JOURNAL OF APPLIED STATISTICS
2024, VOL. 51, NO. 3, 515–533
https://doi.org/10.1080/02664763.2022.2138837

Taylor & Francis
Taylor & Francis Group

New nonparametric measures for instantaneous and granger-causality tail co-dependence

Cees Diks[a,b] and Marcin Wolski[a,c]

[a]Center for Nonlinear Dynamics in Economics and Finance (CeNDEF), University of Amsterdam, Amsterdam, The Netherlands; [b]Tinbergen Institute, Amsterdam, The Netherlands; [c]European Investment Bank, Luxembourg, Luxembourg

ABSTRACT

We propose a new methodology to asses risk spillovers in a time-series framework. Firstly, we introduce an explicit nonparametric measure of cross-sectional conditional tail co-movement, which is intuitively comparable to the Conditional Value-at-Risk (CoVaR). We show that nonlinear CoVaR (NCoVaR) is able to capture even highly nonlinear dependence structures. Secondly, for the purpose of potential contagion analysis, we adapt the measure to be informative about the causality direction between the variables in the Granger causality sense. By showing that the natural estimators of the two metrics are U-statistics, we construct formal nonparametric tests for independence and Granger non-causality. Numerical simulations confirm that in common situations the nonparametric tests have better size and power properties than their parametric counterparts. The methodology is illustrated empirically by assessing risk transmissions between sovereigns and banking sectors in the euro area, which observed highly irregular co-movements between asset prices after the global financial crisis. The new measures seem to be less susceptible to these irregularities than their parametric analogues, providing a clearer overview of the underlying sovereign-bank risk feedback loops.

ARTICLE HISTORY
Received 15 January 2022
Accepted 16 October 2022

KEYWORDS
Financial risk; nonparametric tests; general Granger causality

AMS SUBJECT CLASSIFICATIONS
C1; G1; G20; G32

1. Introduction

Conditional Value-at-Risk, hereafter CoVaR, has become an industry benchmark in analyzing tail co-dependence between financial series [2]. As we show in this paper, standard linear and/or (G)ARCH type parametric approaches, which are the most commonly used CoVaR estimation frameworks, may however lead to inaccurate co-risk assessment as they overlook deviations from linearity and/or normality. They are also unable to determine the

CONTACT Marcin Wolski M.Wolski@eib.org Center for Nonlinear Dynamics in Economics and Finance (CeNDEF), University of Amsterdam, Roetersstraat 11, Amsterdam 1018 WB, The Netherlands; Tinbergen Institute, Gustav Mahlerplein 117, Amsterdam 1082 MS, The Netherlands; European Investment Bank, 98-100 Boulevard Konrad Adenauer, Luxembourg 2950, Luxembourg

Supplemental materials for this article, including the mathematical appendix, data and codes, can be accessed here https://doi.org/10.1080/02664763.2022.2138837

Table 1. Taxonomy of CoVaR-type measures.

	Type of dependence	
Method	Cross-sectional/instantaneous	Causal/dynamic
Parametric	CoVaR	*CoVaR Granger causality*
Nonparametric	*NCoVaR*	*NCoVaR Granger causality*

Note: Entries in italics highlight the contributions of this paper.

direction of the relation between variables, making it more challenging to apply them in contagion analysis.

Motivated by this, we develop a new measure for tail co-movement, which is more flexible than a parametric measure, but still intuitively comparable to CoVaR. Firstly, we adapt the original ΔCoVaR metric to the nonlinear setting (hence the new metric is called ΔNCoVaR). Our preferred specification focuses on the regions around respective quantiles, and as such modestly deviates from the standard CoVaR measure, however we also offer a specification with the regions below quantiles, which is analogous to the original metric. Secondly, we introduce Granger causal effects to the standard CoVaR. While this is rather a model adaptation than a novel risk measure, to the best of knowledge, this is the first attempt to test it in a controlled numerical environment and to compare its performance against other measures empirically. The resulting family of CoVaR measures, highlighting the ones developed in this paper, is summarized in Table 1.

By using U-statistic representations we derive asymptotic normality of the nonparametric estimators and demonstrate numerically that NCoVaR and NCoVaR Granger causality (NCoVaR-GC) are more suitable in the presence of dynamic spillover effects, compared to their parametric counterparts. In the same spirit, Jeong *et al.* [19] proposed a nonparametric test for causality in quantile, extending the work of Hong *et al.* [17] to nonparametric Value-at-Risk estimation. These methodologies are complementary to ours in the sense that whereas these authors focus on the effect of one variable on the conditional quantiles of another, we study conditional tail probability effects.

Our framework is also closely related to the literature on conditional independence testing – a research area in which a range of new nonparametric approaches have been proposed recently, see e.g. [4,22,29,30,32]. While it is only natural that some of the arguments may overlap between these papers and ours, we hope that our approach will spark interest among the readers already familiar with CoVaR.

We apply the proposed methodology to assess the bank-sovereign risk transmission between selected Euro Area (EA) countries. We confirm the differences between the core- and vulnerable-EA dynamics, identified earlier by Paries *et al.* [25], Ohnsorge *et al.* [24] and EIB [11]. We furthermore argue that NCoVaR is a more conservative methodology and to a higher extent it captures the information spanned by extra confounding variables.

We argue that NCoVaR performs well in nonlinear and/or non-Gaussian settings. In fact, the existence of nonlinearities is widely recognized in the financial literature. For instance, von Borstel *et al.* [31] find that the sovereign debt crisis changed the composition of the pass-through, adjusting for indirect effects from lower sovereign risk premia in the EA. [20] find that banks with high shares of relationship lending appear to be characterized by nonlinear pass-through effects. This is somehow in line with a more general finding of Huang *et al.* [18] and He and Krishnamurthy [16], who suggest that banking risk is a nonlinear function of asset exposure.

A major motivation for considering Granger causality lies in its possible applications to networks and contagion analysis (see e.g. [6]). Looking at any pair of institutions, the possible mutual risk transmission effects do not have to be symmetric (as they are implicitly assumed to be in a linear Gaussian setting). For instance, a lender has a different kind of risk exposure to a creditor than vice versa. Granger causality captures that phenomenon explicitly, allowing for a more detailed analysis of network spillover effects, cascades and shock propagation. A parametric Granger causality measure has already been successfully applied as a network mapping tool in financial analysis [12]. The general type of Granger causality employed in this study, i.e. a distribution-wide version of the concept originally proposed by Granger [13], goes beyond the conditional mean effects and spreads over the entire conditional distributions. Such a general notion of causality appears to be particularly relevant to identify feedback loops and propagation mechanisms in complex environments.

The motivation behind the empirical part stems from the strong adverse effects from the bank-sovereign feedback loops, i.e. the interdependence of the banking sectors and corresponding sovereigns, on the real economy and taxpayers. The majority of econometric approaches in these fields focus on co-risk measures, where the risk of one sector is assessed in relation to the risk of the other one. The intuition behind these models lies in negative externalities. As argued by Adrian and Brunnermeier [2], such externalities are a consequence of asset exposures, excessive risk taking and leverage. Given, for instance, that the banking sector is facing a liquidity shock, it liquidates its assets, including sovereign securities, at fire-sale prices as given, affecting the borrowing constraints of the sovereign. On the other hand, sovereign characteristics are often perceived as a country-wide benchmark in credit risk assessment. It is rarely the case that a financial entity 'pierces the sovereign ceiling' in a credit rating context. Furthermore, sovereigns are often the implicit guarantors of the financial system [1].

This paper is organized as follows. In Section 2, we explain the methodology of NCoVaR and NCoVaR-GC. We numerically evaluate the asymptotic properties of the test statistic in Section 3. In Section 4, we test our approach on the EA sample. Section 5 concludes.

2. Methodology

In this section we introduce the main mechanics of the new risk transmission measures. While our methodology can be possibly applied to assess risk more broadly, we benchmark our framework against the literature on financial risk, in which investors and policy makers typically aim to understand probability of financial loss associated with a given scenario. This feeds into their decision making process, hoping to improve the expected outcome or better prepare for a possible disaster. Commonly used financial risk management techniques include measures of dispersion, Sharpe ratio, and so-called beta, which describes sensitivity of individual stocks to shocks in the market benchmark. More sophisticated methods include, for instance, Expected Shortfall or Value-at-Risk metrics, as we describe below. In this spirit, conditional measures, like CoVaR, allow to study how specific risk factors associated with one variable affect other variables, being particularly useful for spillover or contagion analysis. This is also an area which our methodology aims to contribute to.

For convenience, we begin by highlighting the main features of the standard CoVaR methodology. We then introduce an analogue of CoVaR which can capture even nonlinear dependence structures. Finally, we complement both CoVaR methods with their Granger causality versions determined by conditional rather than unconditional dependence.[1]

2.1. CoVaR

To start with the basics, the unconditional Value-at-Risk (VaR_γ^Y) of a continuous random variable Y representing losses of an institution, j, say, is defined as the γ^{th} quantile of Y for a given quantile level $\gamma \in (0, 1)$, that is

$$\mathrm{P}\left(Y \leq \mathrm{VaR}_\gamma^Y\right) = \gamma. \tag{1}$$

To make the presentation transparent and consistent with our further argumentation, let us denote the losses of another institution, i, say, by X. The conditional Value-at-Risk, or CoVaR, proposed by Adrian and Brunnermeier [2], measures the effect of an event $\mathbb{C}(X)$ occurring in institution i on VaR_γ of institution j. To put it formally, one can define implicitly CoVaR as

$$\mathrm{P}\left(Y \leq \mathrm{CoVaR}_\gamma^{Y|\mathbb{C}(X)} \middle| \mathbb{C}(X)\right) = \gamma. \tag{2}$$

To capture risk transmission effects, also referred to as tail dependence, Adrian and Brunnermeier [2] introduced the ΔCoVaR metric, which measures the change in CoVaR of an institution j when the conditioning event of institution i changes. ΔCoVaR measures the effect of a shift in X from the median to the tail quantile (from a safe to a risky state) of institution i on the performance of institution j. Formally, one can define

$$\Delta\mathrm{CoVaR}_\gamma^{Y|X} = \mathrm{CoVaR}_\gamma^{Y|X=\mathrm{VaR}_\gamma^X} - \mathrm{CoVaR}_\gamma^{Y|X=\mathrm{VaR}_{0.5}^X}, \tag{3}$$

measuring the shift in the conditional Y-quantile, in response of a change in X from the median to the γ-th quantile.

2.2. Nonlinear CoVaR (NCoVaR)

To introduce the concept, we reformulate CoVaR from conditional quantiles into conditional tail event probabilities. In other words, we change the perspective from the domain of Y values in Equation (2), to the domain of probabilities. While one can think of the NCoVaR simply as a conditional probability, our main focus is on the nonparametric adaptation of ΔCoVaR metric, called ΔNCoVaR.[2] The approach differs from [28,33] by focusing (like CoVaR) on the shift in conditional probability of Y given X as X changes, rather than on the conditional probability itself.

To keep the exposition general, first consider the events $Y \in A$, and $X \in C$ or $X \in D$, where the sets A, C and D are specific regions of interest. We can quantify the effect of a change in X from region C to D on the tail event probability of Y by defining

$$\Delta\mathrm{NCoVaR} = \mathrm{P}(Y \in A|X \in C) - \mathrm{P}(Y \in A|X \in D).$$

In our applications we let A correspond to a set of extreme events for Y, such as Y being near or above a given unconditional tail quantile $y_\gamma \equiv \mathrm{VaR}_\gamma^Y$, and C and D sets denoting events where X is either near a given unconditional tail quantile (x_γ) of X or near its unconditional median ($x_{0.5}$), respectively. In the simulations and empirical applications described below we focus on the choice

$$A = [y_\gamma - \mu, y_\gamma + \mu], \quad C = [x_\gamma - \mu, x_\gamma + \mu], \quad D = [x_{0.5} - \mu, x_{0.5} + \mu], \tag{4}$$

for some positive parameter μ. Another obvious choice might be $A = [y_{\gamma'}, \infty)$, which has a similar interpretation to the standard CoVaR. In practice we found this to give similar results if γ' is chosen such that $y_{\gamma'}$ approximately equals $y_\gamma - \mu$, if γ is in the range 0.95 to 0.99 and μ is not too small. For ease of exposition, we specify the event A as being near a given tail quantile throughout this paper.

From the definition of ΔNCoVaR, we obtain

$$\Delta\text{NCoVaR} = \frac{\text{P}\,(Y \in A, X \in C)}{\text{P}\,(X \in C)} - \frac{\text{P}\,(Y \in A, X \in D)}{\text{P}\,(X \in D)}.$$

The null hypothesis H_0 of no (instantaneous) NCoVaR relation of the variable X on the tail probability of Y can be defined as $\Delta\text{NCoVaR} = 0$ for all events A, C and D for which $P(X \in C) > 0$ and $P(X \in D) > 0$. Upon multiplication by $\text{P}(X \in D)\text{P}(X \in C)$, H_0 can be seen to imply

$$q \equiv \text{P}\,(Y \in A, X \in C)\,\text{P}\,(X \in D) - \text{P}\,(Y \in A, X \in D)\,\text{P}\,(X \in C) = 0,$$

which then is, in fact, also defined in cases where $P(X \in C)$ and/or $P(X \in D)$ happen to be zero. Note that equivalently, we may write H_0 as

$$q = \text{E}\,[I_{A\times C}(Y_1, X_1)I_D(X_2) - I_{A\times D}(X_1, Y_1)I_C(X_2)] = 0,$$

for two vectors (X_1, Y_1) and (X_2, Y_2), drawn independently from the joint distribution of (X, Y).

Given a sample of the process $\{(X_t, Y_t)\}$, $t = 1, \ldots, n$, a plug-in frequency count-based estimate of q is

$$q_n = \frac{1}{n(n-1)} \sum\sum_{\ell \neq k} [I_{A\times C}(Y_k, X_k)I_D(X_\ell) - I_{A\times D}(Y_k, X_k)I_C(X_\ell)]\,.$$

2.2.1. Asymptotic theory

To develop asymptotics, define $W_t = (X_t, Y_t)$ and write the estimator as a weighted average of a symmetric kernel function (as a U-statistic)

$$q_n = \frac{1}{n(n-1)} \sum\sum_{k \neq \ell} \mathcal{K}(W_k, W_\ell),$$

where $\mathcal{K}(W_k, W_\ell)$ is the symmetric (w.r.t. swapping W_k and W_ℓ) kernel function

$$\mathcal{K}(W_k, W_\ell) = \frac{1}{2}\,[I_{A\times C}(Y_k, X_k)I_D(X_\ell) - I_{A\times D}(Y_k, X_k)I_C(X_\ell) + \ell \leftrightarrow k]\,,$$

where $\ell \leftrightarrow k$ stands for similar terms with ℓ and k swapped. This shows that q_n is in fact a U-statistic estimator of $q = E(\mathcal{K}(W_i, W_j))$, (where W_i and W_j are drawn independently from the stationary distribution of W, provided it exists). Although in a time series process the observed vectors W_t, $t = 1, \ldots, n$ are not independent, as long as the process $\{W_t\}$ is strictly stationary and satisfies some rather mild mixing conditions, the asymptotic theory of U-statistics still apply, provided that a HAC estimator of variance is used [7,8]. This leads to the following theorem.

Theorem 2.1: *Consider a sample $\{W_t\}_{t=1}^n$ from the bivariate random process $\{W_t\} \equiv \{(X_t, Y_t)\}$ with $t \in \mathbb{Z}$ that is strictly stationary and β-mixing with exponential decay rate. Then for the kernel function $\mathcal{K}(\cdot,\cdot)$ as defined above, for fixed A, C and D,*

$$\sqrt{n}\frac{q_n - q}{S_n} \xrightarrow{d} N(0,1),$$

where S_n^2 is any heteroskedasticity and autocovariance consistent (HAC) estimator of the asymptotic variance of $\sqrt{n}(q_n - q)$.

The proof of Theorem 2.1 is provided in Appendix A.

2.3. CoVaR Granger causality (CoVaR-GC)

In Granger causality testing, the goal is to find evidence against the null hypothesis of Granger non-causality. We define it in a general sense as follows.

Definition 2.1 (Granger non-causality (bivariate)): For a strictly stationary bivariate time series process $\{(X_t, Y_t)\}$, $t \in \mathbb{Z}$, $\{X_t\}$ does not Granger cause $\{Y_t\}$ if

$$Y_{t+1}|(\mathcal{F}_{X,t}, \mathcal{F}_{Y,t}) \sim Y_{t+1}|\mathcal{F}_{Y,t} \quad \forall_t,$$

where $\mathcal{F}_{X,t}$ and $\mathcal{F}_{Y,t}$ are information sets spanned by X_s, $s \le t$ and Y_s, $s \le t$, respectively and '$\sim$' denotes that the random variables on both sides are identically distributed.

Throughout we assume that the process $\{(X_t, Y_t), t \in \mathbb{Z}\}$ is strictly stationary and β-mixing with exponential decay rate. Since it is generally impossible to condition on the entire past history spanned by X_s and/or Y_s, $s \le t$, in practice we condition on the last k observations only ($t - k + 1 \le s \le t$) where k is a finite positive integer. For now, we focus on $k = 1$.

Following [9,10] one can represent the null hypothesis of Granger non-causality in terms of equality of conditional probabilities. For the ease of notation we introduce the lead variable $Z_t = Y_{t+1}$. In this notation the null hypothesis is a statement about the invariant distribution evaluated at conditional quantile levels of the 3-dimensional vector $W_t = (X_t, Y_t, Z_t)$. For clarity, since the null hypothesis concerns the invariant distribution of W_t, in formulating the null hypothesis we often drop the time index and refer simply to the distribution of the random variable $W = (X, Y, Z)$.

Under the null hypothesis of Granger non-causality X and Z are conditionally independent given $Y = y^*$, where y^* is a given unconditional quantile of Y. This, in fact, allows to estimate the relation by adding a time lag over Y-dimension to the basic CoVaR quantile regression. We call this model specification by CoVaR-GC. Following Definition 1, the null hypothesis H_0 implies that the conditional quantiles of Z, being the future Y-variable, are independent of X given Y, or, the conditional quantiles of $Z|X = x$, $Y = y$ and $Z|Y = y$ are the same for all (x, y) in the support of (X, Y).

Note that unlike CoVaR, CoVaR-GC is directional also in multivariate Gaussian settings, i.e. the effect of X on future values of Y need not be the same as that of Y on future X-values. As such it is a measure of Granger causality from one variable to another. This provides an adaptation of the linear CoVaR, which for Gaussian processes is correlation-driven, and hence inherently symmetrical under multivariate Gaussian conditions [23].

2.4. NCoVaR Granger causality (NCoVaR-GC)

Recall that for $(X, Y, Z) \sim (X_t, Y_t, Y_{t+1})$ under the null hypothesis H_0 vectors X and Z are conditionally independent given Y. As with NCoVaR, we take a nonparametric approach for estimation, starting by considering events where now A is a subset of the outcome space of Z, and C and D of that of X. For the simulations and applications presented below, we use the specific sets

$$A = [z_\gamma - \mu, z_\gamma + \mu], \quad C = [x_\gamma - \mu, x_\gamma + \mu], \quad D = [x_{0.5} - \mu, x_{0.5} + \mu], \tag{5}$$

however, in general one may also utilize the above-quantile specification of the events (for a discussion see Section 2.2).

For general sets A, C and D, the null hypothesis implies (see Appendix B for the derivation)

$$\frac{P(Z \in A, X \in C|Y = y)}{P(X \in C|Y = y)} = \frac{P(Z \in A, X \in D|Y = y)}{P(X \in D|Y = y)}, \quad \forall A, C, D, y.$$

The analogy with ΔNCoVaR above now suggests testing the null hypothesis

$$\begin{aligned} H_0': \quad & P(Z \in A, X \in C|Y = y_*)P(X \in D|Y = y_*) \\ & - P(z \in A, X \in D|Y = y_*)P(X \in C|Y = y_*) = 0, \end{aligned}$$

for a given past Y_t value y_* (e.g. some unconditional Y-quantile) and given A, C and D. For instance $P(X \in D|Y = y_*)$ is now estimated by counting the frequency of events $X \in D$, among the vectors close to y_*.

A plug-in estimator for $P(X \in D|Y = y_*)$ is the Nadaraya-Watson nonparametric regression function estimator

$$\hat{P}(X \in D|Y = y_*) = \frac{\frac{1}{n}\sum_{k=1}^{n} I_D(Z_k)K_h(y_* - Y_k)}{\frac{1}{n}\sum_{k=1}^{n} K_h(y_* - Y_k)}, \tag{6}$$

where we take $K_h(\cdot)$ to be a density estimation kernel. Although the theory holds more general, in the simulations and applications presented herein we focus on the Gaussian kernel

$$K_h(s) = \frac{1}{\sqrt{2\pi}h} \exp(-s^2/(2h^2)),$$

and its associated higher-order kernels (see e.g. [15]). Note that the denominator in Equation (6), which is preventing us from writing the estimator as a U-statistic, is just a kernel density estimate of $f_Y(y_*)$. We can get rid of the denominator and obtain simple U-statistics estimates[3] if we multiply the probabilities by $f_Y(y_*)$. Therefore, for a given unconditional quantile y_* of Y we define

$$q_* = f_Y^2(y_*)\left(P(Z \in A, X \in C|Y = y_*)P(X \in D|Y = y_*)\right.$$
$$\left. -P(Z \in A, X \in D|Y = y_*)P(X \in C|Y = y_*)\right).$$

By construction, $q_* = 0$ under H_0. Now the term

$$f_Y(y_*)P(X \in D|Y = y_*)$$

e.g. can be simply estimated as

$$\frac{1}{n}\sum_{k=1}^{n} I_D(X_k)K_h(y_* - Y_k).$$

The corresponding U-statistic kernel used for estimation of q_* is

$$\mathcal{K}(W_k, W_\ell; h) = \frac{1}{2}\left[I_A(Z_k)I_C(X_k)K_h(y_* - Y_k)I_D(X_\ell)K_h(y_* - Y_\ell)\right.$$
$$\left. - I_A(Z_k)I_D(X_k)K_h(y_* - Y_k)I_C(X_\ell)K_h(y_* - Y_\ell) + k \leftrightarrow \ell\right], \quad (7)$$

where '$k \leftrightarrow \ell$' represents the same terms with k and ℓ swapped and $K_h(w) = h^{-1}K(w/h)$ a scaled version of the kernel function $K(w)$, satisfying

$$\int |K(w)|dw < \infty, \quad \int K(w)dw = 1 \quad \text{and}|wK(w)| \to 0 \quad \text{as}|w| \to \infty. \quad (8)$$

The next theorem states that the U-statistics estimator of q_*, given by

$$q_{*,n} = \frac{1}{n(n-1)}\sum\sum_{\ell \neq k}\mathcal{K}(W_k, W_\ell; h),$$

is asymptotically normally distributed provided that the bandwidth $h = h_n$ tends to zero with n at an appropriate rate.

Theorem 2.2: *Consider a strictly stationary bivariate random process $\{(X_t, Y_t)\}$ with $t \in \mathbb{Z}$, and a kernel density bandwidth parameter tending to zero at the rate $h_n = cn^{-\beta}$ with $\beta \in (\frac{1}{2\alpha}, \frac{1}{2d_Y})$, where α is the order of the kernel and d_Y the dimension of the conditioning variable Y. Then, given the events A, C and D, and Y-quantile y_*,*

$$\sqrt{n}\frac{q_{*,n} - q_*}{S'_n} \xrightarrow{d} N(0,1),$$

where $S_n'^2$ is any consistent (HAC) estimator of the asymptotic variance of $\sqrt{n}(q_{*,n} - q_*)$

The proof of Theorem 2.2 is provided in Appendix C. The conditions on the bandwidth rate β for a second-order ($\alpha = 2$) density estimation kernel, and $d_Y = 1$ (first Markov-order bivariate process), imply $\beta \in (\frac{1}{4}, \frac{1}{2})$. The MSE-optimal rate in that case is $\beta = \frac{1}{3}$ (see Appendix C for details).

In addition to the MSE optimal bandwidth rate β, one could also find the asymptotically optimal value of c in the sequence $h_n = cn^{-\beta}$. Unfortunately, the optimal value of c will be not independent of the data generating process assumed. We offer a guide how to find the c parameter in Appendix D on an example of stylized VAR and ARCH processes.

2.5. Extensions to higher order processes and/or confounding variables

Most of the remarks above concern the bivariate case, with Markov order $k = 1$ lag and density estimation kernel order $\alpha = 2$. Adding more conditional variables, such as lagged Y-variables or possibly confounding variables puts additional restrictions on the feasible bandwidth rates (β-values). Mathematically, extra lagged Y-values or additional variables can be added to the conditioning variable Y, in $W = (X, Y, Z)$, when testing whether X and Z are conditionally independent given Y. The increase of the dimension of Y places additional conditions on the feasible rates β at which h_n can tend to zero without letting the bias or the variance of $q_{*,n}$ dominate asymptotically. Specifically, α and d_Y put the restrictions $\frac{1}{2\alpha} < \beta < \frac{1}{2d_Y}$, on β. Note that for the usual second order density estimation kernels, which have $\alpha = 2$, there are already no feasible rates β as soon as we increase the dimension d_Y of Y from 1 to 2, e.g. by adding a single extra lag of Y_t or the first lag of a possible extra, confounding, variable, V_t, say.

When addressing this issue, one has a choice between Data Sharpening (DS) on one hand, and the use of higher order kernels on the other. Both these methods reduce the order of the bias, opening up some room for the bandwidth to tend to zero slower, and hence reduce the variance. However, there is a huge practical difference between these methods, in that higher order kernels require only a single bandwidth to be considered, while in DS there is one bandwidth for the data sharpening step and another for the estimation step, and these need to be carefully adjusted to each other. For instance, as noted by Diks and Wolski [10] the DS bandwidth should go to zero at a slower rate than the density estimation bandwidth, to make sure that the gradient of the density is estimated consistently. Therefore, in this study we decided to use higher-order kernels rather than DS when Y is multivariate. In the case described above, with $d_Y = 2$, using a 4th-order ($\alpha = 4$) kernel provides a range of feasible β-values, $\beta \in (\frac{1}{8}, \frac{1}{4})$.

Formally, if we denote the vector of extra conditioning variables by V, one can represent the null hypothesis for the multivariate ΔNCoVaR as

$$H_0: \quad \mathrm{P}(Y \in A|X \in C, V = v_*) = \mathrm{P}(Y \in A|X \in D, V = v_*),$$

and for multivariate NCoVaR-GC as

$$\frac{P(Z \in A, X \in C|Y = y_*, V = v_*)}{P(X \in C|Y = y_*, V = v_*)} = \frac{P(Z \in A, X \in D|Y = y_*, V = v_*)}{P(X \in D|Y = y_*, V = v_*)}.$$

The latter equation illustrates explicitly that adding extra control variables to condition on is mathematically equivalent to increasing the dimension of the conditioning variable Y.

3. Size/power simulations

Our strategy is to simulate processes with stylized transmission channels and to verify statistical power of the methodologies proposed in Table 1. Since CoVaR and NCoVaR focus on instantaneous dependence, whereas CoVaR-GC and NCoVaR-GC considers Granger-type dependence, for the former we simulate processes with simultaneous and for the latter with lagged dependence. We then benchmark the results against standard and lagged parametric CoVaR specifications, respectively.

As argued by Adrian and Brunnermeier [2], ΔCoVaR can be estimated as $\Delta\text{CoVaR} = \hat{\beta}^i_\gamma \times (\text{VaR}^i_\gamma - \text{VaR}^i_{0.5})$, where $\hat{\beta}^i_\gamma$ comes from the quantile regression[4]

$$\hat{Y}^{j|X}_\gamma = \hat{\alpha}^i_\gamma + \hat{\beta}^i_\gamma X, \tag{9}$$

where $\hat{Y}^{j|X}_\gamma$ is the predicted value for γ-quantile of institution j conditional on a return realization X of institution i. In fact, variable $\hat{\beta}^i_\gamma$ captures the tail dependence between the institutions and is the core variable of interest for our further investigation. Under standard distributional assumptions, the estimated coefficient follows a Student's t-distribution with $n-2$ degrees of freedom. The statistical significance of $\hat{\beta}^i_\gamma$ is therefore a direct measure to assess the size and power of the parametric approach. Furthermore, to correct for possible heteroskedasticity we compute the t-statistics for $\hat{\beta}^i_\gamma$ with robust standard errors.

We carry out a one-sided t-test where under the null $H_0 : \beta^i_\gamma \leq 0$ and under the alternative $H_a : \beta^i_\gamma > 0$, to make the test size directly comparable with our further investigation.

To construct CoVaR estimates for lagged dependence structure, we adjust the lag composition of Equation (9) to match Definition 2.1. Besides, the testing framework follows the same principles as described above.

We consider two groups of processes, highlighting the dependence in the first and second conditional moments of the random variables. These processes constitute a natural testing environment, used before by Li and Racine [21] and Diks and Wolski [10].[5] One could also focus on a combination of both types of dependencies simultaneously. However, as we demonstrate later, the sensitivity of parametric and nonparametric methods is vastly different across these two groups of processes. We therefore consider them separately.

3.1. Dependence in the first conditional moment

To give an example of a data-generating process with causality in mean, consider the linear bivariate Vector Autoregressive model of order 1 (VAR(1)) given by

$$\begin{aligned} X_t &= aX_{t-1} + \sqrt{1-a^2}\varepsilon_{1,t}, \\ Y_t &= aX_{t-\tau} + \sqrt{1-a^2}\varepsilon_{2,t}, \end{aligned} \tag{10}$$

where $a \in (0,1)$ is a tuning parameter and $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ represent i.i.d. zero-mean innovations. We restrict the parameter a to be within the unit interval and keep the process stationary. The process is designed so that the causality runs from X to Y, which constitutes a foundation for power assessment. For the size assessment we use the same process, but switch the causality from Y to X, so that the null hypothesis of non-causality holds.

The subscripts denote the time dimension. Here we make an important distinction between instantaneous and lagged dependence spillovers or, by a slight abuse of terminology, 'instantaneous' and 'standard' Granger causality. By adjusting the dependence lag structure, we match the timing of the data generating process to (N)CoVaR (Granger causality) measures, providing an appropriate testing framework for each. In particular, we set $\tau = 0$ for 'instantaneous' Granger causality to assess the size and power of CoVaR and NCoVaR. We investigate the properties of the CoVaR-GC and NCoVaR-GC tests on the process with lagged dependence, i.e. $\tau = 1$. The sets A, C and D correspond with Equation (4) for CoVaR and NCoVaR, and with Equation (5) for CoVaR-GC and NCoVaR-GC, respectively.

For the simulations we set $a = 0.4$ and run 1000 independent realizations of the process in Equation (10), after a burn-in period of 100 time steps. Sample size-dependent bandwidths are set at the calculated MSE-optimal value and the fixed-range parameter is set at $\mu = 0.8$, as this seems to provide consistently good size and power properties (for comparison see Appendix F).

The detailed results for selected nominal size levels are summarized in Table 2. Presentation focuses on the risky quantile of $\gamma = 0.95$ but the results for $\gamma = 0.99$ and for above-quantile specification are available upon request.

There are three main observations that can be made based on this experiment. Firstly, the instantaneous CoVaR and NCoVaR measures seem to be nondirectional, at least in multivariate Gaussian cases, whereas the Granger causality-based specifications are sensitive to the direction of causality. Regarding the former, we note that [2] consider such a property a virtue, as the methodology captures the co-risk effects between variables. Our numerical exercise confirms this feature for both CoVaR and NCoVaR measures. Secondly, regarding the Granger causality measures, it seems that parametric estimation modestly over-rejects under the null, i.e. the methodology finds evidence for CoVaR-GC too often when it is actually absent. This over-rejection is confirmed for conventional nominal size levels and doesn't seem to diminish with increasing sample size (see Table 2). On the contrary, NCoVaR-GC displays much more conservative size properties.

Thirdly, for the VAR(1) process we find a considerable power gain for CoVaR relative to NCoVaR in small samples for the $\gamma = 0.95$ quantile. The differences evaporate as the

Table 2. Performance summary of CoVaR and NCoVaR methodologies in VAR class of models for selected nominal size levels (α).

		Instantaneous dependence				Granger causality			
		Size		Power		Size		Power	
α	n	CoVaR	NCoVaR	CoVaR	NCoVaR	CoVaR	NCoVaR	CoVaR	CoVaR
0.01	100	0.395	0.147	0.462	0.142	0.053	0.007	0.511	0.087
	200	0.621	0.391	0.625	0.374	0.024	0.004	0.681	0.186
	500	0.976	0.850	0.972	0.831	0.016	0.002	0.979	0.432
	1000	1.000	0.991	1.000	0.990	0.030	0.003	1.000	0.702
0.05	100	0.610	0.394	0.633	0.397	0.100	0.037	0.698	0.289
	200	0.837	0.661	0.825	0.675	0.070	0.035	0.857	0.473
	500	0.998	0.957	0.997	0.953	0.058	0.034	0.999	0.730
	1000	1.000	1.000	1.000	1.000	0.074	0.042	1.000	0.908
0.1	100	0.734	0.560	0.741	0.567	0.141	0.081	0.789	0.468
	200	0.907	0.803	0.904	0.803	0.115	0.098	0.923	0.648
	500	0.999	0.979	0.999	0.983	0.114	0.095	1.000	0.849
	1000	1.000	1.000	1.000	1.000	0.130	0.078	1.000	0.958

Notes: Actual rejection rates under $X \to Y$ causality (power) and $Y \to X$ causality (size) for different sample sizes (n), generated under process in Equation (10) with instantaneous dependence ($\tau = 0$) and with lagged dependence ($\tau = 1$). The quantile at which the risk is defined is set to $\gamma = 0.95$. We apply the MSE-optimal bandwidth and we set the fixed-range parameter to $\mu = 0.8$. The results are aggregated over 1000 simulations.

sample size increases and the instantaneous NCoVaR converges faster in power than the Granger causality setup. There are several possible reasons for these results. As confirmed by Rothe [26], parametric models are characterized by higher efficiency (in terms of MSE) for correctly specified models. Our simulation setup assumes the simplest process dynamics, which corresponds to the model specifications. Secondly, the strong power of CoVaR can be partially driven by its over-rejection bias, at least to some extent. Thirdly, the slower convergence for NCoVaR-GC can be attributed to the 'curse of dimensionality', i.e. less precise estimates in higher dimensions. There are several possible remedies to this problem, including data sharpening, principal components, projection pursuit or informative components analysis [14,27]. This topic is, however, beyond the scope of this paper and we leave it for further investigation.

3.2. Dependence in the second conditional moment

The second-moment dependence is analyzed through a prism of a class of (Generalized) Autoregressive Conditional Heteroskedasticity ((G)ARCH) models. In particular, we focus on a stylized bi-variate (G)ARCH process with (possibly lagged) volatility spillovers from $\{X_t\}$ onto $\{Y_t\}$ of the form

$$\begin{aligned} X_t &\sim N(0,1), \\ Y_t|X_t &\sim N(0, 1 + aX_{t-\tau}^2), \end{aligned} \tag{11}$$

where $a > 0$ is again tuning parameter and $N(0, \sigma^2)$ denotes the zero-centered normal distribution with variance σ^2.

The nomenclature and testing procedure are the same as for the VAR(1) process described in Section 3.1. In the simulations we set $a = 0.4$ and we focus on a risky quantile of $\gamma = 0.95$ but the results for alternative specifications are also available upon request. The size-size and size-power results for selected nominal size levels can be found in Table 3.

The experiments on the (G)ARCH process confirm the first two findings reported in Section 3.1, i.e. the limited directionality of the CoVaR and NCoVaR measures, and the over-rejection bias of CoVaR-GC, although the latter seems to be somehow contained. Regarding the power results, parametric CoVaR estimation is unable to detect the volatility spillovers generated by process in Equation (11), irrespective of the sample size. On the contrary, NCoVaR and NCoVaR-GC still capture this type of dependence, however, the power is considerably subdued compared to the VAR experiments, and it is also lower for smaller samples. While the latter can be viewed as a small sample problem of nonparametric setups, such effects are relatively modest in our examples. For instance, for the nominal level of $\alpha = 0.05$, the power gains associated with nonparametric procedures are already visible for $n = 200$.

The poor performance of CoVaR measures in (G)ARCH environment can be explained by its methodological design. As argued by Mainik and Schaanning [23], CoVaR is a correlation-driven measure. Having pointed this out, it misses any type of dependence in the higher moments of the conditional variable distributions. To put it pragmatically, $\hat{\beta}^i_\gamma$ estimates capture the average linear quantile effects between the variables of interest. In the case of (G)ARCH (or volatility spillovers) processes, the left-tail effects are offset by the right-tail equivalents, on average, which escalates the standard errors and reduces the

Table 3. Performance summary of CoVaR and NCoVaR methodologies in (G)ARCH class of models for selected nominal size levels (α).

		Instantaneous dependence				Granger causality			
		Size		Power		Size		Power	
α	n	CoVaR	NCoVaR	CoVaR	NCoVaR	CoVaR	NCoVaR	CoVaR	NCoVaR
0.01	100	0.023	0.006	0.042	0.008	0.042	0.002	0.061	0.008
	200	0.018	0.027	0.020	0.024	0.026	0.004	0.025	0.017
	500	0.018	0.078	0.017	0.066	0.029	0.000	0.025	0.022
	1000	0.009	0.146	0.018	0.120	0.015	0.005	0.011	0.065
0.05	100	0.061	0.080	0.088	0.086	0.096	0.028	0.129	0.070
	200	0.043	0.118	0.062	0.116	0.066	0.026	0.073	0.109
	500	0.045	0.235	0.062	0.229	0.086	0.029	0.063	0.158
	1000	0.045	0.363	0.063	0.324	0.066	0.043	0.060	0.240
0.1	100	0.114	0.176	0.139	0.169	0.150	0.070	0.174	0.141
	200	0.084	0.220	0.096	0.211	0.106	0.072	0.129	0.207
	500	0.084	0.381	0.115	0.351	0.141	0.071	0.107	0.290
	1000	0.091	0.505	0.105	0.485	0.108	0.083	0.117	0.380

Notes: Actual rejection rates under $X \to Y$ causality (power) and $Y \to X$ causality (size) for different sample sizes (n), generated under process in Equation (11) with instantaneous dependence ($\tau = 0$) and with lagged dependence ($\tau = 1$). The quantile at which the risk is defined is set to $\gamma = 0.95$. We apply the MSE-optimal bandwidth and we set the fixed-range parameter to $\mu = 0.8$. The results are aggregated over 1000 simulations.

statistical power of the method. Under such circumstances, NCoVaR framework provides a robust alternative.

4. Empirical illustration

To demonstrate the performance of NCoVaR and NCoVaR-GC we choose the Euro Area (EA) financial environment. In particular, we investigate the feedback loops (after [24]) between sovereigns and banks in selected EA Member States. Feedback loops are of particular importance for policy makers and regulators as they serve as a shock transmission channel during distress times. Banks, as key sovereign debt holders, are directly exposed to debt valuation and sovereign risk. On the other side, sovereigns are implicit guarantors of the banking sector and they took a huge hit on their debt accounts during the financial and subsequent sovereign debt crises. In essence, the prices of both instruments showed a high degree of co-movement in the recent history across different parts of the Europe [5].

The data used in the empirical analysis covers seven countries, i.e. two core EA Member States: Germany and France, and five vulnerable EA Member States: Spain, Portugal, Italy, Ireland and Greece. The Sovereign Price Index (SPI) is calculated from the price-yield relation of a 1-year zero-coupon bond, on the basis of a generic 1-year sovereign bond yield for each country. The Banking Price Index (BPI) is taken as the FTSE banking price index for each country. SPI come from Bloomberg and BPI come from Datastream. The missing observations are interpolated using linear interpolation technique but the results fully hold when excluding the missing values. We focus on daily observations between January 1994 and September 2016, however, due to data availability the precise ranges differ across countries. The exact coverage together with basic summary statistics are depicted in Appendix G.

In the empirical analysis, we look at the log returns of respective variables. We also standardize the data magnitude by the standard normal transformation. The fixed-bandwidth are set to $\mu = 0.8$ and the size-dependent window is chosen as indicated for the VAR(1) process.

The goal of the exercise is to quantify the sovereign-bank feedback loops on the sample countries and to compare the NCoVaR and NCoVaR-GC estimates against their parametric CoVaR equivalents. The main results for $\gamma = 0.95$ are depicted in Table 4.[6]

It can be readily observed that the NCoVaR estimates are somehow more conservative than the CoVaR estimates. This finding holds for both quantile specifications as well as across the variables and countries. In fact, this evidence is in line with our numerical conclusion that the linear CoVaR framework over-rejects under the null (see 3 for more detailed discussion). Looking at the results, the size of over-rejection is quite substantial. They strongly suggest CoVaR-GC spillovers from sovereign onto banks in all sample countries, with moderate support from NCoVaR-GC only in the case of Spain and Italy. Consequently, we consider the CoVaR-GC results to be inconclusive.

Looking at the instantaneous NCoVaR results, we find evidence for feedback loops in Spain and risk spillovers from sovereigns onto banks in Italy. NCoVaR-GC results appear to confirm the directional dependence from sovereigns onto banks in both countries, however, the spillovers from banks onto sovereigns in Spain disappear.

The parametric CoVaR results largely support the findings of simultaneous NCoVaR, suggesting also further bi-directional effects in Portugal and bank-to-sovereign spillovers in Greece.[7] Overall, with the exception of Ireland, the exercise confirms the differences in bank-sovereign feedback loops between vulnerable and core EA countries [24].

Interestingly, we find a surprisingly strong similarity between the CoVaR results and the simple linear OLS regressions between the variables (the linear results are depicted in Appendix G). We explain that by the proximity of both methodologies, which are designed to capture the dependence in the first moments, due to the fact that the quantile regression method puts a higher weight on central observations in estimating the tail co-dependence.

We also investigate the performance of both methodologies on two sub-samples, i.e. during the global financial crisis and the sovereign debt crisis. Following a stylized timeline given by St. Louis FED, we take that the former started on 27 February 2007 when Freddie Mac announced that it would no longer buy the most risky sub-prime mortgages and mortgage-related securities, and it finished on 13 April 2011 with the publication of the final report on the key causes of the crisis by the US Senate Permanent Subcommittee. Similarly, we assume that the sovereign debt crisis begun on 4 October 2009 with the

Table 4. Bank-sovereign feedback loops in selected euro area countries.

			CoVaR		NCoVaR		CoVaR-GC		NCoVaR-GC	
	X	Y	$X \to Y$	$Y \to X$	$X \to Y$	$Y \to X$	$X \to Y$	$Y \to X$	$X \to Y$	$Y \to X$
Germany	BPI	SPI						***		
France	BPI	SPI						***		
Spain	BPI	SPI	***	**	***	***	**	***		**
Italy	BPI	SPI	***	***		**	***	***		*
Portugal	BPI	SPI	***	***			***	***		
Ireland	BPI	SPI					***	***		
Greece	BPI	SPI	***				**	***		

Notes: BPI and SPI denote the Banking Price Index and Sovereign Price Index, respectively. Columns CoVaR and NCoVaR denote the instantaneous specifications, whereas columns CoVaR-GC and NCoVaR-GC correspond to Granger causality setups. ***, **, ** denote 1%, 5% and 10% significance levels. For NCoVaR and NCoVaR Granger causality tests we set $\mu = 0.8$. Risky quantiles are estimated at $\gamma = 0.95$.

PASOK's victory in the Greek Parliamentary elections, and lasted until the announcement of the new economic recovery plan for Europe on 30 May 2013. For transparency, the exact results are depicted in Appendix G.

It turns out that many of the EA bank-sovereign dependencies, in particular in the crisis-hit countries, are brought to light by the sovereign debt crisis. This somehow confirms different characteristics and propagation mechanisms between the two crises, exemplified by increased sovereign debt holdings of banking sectors in Spain, Italy, Portugal and Ireland [3].

4.1. Extra controls

As a robustness check we tests whether the nonlinear dependencies between banks and sovereigns found above can be explained by potential common-factor effects. As pointed out in Section 2.5, Theorems 2.1 and 2.2 require higher-order kernel smoothing and slower convergence rate of the bandwidth. In our example, we take the 4th-order Gaussian kernel, allowing to include two extra conditioning variables, which makes the bandwidth rates equal to $n^{-1/6}$ for NCoVaR and $n^{-1/7}$ for NCoVaR-GC. We take the bandwidth constants consistent with the higher-order kernels and under the no-dependence condition against extra co-variates. As common-factor benchmarks, we take the daily changes of the USD/EUR exchange rate and the STOXX Europe 600 equity index. Both time series cover the entire time span of the main variables of interest so that the number of observations matches for each country. The conditioning densities are evaluated around the median of the conditioning variables so that the results are indicative of the bank-sovereign dependencies in the absence of substantial shocks in the extra control variables, which in our example represent currency and equity markets.

The results are given in Table 5. The framework assumes $\gamma = 0.95$ but again the more conservative quantiles confirm the main findings. It can be observed that both linear CoVaR and CoVaR-GC vastly resemble the structure observed in the basic specification in Table 4. However, the statistical significance of dependence between variables weakens in the nonparametric results after controlling for the confounding variables. Again, the CoVaR results are vastly similar to the evidence delivered by a standard linear framework (see Appendix G for comparison).

The cross-sectional NCoVaR measure does suggest some evidence for a bank-sovereign feedback loop in Italy. Yet, it seems that the dynamics behind the Spanish feedback loop discovered in Table 4 is fully captured by the information present in the extra variables. Similarly, NCoVaR-GC detects weak evidence for a bank-onto-sovereign risk spillovers in Spain when conditioning for confounding variables.

The results indicate a clear difference how the parametric and nonparametric setups incorporate extra information from the confounding variables. It seems that the former remains intact whereas the latter is more agile. It may be that the extra information is present at higher moments of distribution of the confounding variables, which may be difficult to be discovered by linear frameworks. Similarly, the linear and nonlinear results may yield different predictive power. Although it is difficult to measure their exact performance, it seems that the lagged CoVaR results are largely driven by the central observations, and therefore deliver relatively less convincing results about the tail co-dependence. The exact nature of these phenomena is, however, beyond the scope of this paper.

Table 5. Bank-sovereign feedback loops in selected euro area countries correcting for common-factor effects.

			CoVaR		NCoVaR		CoVaR-GC		NCoVaR-GC	
	X	Y	$X \to Y$	$Y \to X$	$X \to Y$	$Y \to X$	$X \to Y$	$Y \to X$	$X \to Y$	$Y \to X$
Germany	BPI	SPI						**		
France	BPI	SPI						**		
Spain	BPI	SPI	***	***			**	***	*	
Italy	BPI	SPI	***	***	*	*	**			
Portugal	BPI	SPI	***	***			***	***		
Ireland	BPI	SPI					***	***		
Greece	BPI	SPI	***				**	***		

Notes: BPI and SPI denote the Banking Price Index and Sovereign Price Index, respectively. Columns CoVaR and NCoVaR denote the instantaneous specifications, whereas columns CoVaR-GC and NCoVaR-GC correspond to Granger causality setups. ***, **, * denote 1%, 5% and 10% significance levels. For NCoVaR and NCoVaR Granger causality tests we set $\mu = 0.8$. Risky quantiles are estimated at $\gamma = 0.95$.

5. Conclusions and discussion

NCoVaR and NCoVaR Granger causality (NCoVaR-GC) build a new methodological framework to assess co-risk relations, designed to capture the possible nonlinear effects. We derive the regular asymptotic properties of the NCoVaR tests and we confirm them numerically. Importantly, the framework is able to capture risk dependencies even in highly nonlinear environments, mimicking for instance volatility spillovers, which the standard CoVaR methodology is unable to capture. Moreover, we demonstrate that the CoVaR-GC measure is vulnerable to a false positive error.

We apply our methodology to assess the bank-sovereign co-risk relations in the Euro Area (EA). Our findings suggest substantial differences between core and vulnerable EA countries, as often highlighted in the literature [3,24]. The findings are preserved when conditioning for common-factor effects, which include currency and equity markets' dynamics. In particular, our findings suggest substantial instantaneous and lagged co-movement between bank and sovereign asset returns in Spain, Italy, Portugal and Greece, with negligible effects in Germany, France and Ireland. Furthermore, the evidence suggests that the bank-sovereign co-risk spillovers were much stronger during the sovereign debt crisis rather than during the global financial crisis, exemplifying a different nature between the two.

The NCoVaR framework can be of great use for macroprudential policy makers. Our extensive numerical and empirical studies suggest that NCoVaR tests provide more conservative estimates, compared to their parametric equivalents. In other words, standard CoVaR estimates may overprice the co-risk relevance between given entities or asset classes, possibly leading to inefficient allocation of macroprudential attention. To further test these predictions, the NCoVaR and NCoVaR-GC may be extended in the directions dictated by the semi-parametric extreme value theory.

The novel methodology reveals some intriguing phenomena on the nonlinear nature of the co-risk relations. A tempting idea is to investigate the underlying structures analytically in models of the aggregate economy. Such settings would allow to capture not only the risk contribution of relevant sectors but also measure the dynamics of aggregate disturbances. For practical applications, it would be also relevant to analyze algebraic similarities and differences between CoVaR and NCoVaR. Finally, we would like to mention that, as one of the reviewers noted, Theorems 1 and 2 depend on the metric used, and it might be worthwile to investigate the use of other metrics such as the Wasserstein distance.

Notes

1. Throughout the paper by causality we refer to causality in Granger's sense.
2. While the new risk metric is formally denoted by ΔNCoVaR, for the ease of exposition we will often skip the Δ symbol when talking about the concept rather than an exact formulation. We will keep the Δ when refering to the specific metric though.
3. Note that we slightly abuse language here. Strictly speaking, we should call the resulting estimators sample averages of kernel functions rather than U-statistics, since they are only asymptotically unbiased.
4. An alternative approach would be to estimate ΔCoVaR via a multivariate GARCH model. Under the Gaussian case, the main difference between two techniques is that whereas in the quantile regression case the estimate is proportional to the overall correlation between variables, the multivariate GARCH estimate is proportional to the instantaneous correlation. Consequently, neither method can capture the dependence structure which is not related to correlation, which builds an argument for the ΔNCoVaR metric proposed herein.
5. We also carry out an extra simulation study on an example of a GARCH-BEKK model, which is more widely applied by practitioners. The results are fully confirmed and desribed in detail in Appendix E.
6. The results for $\gamma = 0.99$ are largely in line but their statistical significance is weaker. For transparency reasons we do not report them in this paper but they are available from the authors upon request.
7. The fact that CoVaR results are asymmetric for Greece is a result of its subdued sovereign debt prices during the crisis period, which makes the SPI dynamics non-Gaussian. After removing the outliers, the results are bi-directional, as expected.

Acknowledgements

The authors would like to thank participants of seminars at the Federal Reserve Bank of Boston, Bank of Italy, University of Amsterdam, Bielefeld University and City University of New York (during the 22nd Annual Symposium of the Society for Nonlinear Dynamics and Econometrics) for their useful comments. Numerical simulations, reported in this paper, were run on the Lisa Compute Cluster. The authors also acknowledge financial support from the European Doctorate in Economics - Erasmus Mundus. The views and opinions expressed in this paper are those of the authors and do not necessarily reflect the official policy or position of the EIB or any other affiliated institution.

Disclosure statement

No potential conflict of interest was reported by the author(s).

References

[1] V. Acharya, I. Drechsler, and P. Schnabl, *A pyrrhic victory? Bank bailouts and sovereign credit risk*, J. Finance. 69 (2014), pp. 2689–2739.
[2] T. Adrian and M.K. Brunnermeier, *CoVaR*, Am. Econ. Rev. 106 (2016), pp. 1705–1741.
[3] P. Angelini, G. Grande, and F. Panetta, *The negative feedback loop between banks and sovereigns*. Questioni di Economia e Finanza (Occasional Papers) 213, Bank of Italy, Economic Research and International Relations Area, 2014.
[4] T. Bouezmarni, J.V. Rombouts, and A. Taamouti, *Nonparametric copula-based test for conditional independence with applications to granger causality*, J. Bus. Econ. Stat. 30 (2012), pp. 275–287.

[5] M. Caporin, L. Pelizzon, F. Ravazzolo, and R. Rigobon, *Measuring sovereign contagion in Europe*. Working Paper 05, Norges Bank, 2012.
[6] M. Chinazzi and G. Fagiolo, *Systemic risk, contagion, and financial networks: A survey*. LEM Working Paper 2013/08, Sant'Anna School of Advanced Studies, Pisa, 2013.
[7] M. Denker and G. Keller, *On U-statistics and v. mises' statistics for weakly dependent processes*, Zeitschrift Für Wahrscheinlichkeitstheorie Und Verwandte Gebiete 64 (1983), pp. 505–522.
[8] M. Denker and G. Keller, *Rigorous statistical procedures for data from dynamical systems*, J. Stat. Phys. 44 (1986), pp. 67–93.
[9] C. Diks and V. Panchenko, *A new statistic and practical guidelines for nonparametric granger causality testing*, J. Econ. Dyn. Control 30 (2006), pp. 1647–1669.
[10] C. Diks and M. Wolski, *Nonlinear granger causality: Guidelines for multivariate analysis*, J. Appl. Econom. 31 (2016), pp. 1333–1351.
[11] EIB, *Investment and investment finance in Europe*, European Investment Bank, Luxembourg, 2016.
[12] B. Gao and R.e. Ren, *The topology of a causal network for the Chinese financial system*, Phys. A: Stat. Mech. Appl. 392 (2013), pp. 2965–2976.
[13] C.W.J. Granger, *Investigating causal relations by econometric models and cross-spectral methods*, Econometrica 37 (1969), pp. 424–438.
[14] P. Hall and M.C. Minnotte, *High order data sharpening for density estimation*, J. R. Stat. Soc. Series B Stat. Methodol. 64 (2002), pp. 141–157.
[15] B. Hansen, *Lecture notes on nonparametrics*. Working Paper, University of Wisconsin – Madison, 2009.
[16] Z. He and A. Krishnamurthy, *A macroeconomic framework for quantifying systemic risk*. Working Paper 233, National Bank of Belgium, 2012.
[17] Y. Hong, Y. Liu, and S. Wang, *Granger causality in risk and detection of extreme risk spillover between financial markets*, J. Econom. 150 (2009), pp. 271–287.
[18] X. Huang, H. Zhou, and H. Zhu, *Systemic risk contributions*. Working Paper 60, Bank for International Settlements, 2010.
[19] K. Jeong, W.K. Härdle, and S. Song, *Nonparametric test for causality in quantile*, Econ. Theory. 28 (2012), pp. 861–887.
[20] T. Kitamura, I. Muto, and I. Takei, *Loan interest rate pass-through and changes after the financial crisis: Japan's evidence*, J. Jpn. Int. Econ. 42 (2016), pp. 10–30.
[21] Q. Li and J.S. Racine, *Nonparametric Econometrics: Theory and Practice*, Princeton University Press, Princeton and Oxford, 2007.
[22] O. Linton and P. Gozalo, *Testing conditional independence restrictions*, Econom. Rev. 33 (2014), pp. 523–552.
[23] G. Mainik and E. Schaanning, *On dependence consistency of coVaR and some other systemic risk measures*, Stat. Risk Model. 31 (2014), pp. 49–77.
[24] F.L. Ohnsorge, M. Wolski, and S.Y. Zhang, *Safe havens, feedback loops, and shock propagation in global asset prices*. IMF Working Paper 14/81, International Monetary Fund, 2014.
[25] M. Paries, D. Moccero, E. Krylova, and C. Marchini, *The retail bank interest rate pass-through: The case of the euro area during the financial and sovereign debt crisis*. Occasional Paper Series of the European Central Bank No 154, European Central Bank, 2014.
[26] C. Rothe, *Nonparametric estimation of distributional policy effects*, J. Econom. 155 (2010), pp. 56–70.
[27] D.W. Scott, *Multivariate Density Estimation. Wiley Series in Probability and Mathematical Statistics*, Wiley-Interscience, New York, 1992.
[28] M. Segoviano and C. Goodhart, *Banking stability measures*. IMF Working Paper No. 09/04, IMF, 2009.
[29] L. Su and H. White, *A nonparametric hellinger metric test for conditional independence*, Econ. Theory. 24 (2008), pp. 829–864.
[30] L. Su and H. White, *Testing conditional independence via empirical likelihood*, J. Econom. 182 (2014), pp. 27–44.

[31] J. von Borstel, S. Eickmeier, and L. Krippner, *The interest rate pass-through in the euro area during the sovereign debt crisis*, J. Int. Money. Finance. 68 (2016), pp. 386–402.
[32] X. Wang and Y. Hong, *Characteristic function based testing for conditional independence: A nonparametric regression approach*, Econ. Theory. 34 (2018), pp. 815–849.
[33] C. Zhou, *Are banks too big to fail? Measuring systemic importance of financial institutions*, Int. J. Cent. Bank. 6 (2010), pp. 205–250.