary Transactions (OMT), a new conditional asset purchase program, followed shortly afterwards on 02 August; see ECB (2012). The OMT technical details were announced on 06 September 2012. The OMT replaced the SMP. Within the OMT, the ECB could potentially undertake purchases ("outright transactions") in secondary euro area sovereign bond markets provided certain conditions are met. OMT interventions were stipulated to be potentially limitless, to focus on short-maturity bonds, and were conditional on the bond-issuing countries agreeing to and complying with certain domestic economic measures. In the years since its inception, the OMT never had to be used. Nevertheless, its announcement is widely credited for ending the acute phase of the sovereign debt crisis by restoring confidence; see e.g. Wessel (2013) and ECB (2014).

2.2 Risk data

We rely on expected default frequency (EDF) data from Moody's Analytics, formerly Moody's KMV, when assigning point-in-time probabilities of default (pds) to ECB bank counterparties. EDFs are point-in-time forecasts of physical default hazard rates, and are based on a proprietary firm value model that takes firm equity values and balance sheet information as inputs; see Crosbie and Bohn (2003) for additional details. EDF data are standard default risk measurements and are routinely used in the financial industry and credit risk literature; see for example Lando (2003), Duffie et al. (2007) and Duffie et al. (2009).

EDF measures are available for listed banks only. On the other hand, many ECB bank counterparties are not listed. We overcome this problem by using EDF-based country-level indexes of point-in-time risk. When measuring bank risk at the country level, we thus implicitly assume that the (weighted average of) point-in-time risks as calculated for listed banks are also approximately appropriate for non-listed banks.

Unfortunately, EDF measures do not exist for sovereigns. We therefore need to infer physical probabilities of default from observed sovereign CDS spreads. Web Appendix A provides the details of our approach. We first invert a CDS pricing formula to calculate one-year ahead point-in-time risk-neutral default probabilities. We then convert these into physical measures using a nonlinear mapping from the literature.

3 Statistical model

3.1 Portfolio risk measures

Credit losses at time $t = 1, \ldots, T$ over a one-year-ahead horizon are only known with certainty after that year has passed, and thus uncertain (random) at time t. The probability distribution of ex-ante credit losses is a key concern for risk measurement. We model total credit losses $\ell_t(k)$ associated with potentially many counterparties $i = 1, \ldots, N_t(k)$ as

$$\ell_t(k) = \sum_{i=1}^{N_t(k)} \ell_{it}(k) = \sum_{i=1}^{N_t(k)} \mathrm{EAD}_{it} \cdot \mathrm{LGD}_{it} \cdot 1(\tilde{y}_{it} < \tau_{it}), \qquad (1)$$

where $k = 1, \ldots, K$ denotes a certain monetary policy operation (e.g., MRO lending or SMP asset holdings), ℓ_{it} is the counterparty-specific loss, EAD_{it} is the exposure-at-default associated with counterparty i, $\mathrm{LGD}_{it} \in [0, 1]$ is the loss-given-default as a fraction of EAD_{it}, and $1(\tilde{y}_{it} < \tau_{it})$ is the default indicator function that takes the value of one if and only if the log-asset value of firm i falls below a counterparty-specific default threshold τ_{it}; see e.g. Merton (1974) and CreditMetrics (2007). The loss $\ell_{it}(k)$ is random because it is a function of three random terms, EAD_{it}, LGD_{it}, and the default indicator. Total losses from monetary policy operations are given by $\ell_t = \sum_{k=1}^{K} \ell_t(k)$.

Portfolio risk measures are typically moments or (average) quantiles of the ex-ante loss distribution. We focus on standard risk measures such as the Expected Loss $\mathrm{EL}(k)_t$, Value-at-Risk at a given confidence level γ VaR_t^γ, and Expected Shortfall ES_t^γ at a one-year-ahead horizon. These risk measures are given by

$$\begin{aligned}
&\mathrm{EL}(k)_t = \mathrm{E}\left(\ell_t(k)\right) \\
&\Pr\left(\ell_t(k) \geq \mathrm{VaR}^\gamma(\ell_t(k))\right) = \gamma \\
&\mathrm{ES}(k)_t^\gamma = \mathrm{E}\left(\ell_t(k) \,|\, \ell_t(k) \geq \mathrm{VaR}^\gamma(\ell_t(k))\right),
\end{aligned}$$

where $\mathrm{E}(\cdot)$ is the expectation over all sources of randomness (e.g., defaults and LGDs). $\mathrm{ES}(k)_t^\gamma$ is often interpreted as the "average VaR in the tail", and is typically more sensitive to the shape of the tail of the loss distribution. The subscript t indicates that the time series of portfolio risk measures is available at a much higher than annual frequency (e.g., weekly).

The remainder of this section reviews the modeling of the ingredients of (1) from right to left: dependent defaults, LGD, and EAD.

3.2 Copula model for dependent defaults

Our model for dependent defaults follows closely from the models developed in Lucas, Schwaab, and Zhang (2014, 2017). To tailor these frameworks to the problem at hand, however, we need to extend the model to accommodate a potentially large number of bank and sovereign counterparties and seek to capture any asymmetry in the copula in a computationally straightforward way.

Following the seminal Merton (1974) and CreditMetrics (2007) frameworks, we assume that a counterparty i defaults if and only if its log asset value $\tilde{y}_{it}$ falls short of a default threshold τ_{it}. This event happens with a time-varying default probability p_{it}, where

$$p_{it} = P(\tilde{y}_{it} < \tau_{it}) = F(\tau_{it}) \Leftrightarrow \tau_{it} = F^{-1}(p_{it}), \tag{2}$$

where F is the CDF of $\tilde{y}_{it}$. We stress that, unlike in the Merton (1974) model, p_{it} is an *observed* input in our model. In fact, τ_{it} is chosen at each point in time and for each counterparty such that the marginal default probability implied by the multivariate (copula) model coincides with the observed market-implied default probability for that counterparty at the time; see the last equality in (2).

When modeling dependent defaults, we assume that (changes in) one-year ahead log-asset values $\tilde{y}_{it}$ are generated by a high-dimensional multivariate Student's t-distribution $\bar{p}(\tilde{y}_t; \theta_t)$, where

$$\tilde{y}_{it} = \mu_{it} + \sqrt{\zeta_t} L_{it}^{(k)} z_{it}, \quad i = 1, \ldots, N_t(k), \tag{3}$$

and where $z_t^i \sim \mathrm{N}(0, \mathrm{I}_N)$ is a vector of standard normal risk terms, $L_{it}^{(k)}$ is the ith row of $L_t^{(k)}$, $L_t^{(k)}$ is the Choleski factor of the Student's t covariance matrix $\Omega_t^{(k)} = L_t^{(k)} L_t^{(k)\prime}$, $\zeta_t \sim \mathrm{IG}(\frac{\nu}{2}, \frac{\nu}{2})$ is an inverse-gamma distributed scalar mixing variable that generates the fat tails in the copula, and ν is a degrees of freedom parameter to be estimated. The covariance matrix $\Omega_t^{(k)}$ depends on k because different counterparties participate in different monetary policy operations. We fix $\mu_{it} = 0$ in (3) without loss of generality since quantiles shift linearly with the mean.

Rather than modeling a potential asymmetry in default dependence via the copula (3), we introduce asymmetry in a novel way via the transition equation governing the correlation parameters. This has the advantage that the quantiles of a standard t-density can be used in the estimation of the copula parameters. Student-t quantiles are almost always tabulated and thus quickly available in standard software packages. By contrast, quantiles of skewed densities such as the Generalized Hyperbolic skewed-t density used in Lucas, Schwaab, and Zhang (2014) are not usually tabulated and need to be solved for numerically. Repeated numerical integration within a line search is time-consuming in high-dimensional applications, and in practise not always reliable for counterparties with very low default probabilities (when τ_{it} is far in the tail).

3.3 Score-driven copula dynamics

The covariance matrix $\Omega_t^{(k)}$ is typically of a high dimension. For example, $N_{it}(k) > 800$ banks participated in the ECB's second VLTRO program. The high dimensions and time-varying size of $\Omega_t^{(k)} \in \mathbb{R}^{N_{it}(k) \times N_{it}(k)}$, however, imply that it is difficult to model directly.

Our empirical application below uses time-varying block equi-correlations within and across countries when constructing $\Omega_t^{(k)}$. For other applications based on dynamic (block-) equicorrelation models see e.g. Engle and Kelly (2012) and Lucas, Schwaab, and Zhang (2017). The block equicorrelation assumption allows us to specify $\Omega_t^{(k)}$ as a function of a much smaller covariance matrix Σ_t which is independent of k. The covariance matrix Σ_t in turn depends on a vector of latent correlation factors f_t. Specifically, $\Omega_t^{(k)} = \Omega_t^{(k)}\left(\Sigma_t(f_t)\right)$, where $\Sigma_t\left(f_t\right) \in \mathbb{R}^{D \times D}$, and $D < N_t(k)$.

The mapping of matrix elements $\Omega_t^{(k)}(i,j) = \Sigma_t\left(l(i), m(j)\right)$ is surjective but not injective. I.e., any element of Σ_t typically appears multiple times in $\Omega_t^{(k)}$. All bank correlation pairs across countries can be take from Σ_t. The within-country correlation pairs (the off-diagonal elements in the diagonal blocks of $\Omega_t^{(k)}$), however, cannot be read off Σ_t. We therefore need to make an assumption, and take the within-country bank correlations as equal to the maximum row entry of Σ_t. I.e., banks within a country are as correlated as the maximum estimated correlation pair across borders at that time.

The positive definiteness of $\Omega_t^{(k)}$ thus constructed is not guaranteed. However, a close approximation to $\Omega_t^{(k)}$ almost always is positive definite. In case the Choleski decomposition of $\Omega_t^{(k)}$ fails we take the Schur decomposition matrix and set the upper-triangular elements to zero. We check that these entries are numerically small (below 10^{-7}).

Our approach for modeling $\Sigma_t(f_t)$ mirrors the approach of Creal, Koopman, and Lucas (2011). In this framework, factors $f \in R^{D(D-1)/2\times 1}$ have an interpretation as 'angles' in a hyper-geometric space. This setup ensures that Σ_t is always positive definite and symmetric. Each element of f_t maps into exactly one correlation pair in Σ_t, and thus typically into multiple elements in $\Omega_t^{(k)}$.

The dynamics of f_t are specified by the transition equation

$$f_{t+1} = \omega + A \cdot S_t \nabla_t + B \cdot f_t + C \cdot \mathrm{Lev}_t \cdot s_t, \tag{4}$$

where ω, $A = A(\theta)$, $B = B(\theta)$ and $C = C(\theta)$ are parameters and matrices to be estimated, $\nabla_t = \partial \ln \bar{p}(y_t|f_t;\theta)/\partial f_t$ is the score of a multivariate Student's t density, the scaling matrix is chosen as the inverse conditional Fisher information matrix $S_t = \mathrm{E}_{t-1}\left[\nabla_t \nabla_t'\right]^{-1}$, and

$$\mathrm{Lev}_t = \mathrm{vech}\left[1(y_{t\,\cdot} > 0) \cdot 1(y_{t\,\cdot} > 0)'\right], \tag{5}$$

is a "leverage" term that allows for an asymmetric response in the correlation dynamics. To see this, note that when $C > 0$ and element $\left(\Sigma_t(f_t)\right)_{jk}$ is monotonically increasing in f_t for all j, k, then the dependence in the copula increases more in response to unexpectedly rising marginal risks ("bad" shocks) than falling marginal risks ("good" shocks). The asymmetry in the conditional correlations carries over to skewness in the unconditional distribution of $\tilde{y}_{it}$. This is analogous to a GARCH model with leverage; see e.g. Glosten, Jaganathan, and Runkle (1993).

All univariate (volatility) modeling precedes the fitting of the copula. We refer to Web Appendix B for the univariate modeling of the marginal risks. The score ∇_t and scaling function S_t in (4) are available in closed form. We refer to Web Appendix C for the detailed expressions.

3.4 Loss-given-default

Portfolio risk levels depend substantially on the assumptions made in the modeling of the loss (fraction) given default. We distinguish two separate cases: bank and sovereign counterparties.

Collateralized lending to banks within the ECB's liquidity facilities implies a double recourse. If a bank defaults, the central bank can access the pledged collateral and sell it in the market to cover its losses. Conservatively calibrated haircuts on the market value of pledged assets ensure that a sufficient amount of collateral is almost always available to cover losses. Haircuts are higher for more volatile (longer duration) and more credit-risky claims. For example, non-marketable assets carry valuation haircuts of up to 65%. As a result, historical counterparty-level LGDs have been approximately zero for most central banks, owing to conservative risk management frameworks and haircuts.

The case of Lehman brothers can serve as an (extreme) example. Its German subsidiary, Lehman Brothers Bankhaus, defaulted on the Eurosystem on 15 September 2008. In the weeks leading up to the default, complicated asset-backed-securities had been posted as collateral. These were highly non-liquid and non-marketable at the time. In addition, an untypically large amount of central bank liquidity had been withdrawn just just prior to the default. Even so, the posted collateral was ultimately sufficient to recover all losses. The workout-LGD was zero as a result; see e.g. Dombret (2017).

A substantial loss to the central bank may occur, however, in extreme scenarios when both the bank and the pledged collateral default simultaneously. This was a valid concern during the sovereign debt crisis, when a subset of banks pledged bonds issued by their domestic government, or bank bonds that were eligible only because they were also government-guaranteed. This exposed the cental bank to "wrong way risk", as bank and sovereign risks are typically positively dependent in the data.

We capture this concern as follows. For a bank counterparty i, we model LGD stochastically as

$$\text{LGD}_{it} = 0.05 + 0.55 \cdot 1\left\{\tilde{y}_{jt} < y^*_{jt};\text{ for some country } j\right\}.$$

I.e., $\text{LGD}_{it} = 0.05$ if bank counterparty i defaults but no sovereign j defaults. The LGD increases to 60% if bank i defaults and a euro area sovereign were to default as well (i.e., in the same simulation) and a sovereign debt crisis were to ensue as a result. The 5% value for LGD is on the high side, as explained above. The 60% stressed LGD is conservative as well, and chosen in line with international evidence on sovereign bond haircuts; see e.g. Cruces and Trebesch (2013, Table 1).

In case of a sovereign counterparty, e.g. for government bonds acquired within the SMP, only a single recourse applies. We therefore set the LGD to 60% should such a default be observed.

More elaborate specifications for LGD are clearly possible. The present approach, however, is parsimonious and transparent, while still sufficiently flexible to capture the issue of systematic variation in LGDs; see also Creal et al. (2014).

3.5 Exposures-at-default

Exposures-at-default $\mathrm{EAD}_{it}(k)$ in (1) do not have to coincide with currently observed exposure $\mathrm{EXP}_{it}(k)$. EADs may in fact be higher, as bank counterparties in distress tend to draw significantly on the central bank's liquidity facilities in the days and weeks leading up to a default. We capture this concern to some extent by specifying

$$\mathrm{EAD}_{it}(k) = \mathrm{EXP}_{it}(k) \cdot G_{it}(k), \tag{6}$$

where $G_{it}(k)$ is gamma-distributed random variable with parameters such that $\mathrm{E}[G_{it}] = 1$ and $\mathrm{V}[G_{it}] = 6$. These values imply that exposures-at-default equal current exposures in expectation, but with a pronounced skewness towards larger values in the simulations. We matched the variance and skewness of $G_{it}(k)$ to a small number of observed euro area defaults in a confidential dataset.

We do not have access to all counterparty-specific exposures (loan amounts) over time in our sample. Instead, we have access only to the aggregate exposures at the country $\times$ operation $\times$ week level. In addition, we have access to the number of banks $N_t(j, k)$ that accessed a monetary policy operation at any week in country j. We therefore proceed under the assumption that exposures $a_{i,kt}$ for $i = 1, \ldots, N_t(j, k)$ within country j are Pareto-distributed, see e.g. Gabaix (2009), and draw these exposures according to $P(a_{i,jkt}) \propto (a_{i,jkt})^{-r}$ for a given value of r. We obtain counterparty-specific current exposures as

$$\text{EXP}_{i,jkt} = \frac{a_{i,jkt}}{\sum_{i=1}^{N_t(j,k)} a_{i,jkt}} \cdot \text{CountryExp}_{jkt}^{\text{observed}}, \tag{7}$$

thus dividing up aggregate bank lending volumes per country and policy operation to $N_t(j,k)$ banks. We infer r from the cross-section of total bank assets in the euro area between 2009 and 2015 using the Hill (1976) estimator, and fix $r = 2.3$ to construct the relative shares.

3.6 Parameter estimation

Observation-driven multivariate time series models are attractive because the log-likelihood is known in closed form. Parameter estimation is standard as a result. This is a key advantage over alternative parameter-driven risk frameworks, as e.g. considered in Koopman et al. (2011), Koopman et al. (2012), and Azizpour et al. (2017), for which the log-likelihood is not available in closed form and parameter estimation is non-standard. For a given set of observations $y_1, \ldots, y_T$, the vector of unknown parameters $\theta = \{\omega, A, B, C, \nu\}$ can be estimated by maximizing the log-likelihood function with respect to θ, that is

$$\hat{\theta} = \arg \max_{\theta} \sum_{t=1}^{T} \ln \bar{p}(y_t | f_t; \theta), \tag{8}$$

where $\bar{p}(y_t|f_t;\theta)$ is implied by (3), and where y_t are observed log-changes in bank and sovereign EDF measures. The evaluation of $\ln \bar{p}(y_t|f_t;\theta)$ is easily incorporated in the filtering process for f_t as described in Section 3.3.

The maximization in (8) can be carried out using a conveniently chosen quasi-Newton optimization method that is based on score information. The score here is defined as the first derivative of the log-likelihood function in (8) with respect to the constant parameter vector θ. Analytical expressions for the score function can be developed, but typically lead to a collection of complicated equations. In practice, the maximization of the log-likelihood function is therefore carried out using numerical derivatives.

For the empirical application, we reduce the computational burden of parameter estimation in two ways. First, we proceed in two steps and estimate the parameters of all D

marginal models before estimating the parameters of the copula model. Copula model parameters are estimated based on the probability integral transforms of the de-volatized data from the first step. Lucas et al. (2014) discuss one-step and two-step estimation approaches in a related setting, and find that they lead to similar estimates if T is large. Second, we assume that matrices A, B, and C in the factor transition equation (4) are scalars, such that $\theta = \{\omega, A, B, C, \nu\}$ is of a relatively low dimension.

4 Empirical results

Our empirical application to weekly ECB balance sheet data revolves around three questions. First, what were the expected 'risk cost associated with ECB unconventional monetary policy measures during the crisis? Addressing this question allows us to compare different policy operations in terms of their risk efficiency. Second, have Eurosystem financial buffers been sufficiently high to withstand the materialization of a bad outcome at high confidence levels? Finally, how economically significant were spillovers across different unconventional monetary policies? Addressing the last question allows us to study the extent to which central banks can influence their balance sheet risks by impacting point-in-time risks and risk correlations.

4.1 Parameter estimates and model selection

[To be added.]

4.2 Expected losses

A large literature studies the impact of ECB unconventional monetary policy measures between 2008 and 2012 on financial markets as well as macroeconomic outcomes; see e.g. ECB (2014) for a survey. The 'costs' of these measures, e.g. in terms of increased balance sheet size and risk, have received relatively less attention.

Figure 3 plots the one-year-ahead expected losses from ECB collateralized lending operations (top panel) and SMP asset purchases (bottom panel). The expectation of the loss density is calculated by simulation, using 200,000 draws at each time t. For each simulation, we keep track of exceedances of $\tilde{y}_{i,t}$ below their respective calibrated thresholds at time t as

Figure 3: Expected losses from collateralized lending and SMP purchases

The top and bottom panels plots the expected losses from liquidity providing operations and the SMP asset holdings, respectively. The vertical axis is in billion euro. Data is weekly between 2009 and 2015. The vertical lines mark the events described in Figure 2.1 and Section 2.1.

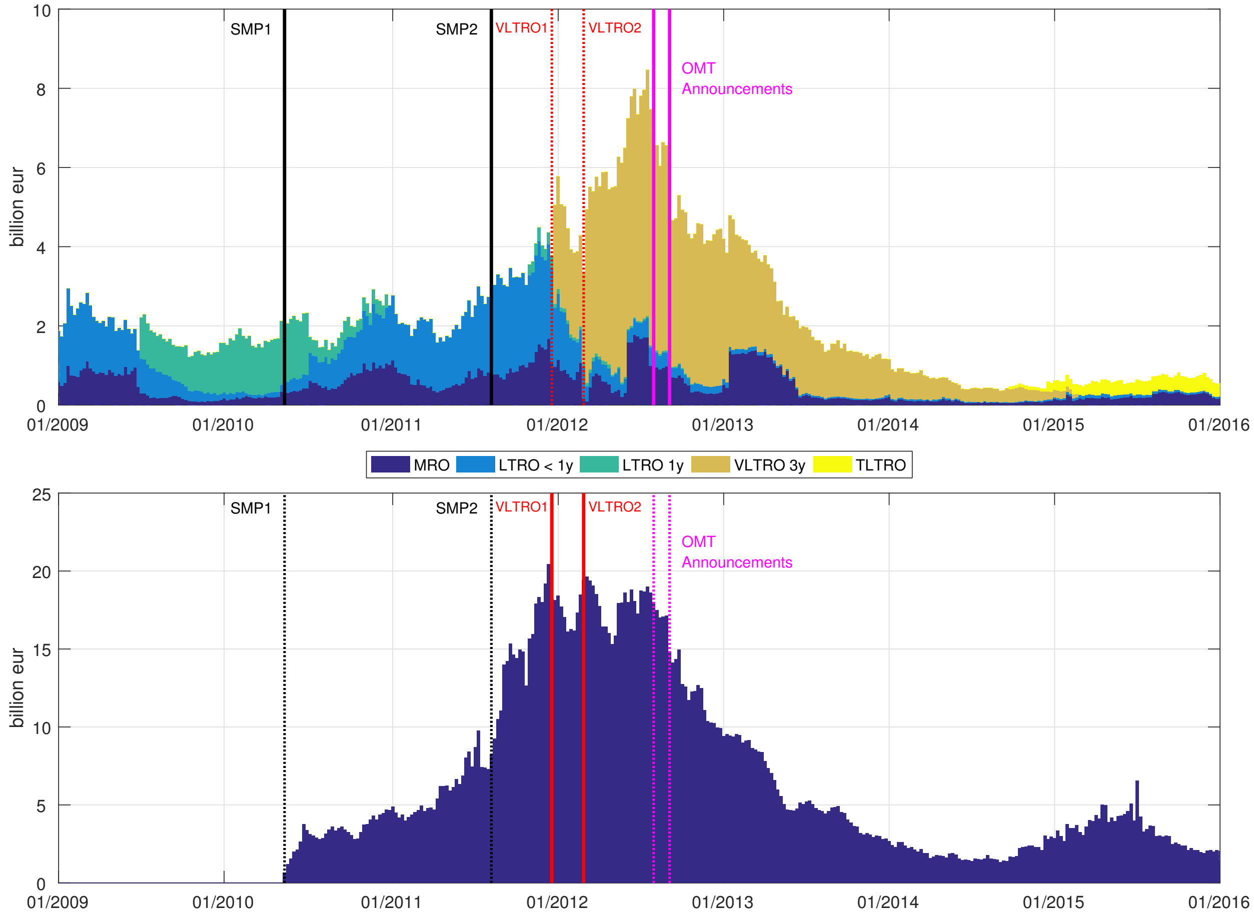

well as outcomes for LGD and EXP, as described in Section 3. The mean estimates combine all marginal risk estimates and all 14(14-1)/2=109 time-varying correlation estimates into a single time series plot.

The expected losses reflect, first, a clear deterioration of debt conditions since the beginning of the global financial crisis in 2008 and the euro area sovereign debt crisis in the spring of 2010, and second, a clear turning of the tide around mid-2012. Vertical lines indicate the announcement of the SMP on 10 May 2010 and its cross-sectional extension on 08 August 2011, the allotment of two VLTROs on 21 December 2011 and 28 February 2012, and two announcements regarding the Outright Monetary Transactions (OMT) in August and September 2012. Expected losses are additive across monetary policy operations, and therefore stacked vertically in Figure 3. Expected losses for both collateralized lending and SMP exposures peak in mid-2012 at around approximately €8 bn and €20 bn, respectively.

Figure 3 already hints at the presence of beneficial spillovers across monetary policy operations. For example, the announcement of each VLRTO appears to have had a beneficial but temporary impact on the expected losses associated with the SMP portfolio. Similarly, the OMT announcements appear to have had a pronounced impact on the expected losses associated with the collateralized lending operations and the SMP. Section 4.4 below studies spillovers in an event-study setting.

Risk estimates are also a prerequisite for evaluating policy measures in terms of their 'risk efficiency'. Risk efficiency is the notion that a certain amount of expected policy impact should be achieved with the minimum level of risk possible. Given an estimate of policy impact (e.g., a change in break-even inflation rates around the time of a policy announcement) and an estimate of risk (e.g., the expected losses as plotted in Figure 3), policy measures could be evaluated ex-post by scaling the former by the latter. [More to be added.]

4.3 Tail risk and financial buffers

This section discusses our estimates of portfolio credit risk. Figures 4 and Figure 5 plot the 95% ES associated with five collateralized lending operations and the SMP portfolio, respectively. The ES peaks at approximately €120 bn for the collateralized lending operations and approximately €110 bn for the SMP holdings. Absolute levels of risk correlate with the magnitude of exposures as reported in Figure 2.1.

Tail risks depend on balance sheet composition as well. For example, Figure 5 suggests that the 95%-ES for SMP assets increased in absolute terms (top panel) but decreased relative to total holdings (bottom panel) after the program was extended in 2011 to include Italian and Spanish sovereign bonds. These bonds were perceived by markets to be relatively less credit-risky than the Irish, Portuguese, and Greek bonds purchased earlier.

Given the time variation in the 95% ES it is natural to ask whether the Eurosystem – the ECB together with its 19 national central banks – was at all times sufficiently able to withstand a materialization of a ES95%-sized bad outcome of the magnitude reported in Figure 5. For a commercial bank, financial buffers typically include items such as the current year's (projected) annual income, revaluation reserves in the balance sheet, general risk provisions, and paid-in equity capital. We adopt the same notion of financial buffers for a central bank, even though the notion of buffers against default are somewhat inappropriate in this case. (Recall that a central bank is never liquidity constrained in the currency they issue.) For central banks, revaluation (silent) reserves are an important term. For example, gold holdings are typically valued at historical cost and are worth more at current market prices.

Based on publicly available data, Buiter (2013) calculates that the Eurosystem revaluation reserves alone exceed €300 bn at any time between 2009 and 2012. This suggests that the Eurosystem's buffers were sufficient to withstand the portfolio tail risks as reported in Figures 4 and 5 at any time during both the financial crisis as well as the euro area sovereign debt crisis. [More to be added.]

Figure 4: 95% Expected Shortfall for collateralized lending exposures

The top panel plots the 95% ES for five ECB liquidity-providing monetary policy operations. The bottom panel plots the 95% ES scaled by the respective exposures for each policy. Scaled expected losses are plotted for comparison. The vertical axes are in billion euro. Data is weekly between 2009 and 2015. The vertical lines mark the events described in Figure 2.1 and Section 2.1.

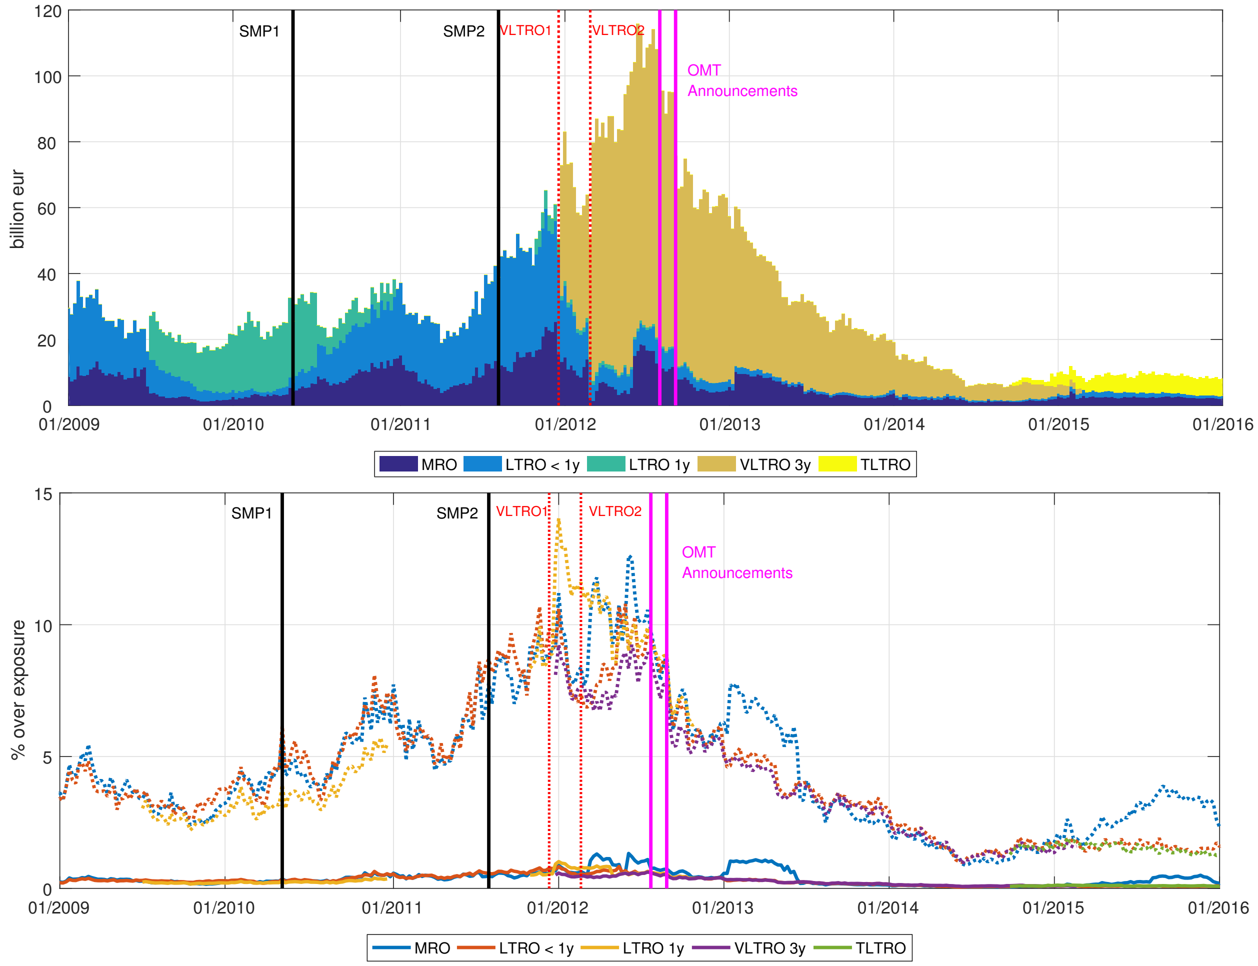

Figure 5: 95%-expected shortfall for SMP portfolio, in absolute values and in per cent of total holdings

The top panel plots the 95% ES for the SMP asset holdings, respectively. The vertical axis is in billion euro. Data is weekly between 2009 and 2015. The vertical lines mark the events described in Figure 2.1 and Section 2.1. The bottom panel plots the 95% ES and EL in per cent of SMP holdings at the time.

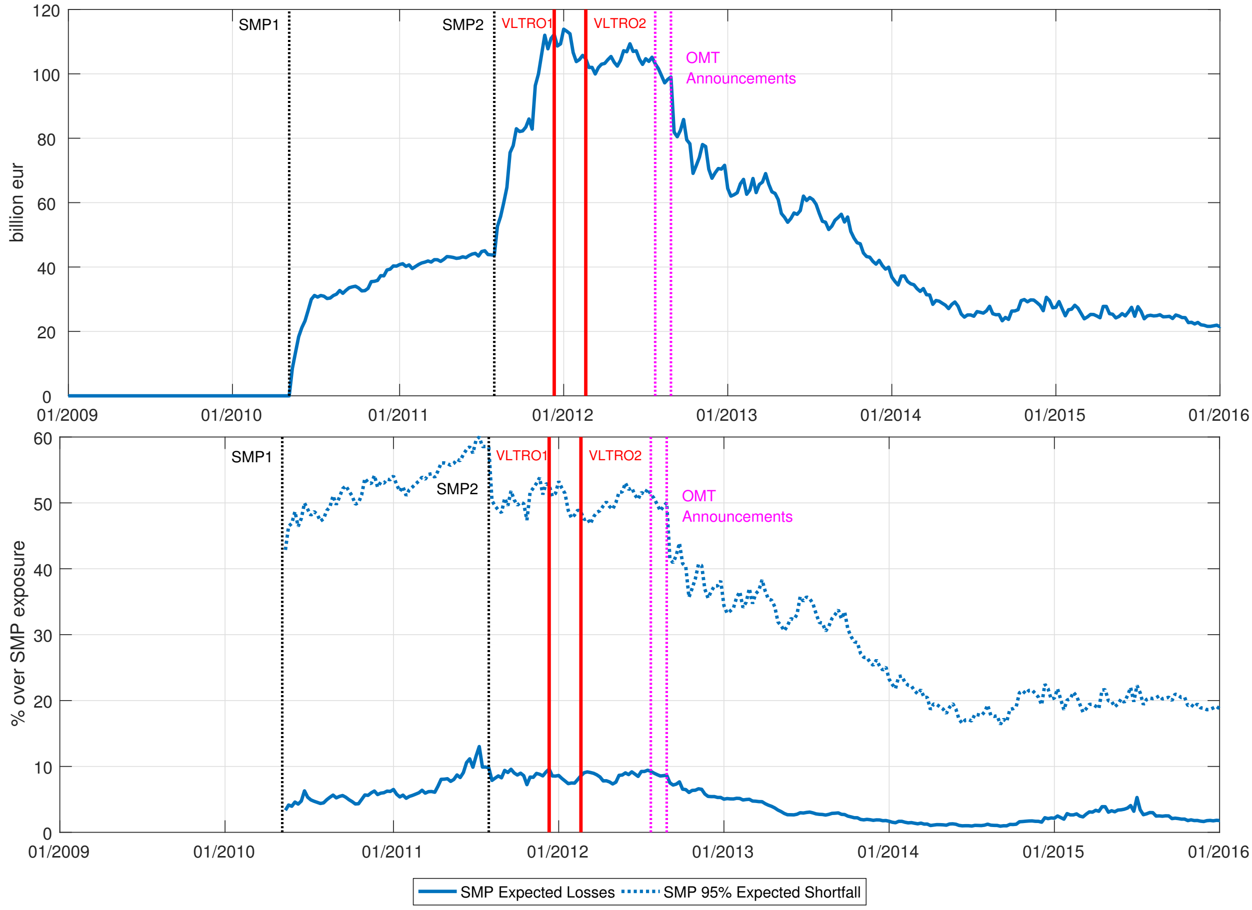

4.4 Risk spillovers from non-standard monetary policies

The riskiness of any central bank's balance sheet depends on the financial health of its counterparties, which in turn depends on the central bank's liquidity provision and asset purchases, particularly during a liquidity crisis. This section applies our modeling framework to five key ECB monetary policy operations, with a focus on the portfolio risks shortly before and after the announcements. This 'high-frequency' (weekly) assessment allows us to identify the impact of each announcement on the portfolio risk associated with each operation.

Table 1 reports one-year ahead portfolio credit risk estimates (EL and 95% ES) around six policy announcements: two SMP-related announcements on 10 May 2010 and 08 August 2011, two VLTRO-related announcements on 20 December 2011 and 20 February 2012, and two OMT-related announcement on 02 August 2012 and 06 September 2012. We find that LOLR- and IOLR-implied risks are almost always negatively related. Specifically, taking risk in one part of the central bank's balance sheet (e.g., the announcement of asset purchases within the SMP) de-risked other positions (e.g., collateralized lending within previous LTROs). Vice versa, the allotment of each VLTRO in 2011 and 2012 reduced the portfolio credit risk estimates associated with the SMP portfolio.

A subset of unconventional monetary policies were (almost) self-financing in net risk terms. For example, the 2010 SMP announcement and the initial SMP purchases raised the corresponding 95% ES by €8.5 bn, but de-risked the collateralized lending book to such an extend that the net increase in 95% ES was much less, at approximately €1.3 bn. The announcement of OMT de-risked all other portfolios by an economically significant amount. The first OMT announcement de-risked the Eurosystem's balance sheet by €-9.1 bn (95%-ES). The annoucement of OMT technical details in September 2012 were associated with an even stronger reduction of €-36.5 bn in 95% ES.

Finally, the risk spillovers between monetary policy operations are remarkably similar in sign (negative) and magnitude. Spillovers were almost always statistically (see Monte Carlo standard error bands) and economically significant.

Table 1: Portfolio credit risks around key policy announcements

Portfolio credit risks for different monetary policy operations around six policy announcements: the SMP announcement on 10 May 2010, the cross-sectional extension of the SMP on 08 August 2011, the allocation of the first VLTRO on 20 December 2011 and of the second VLTRO on 20 February 2012, OMT announcement on 02 August 2012, and the announcement of the OMT's technical details on 06 September 2012.

SMP1	07/05/2010		14/05/2010			
	EL	ES(95%)	EL	ES(95%)	ΔEL	ΔES(95%)
SMP	-	-	0.7	8.5	**0.7**	**8.5**
			[0.7 0.7]	[8.5 8.6]		
MRO	0.3	5.2	0.3	4.3	-0.0	-1.0
	[0.3 0.3]	[5.1 5.6]	[0.3 0.3]	[4.3 4.8]		
LTRO<1y	0.2	3.3	0.2	2.6	-0.0	-0.7
	[0.2 0.2]	[3.1 3.4]	[0.2 0.2]	[2.6 2.8]		
LTRO1y	1.6	24.1	1.3	18.6	-0.3	-5.6
	[1.6 1.8]	[23.8 26.9]	[1.3 1.4]	[18.0 20.5]		
Total	2.1	32.7	2.4	34.0	**0.3**	**1.3**

SMP2	05/08/2011		12/08/2011			
	EL	ES(95%)	EL	ES(95%)	ΔEL	ΔES(95%)
SMP	7.3	43.6	8.2	52.6	**0.9**	**9.0**
	[7.3 7.4]	[43.6 44.2]	[8.2 8.4]	[52.6 53.3]		
MRO	0.8	13.4	0.7	11.0	-0.1	-2.4
	[0.8 0.9]	[12.9 14.2]	[0.7 0.7]	[10.8 11.8]		
LTRO<1y	1.9	29.0	1.9	27.9	-0.0	-1.1
	[1.9 2.0]	[28.1 30.3]	[1.8 2.0]	[27.1 29.6]		
Total	10.0	86.0	10.8	91.5	**0.8**	**5.5**

VLTRO1	16/12/2011		30/12/2011			
	EL	ES(95%)	EL	ES(95%)	ΔEL	ΔES(95%)
SMP	20.4	112.2	18.1	109.3	-2.3	-2.9
MRO	1.7	25.3	1.0	13.2	-0.7	-12.1
LTRO<1y	2.4	30.7	1.5	18.6	-0.9	-12.1
LTRO1y	0.3	5.0	0.1	1.4	-0.2	-3.6
VLTRO3y	-	-	2.5	39.6	**2.5**	**39.6**
Total	24.8	173.3	23.2	182.2	**-1.6**	**8.9**

VLTRO2	24/02/2012		09/03/2012			
	EL	ES(95%)	EL	ES(95%)	ΔEL	ΔES(95%)
SMP	18.5	104.7	19.6	102.0	1.2	-2.7
MRO	1.1	14.0	0.1	1.4	-1.0	-12.5
LTRO<1y	0.8	11.2	0.4	4.8	-0.5	-6.4
LTRO1y	0.1	1.3	0.1	1.3	-0.0	-0.0
VLTRO3y	2.3	37.4	4.4	72.1	**2.1**	**34.7**
Total	22.8	168.6	24.6	181.6	**1.8**	**13.0**

Table 1: Portfolio credit risks around key policy announcements; ctd.

Portfolio credit risks for different monetary policy operations around six policy announcements: the SMP announcement on 10 May 2010, the cross-sectional extension of the SMP on 08 August 2011, the allocation of the first VLTRO on 20 December 2011 and of the second VLTRO on 20 February 2012, OMT announcement on 02 August 2012, and the announcement of the OMT's technical details on 06 September 2012.

OMT1	27/07/2012		03/08/2012			
	EL	ES(95%)	EL	ES(95%)	ΔEL	ΔES(95%)
SMP	18.6	103.2	18.0	101.6	**-0.6**	**-1.5**
MRO	1.0	12.7	1.0	12.2	0.0	-0.5
LTRO<1y	0.5	7.2	0.4	6.1	-0.1	-1.1
LTRO1y	0.0	0.7	0.0	0.7	-0.0	-0.0
VLTRO3y	6.0	87.3	5.5	81.4	**-0.4**	**-6.0**
Total	26.1	211.2	25.0	202.1	**-1.1**	**-9.1**

OMT2	31/08/2012		07/09/2012			
	EL	ES(95%)	EL	ES(95%)	ΔEL	ΔES(95%)
SMP	17.2	99.0	14.8	81.9	**-2.3**	**-17.1**
MRO	1.0	11.7	0.8	8.8	-0.2	-2.8
LTRO<1y	0.4	5.6	0.3	4.9	-0.1	-0.8
LTRO1y	0.0	0.7	0.0	0.6	-0.0	-0.1
VLTRO3y	5.2	77.1	4.2	61.4	**-1.0**	**-15.7**
Total	23.7	194.1	20.1	157.6	**-3.6**	**-36.5**

5 Concluding discussion

We introduced a novel tractable non-Gaussian framework to quantify central bank portfolio credit risks over time. An application of the framework to a subset of ECB monetary policy operations during the financial and sovereign debt crisis between 2009 and 2015 suggests that LOLR- and IOLR-implied portfolio credit risks depend on the central bank's own actions and tend to be negatively related. Taking risk in one way can spill over and de-risk other part of the central bank's balance sheet.

Our findings can have important implications for the design of central banks' post-crisis operational frameworks. In particular, there is a debate to what extent there remains a need for a lender-/investor-of-last-resort given substantial crisis-related bank regulations that aim to lower the frequency of such crises. Liquidity and credit regulations can never fully eliminate illiquidity risks from the bank system. Our results suggest that lender/investor-of-last-resort policies can have a substantial impact on point-in-time risks and risk correlations. The risk-related consequences for the lender-of-last-resort's balance sheet can be assessed in real time and appeared manageable. Lender-of-last-resort policies can thus effectively complement liquidity-related regulations.

References

Abadir, K. and J. Magnus (2005). *Matrix Algebra*. Cambridge University Press.

Acharya, V. and S. Steffen (2015). The "greatest" carry trade ever? Understanding eurozone bank risks. *Journal of Financial Economics 115*, 215–236.

Allen, F. and D. Gale (2000). Financial contagion. *Journal of Political Economy 108 (1)*, 133.

Andres, P. (2014). Computation of maximum likelihood estimates for score driven models for positive valued observations. *Computational Statistics and Data Analysis*, forthcoming.

Azizpour, S., K. Giesecke, and G. Schwenkler (2017). Exploring the sources of default clustering. *Journal of Financial Economics*.

Bagehot, W. (1873). *Lombard Street: A Description of the Money Market*. London: Henry S. King & Co.

Bindseil, U. (2014). *Monetary Policy Operations and the Financial System*. Oxford University Press.

Bindseil, U. and L. Laeven (2017). Confusion about the lender-of-last-resort. *Vox.Eu article on 13 January 2017*.

Blasques, F., S. J. Koopman, and A. Lucas (2015). Information theoretic optimality of observation driven time series models for continuous responses. *Biometrika 102*(2), 325–343.

Carpenter, S. B., J. E. Ihrig, E. C. Klee, D. W. Quinn, and A. H. Boote (2013). The Federal Reserves balance sheet and earnings: A primer and projections. *Finance and Economics Discussion Series no. 2013-1, Federal Reserve Board, January.*, 1–30.

Christensen, J. H. E., J. A. Lopez, and G. D. Rudebusch (2015). A probability-based stress test of Federal Reserve assets and income. *Journal of Monetary Economics 73*, 26–43.

Constâncio, V. (2011). Contagion and the European debt crisis. Keynote lecture at Bocconi University Milan on 10 October 2011.

Creal, D., S. Koopman, and A. Lucas (2011). A dynamic multivariate heavy-tailed model for time-varying volatilities and correlations. *Journal of Business & Economic Statistics 29*(4), 552–563.

Creal, D., S. Koopman, and A. Lucas (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics 28*(5), 777–795.

Creal, D., B. Schwaab, S. J. Koopman, and A. Lucas (2014). An observation driven mixed measurement dynamic factor model with application to credit risk. *The Review of Economics and Statistics 96(5)*, 898–915.

CreditMetrics (2007). CreditMetrics (TM) - Technical Document, RiskMetrics Group. www.riskmetrics.com/pdf/dnldtechdoc/CMTD1.pdf.

Crosbie, P. and J. R. Bohn (2003). Modeling default risk. *Moody's KMV Company, available online.*.

Cruces, J. J. and C. Trebesch (2013). Sovereign defaults: The price of haircuts. *American Economic Journal: Macroeconomics 5*(3), 85–117.

De Pooter, M., R. Martin, and S. Pruitt (2018). The liquidity effects of official bond market intervention in Europe.

Demarta, S. and A. J. McNeil (2005). The t copula and related copulas. *International Statistical Review 73(1)*, 111–129.

Duffie, D., A. Eckner, G. Horel, and L. Saita (2009). Frailty Correlated Default. *Journal of Finance 64(5)*, 2089–2123.

Duffie, D., L. Saita, and K. Wang (2007). Multi-Period Corporate Default Prediction with Stochastic Covariates. *Journal of Financial Economics 83(3)*, 635–665.

ECB (2012). Technical features of Outright Monetary Transactions. ECB press release, 6 September 2012.

ECB (2013). European Central Bank Annual Report 2012.

ECB (2014). The determinants of euro area sovereign bond yield spreads during the crisis. ECB Monthly Bulletin article, May 2014.

Engle, R. (2002). Dynamic conditional correlation. *Journal of Business and Economic Statistics 20*(3), 339–350.

Engle, R. and B. Kelly (2012). Dynamic equicorrelation. *Journal of Business & Economic Statistics 30*(2), 212–228.

Eser, F. and B. Schwaab. Evaluating the impact of unconventional monetary policy measures: Empirical evidence from the ECB'sSecurities Markets Programme. *Journal of Financial Economics 119*.

Freixas, X., C. Giannini, G. Hoggarth, and F. Soussa (2004). Lender of last resort: What have we learned since Bagehot. *Journal of Financial Services Research 18(1)*, 63–84.

Ghysels, E., J. Idier, S. Manganelli, and O. Vergote (2017). A high frequency assessment of the ECB Securities Markets Programme. *Journal of European Economic Association 15*, 218–243.

Giesecke, K., K. Spiliopoulos, R. B. Sowers, and J. A. Sirignano (2014). Large portfolio loss from default. *Mathematical Finance, forthcoming*.

Glosten, L. R., R. Jaganathan, and D. E. Runkle (1993). Relationship between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance 48*, 1779–1801.

González-Páramo, J.-M. (2011). The ECB's monetary policy during the crisis. Closing speech at the Tenth Economic Policy Conference, Málaga, 21 October 2011.

Gordy, M. (2000). A comparative anatomy of credit risk models. *Journal of Banking and Finance 24*, 119–149.

Gordy, M. B. (2003). A risk-factor model foundation for ratings-based bank capital rules . *Journal of Financial Intermediation 12(3)*, 199–232.

Greenlaw, D., J. D. Hamilton, P. Hooper, and F. Mishkin (2013). Crunch time: Fiscal crises and the role of monetary policy. *Finance and Economics Discussion Series no. 2013-1, Federal Reserve Board, January.*, 1–30.

Harvey, A. and A. Luati (2014). Filtering with heavy tails. *Journal of the American Statistical Association, forthcoming*.

Harvey, A. C. (2013). *Dynamic Models for Volatility and Heavy Tails*. Cambridge University Press.

Koopman, S. J., A. Lucas, and B. Schwaab (2011). Modeling frailty correlated defaults using many macroeconomic covariates. *Journal of Econometrics 162 (2)*, 312–325.

Koopman, S. J., A. Lucas, and B. Schwaab (2012). Dynamic factor models with macro, frailty, and industry effects for u.s. default counts: the credit crisis of 2008. *Journal of Business and Economic Statistics 30(4)*, 521–532.

Lando, D. (2003). *Credit Risk Modelling - Theory and Applications*. Princeton University Press.

Lucas, A., P. Klaassen, P. Spreij, and S. Straetmans (2001). An analytic approach to credit risk of large corporate bond and loan portfolios. *Journal of Banking & Finance 25*(9), 1635–1664.

Lucas, A., P. Klaassen, P. Spreij, and S. Straetmans (2003). Tail behaviour of credit loss distributions for general latent factor models. *Applied Mathematical Finance 10*(4), 337–357.

Lucas, A., B. Schwaab, and X. Zhang (2014). Conditional euro area sovereign default risk. *Journal of Business and Economic Statistics 32*(2), 271.

Lucas, A., B. Schwaab, and X. Zhang (2017). Modeling financial sector joint tail risk in the euro area. *Journal of Applied Econometrics 32*(1), 171.

Merton, R. (1974). On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance 29(2)*, 449–470.

Oh, D. H. and A. Patton (2014). Time-varying systemic risk: Evidence from a dynamic copula model of cds spreads. *Working paper*.

Rochet, J.-C. and X. Vives (2004). Coordination failures and the lender of last resort: Was Bagehot right after all? *Journal of the European Economic Association 2(6)*, 1116–1147.

Rodriguez, M. J. and E. Ruiz (2012). Revisiting several popular GARCH models with leverage effect: Differences and similarities. *Journal of Financial Econometrics 10(4)*, 637–668.

Vasicek, O. (1987). Probability of loss on loan portfolio. *Working paper, Moody's Analytics Corporation*.

Wessel, D. (2013). How jawboning works. *Wall Street Journal, January 9*.

Zhang, X., D. Creal, S. Koopman, and A. Lucas (2011). Modeling dynamic volatilities and correlations under skewness and fat tails. *Tinbergen Institute Discussion Paper*.

Web Appendix to

"Risk endogeneity at the lender/investor-of-last-resort"

Diego Caballero, André Lucas,

Bernd Schwaab, and Xin Zhang

Web Appendix A: CDS-implied sovereign EDFs

[To be added]

Web Appendix B: Univariate modeling

This section summarizes our univariate marginal modeling strategy. We estimate the parameters of univariate dynamic Student's t models using log-changes in country j-specific banking sector EDFs as inputs. For sovereigns, we use log-changes in CDS-implied-EDFs. Parameters are estimated based on maximum likelihood. The univariate models allow us to transform the observations into their probability integral transforms $\hat{u}_{jt} \in [0,1]$ based on parameter estimates for σ_{jt} and ν_j. The univariate t density is given by

$$p(y_t;\sigma_t^2,\nu) = \frac{\Gamma((\nu+1)/2)}{\Gamma(\nu/2)\sqrt{(\nu-2)\pi\sigma_t^2}} \cdot \left(1 + \frac{y_t^2}{(\nu-2)\sigma_t^2}\right)^{-\frac{\nu+1}{2}},$$

where σ_t is chosen as $\sigma_t = \sigma(f_t) = \exp(f_t)$.

We would like to allow for a leverage (asymmetry) effect also in the univariate score dynamics for volatility. We do so by defining a leverage term that takes nonzero values whenever $y_t > \tilde{\mu}_t$ with $\tilde{\mu}_t = 0$. The score-driven transition equation for this specification are given by

$$\begin{aligned}
f_{t+1} &= \tilde{\omega} + \sum_{i=0}^{p-1} A_i s_{t-i} + \sum_{j=0}^{q-1} B_j f_{t-j} + C(s_t - s_t^{\tilde{\mu}})1\{y_t > \tilde{\mu}_t\}, \\
s_t &= \mathcal{S}_t \nabla_t, \quad \nabla_t = \partial \ln p(y_t|\mathcal{F}_{t-1}; f_t, \theta)/\partial f_t, \\
s_t^{\tilde{\mu}} &= \mathcal{S}_t \nabla_t^{\tilde{\mu}}, \quad \nabla_t^{\tilde{\mu}} = \partial \ln p(\tilde{\mu}_t|\mathcal{F}_{t-1}; f_t, \theta)/\partial f_t, \\
\mathcal{S}_t &= \frac{\nu+3}{2\nu}, \\
\nabla_t &= \frac{(\nu+1)y_t^2}{(\nu-2)\sigma_t^2 + y_t^2} - 1.
\end{aligned}$$

The above expressions allow for a an asymmetric volatility response and skewness in the unconditional data density, see e.g. Rodriguez and Ruiz (2012) for a survey. Univariate filtered volatilities σ_{jt} can be obtained in this way as well.

Web Appendix C: Score and scaling function for the multivariate t-copula

We estimate the parameters of the multivariate dynamic t-copula model using the pseudo-observations $\hat{u}_{it} \in [0, 1]$ from on the marginal univariate models as inputs. These pseudo-observations are later transformed to t distributed random variables $y_t = F^{-1}(\hat{u}_t; \nu)$. The D-dimension multivariate t density is given by

$$p(y_t; \Sigma_t, \nu) = \frac{\Gamma((\nu + D)/2)}{\Gamma(\nu/2)[(\nu - 2)\pi]^{D/2}|\Sigma_t|^{1/2}} \cdot \left[1 + \frac{y_t' \Sigma_t^{-1} y_t}{(\nu - 2)}\right]^{-\frac{\nu+D}{2}},$$

where Σ_t is the covariance matrix of y_t and $\nu > 2$ is the degree of freedom parameter for the multivariate density. One could rewrite the model in terms of the scaling matrix $\tilde{\Sigma}_t$ as well. Doing so relaxes the parameter restriction on ν to $\nu > 0$. Note that the variance is the identity matrix in our copula setting because the univariate models effectively de-volatized the log-changes in bank and sovereign EDFs. As a result, the covariance matrix Σ_t is equivalent to the correlation matrix R_t.

The score-driven dynamics are incorporated in the covariance matrix in a similar fashion as in Creal et al. (2011), where $\Sigma_t = \Sigma(f_t)$. We also include an asymmetry/leverage term in the copula, defined for the adverse case that risks go up, i.e. that $y_t > \tilde{\mu}_t$. The dynamic system is given by

$$\begin{aligned}
f_{t+1} &= \tilde{\omega} + \sum_{i=0}^{p-1} A_i s_{t-i} + \sum_{j=0}^{q-1} B_j f_{t-j} + C \cdot vech(1\{y_t > \tilde{\mu}_t\} \cdot 1\{y_t > \tilde{\mu}_t\}')(s_t - s_t^{\tilde{\mu}}), \\
s_t &= \mathcal{S}_t \nabla_t, \quad \nabla_t = \partial \ln p(y_t | \mathcal{F}_{t-1}; f_t, \theta)/\partial f_t, \\
s_t^{\tilde{\mu}} &= \mathcal{S}_t \nabla_t^{\tilde{\mu}}, \quad \nabla_t^{\tilde{\mu}} = \partial \ln p(\tilde{\mu}_t | \mathcal{F}_{t-1}; f_t, \theta)/\partial f_t, \\
\mathcal{S}_t &= \left(\frac{1}{4} \Psi_t' \mathcal{D}_k' (\mathcal{J}_t' \otimes \mathcal{J}_t')[gG - vec(\mathrm{I})vec(\mathrm{I})'](\mathcal{J}_t \otimes \mathcal{J}_t)\mathcal{D}_k \Psi_t\right)^{-1}, \\
\nabla_t &= \frac{1}{2} \Psi_t' \mathcal{D}_k' (\Sigma_t^{-1} \otimes \Sigma_t^{-1}) \left[\frac{(\nu + k)}{\nu - 2 + y_t' \Sigma_t^{-1} y_t}(y_t \otimes y_t) - vec(\Sigma_t)\right].
\end{aligned}$$

where Ψ_t is the derivative $\partial vech(\Sigma_t)/\partial f_t'$, $\mathcal{D}_k$ is the duplication matrix and $\mathcal{J}_t$ is defined as the square root matrix such that $\Sigma_t^{-1} = \mathcal{J}_t' \mathcal{J}_t$. The scalar g is $(\nu + k)/(\nu + k + 2)$. The $k^2 \times k^2$ matrix

G is a particular matrix whose element $G[\cdot,\cdot]$ is given by

$$G[(i-1)\cdot k+\ell,(j-1)\cdot k+m]=\delta_{ij}\delta_{\ell m}+\delta_{i\ell}\delta_{jm}+\delta_{im}\delta_{j\ell},$$

for $i,j,\ell,m=1,\cdots,k$. The Kronecker delta δ_{ij} is unity if $i=j$ or 0 otherwise.

There are a few alternative ways to map f_t into the covariance matrix Σ_t, and the specific form of Ψ_t varies accordingly. We adopt a correlation structure based on hyperspherical coordinates $R_t=X_t'X_t$ where $X_t(\phi_t)$ is an upper-triangular matrix,

$$X_t=\begin{pmatrix} 1 & c_{12t} & c_{13t} & \cdots & c_{1kt} \\ 0 & s_{12t} & c_{23t}s_{13t} & \cdots & c_{2kt}s_{1kt} \\ 0 & 0 & s_{23t}s_{13t} & \cdots & c_{3kt}s_{2kt}s_{1kt} \\ 0 & 0 & 0 & \cdots & c_{4kt}s_{3kt}s_{2kt}s_{1kt} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & c_{k-1,kt}\prod_{\ell=1}^{k-2}s_{\ell kt} \\ 0 & 0 & 0 & \cdots & \prod_{\ell=1}^{k-1}s_{\ell kt} \end{pmatrix},$$

where $c_{ijt}=\cos(\phi_{ijt})$ and $s_{ijt}=\sin(\phi_{ijt})$. This setup ensures that Σ_t is a proper covariance matrix regardless of f_t. For example, in the 2-dimensional case, we have the root and correlation matrices given by

$$X_t=\begin{pmatrix} 1 & \cos(\phi_{12,t}) \\ 0 & \sin(\phi_{12,t}) \end{pmatrix}, \qquad R_t=X_t'X_t=\begin{pmatrix} 1 & \cos(\phi_{12,t}) \\ \cos(\phi_{12,t}) & 1 \end{pmatrix}, \tag{C1}$$

with the correlation given by $\cos(\phi_{12,t})$.

Under such a parameterization, we can complete the system using the result that

$$\Psi_t=\mathcal{B}_k[(\mathrm{I}\otimes X_t')+(X_t'\otimes\mathrm{I})\mathcal{C}_k]Z_t\mathcal{S}_\phi,$$

where $\mathcal{B}_k$ is the elimination matrix, $\mathcal{C}_k$ is the commutation matrix, $\mathcal{S}_\phi$ is the selection matrix $\phi_t=\mathcal{S}_\phi f_t$, and $Z_t=\partial vec(X_t)/\partial\phi_t'$; see Abadir and Magnus (2005).

Web Appendix D: Risks in % of size around policy announcements

[To be added.]

Web Appendix E: Expected losses with and without the leverage term

Figure 6 provides evidence of the economic significance of the leverage term. The leverage term matters moderately. Instead, portfolio risk is most sensitive to the changes in the marginal probabilities of default p_{it}, as inferred from the EDFs in levels.

Figure 6: Expected losses from collateralized lending
The top panel plots the expected losses from liquidity providing operations with the leverage term restricted to zero and re-estimated remaining parameters. The bottom panel plots plots the expected losses with non-zero leverage terms. Vertical axes are in billion euro. Data is weekly between 2009 and 2015.