

TI 2024-036/III
Tinbergen Institute Discussion Paper

A Novel Test for the Presence of Local Explosive Dynamics

*F. Blasques*¹

*S.J. Koopman*¹

*G. Mingoli*¹

*S. Telg*¹

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and Vrije Universiteit Amsterdam.

Contact: discussionpapers@tinbergen.nl

More TI discussion papers can be downloaded at <https://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900

A Novel Test for the Presence of Local Explosive Dynamics*

F. Blasques, S.J. Koopman, G. Mingoli[†], and S. Telg

Vrije Universiteit Amsterdam

Tinbergen Institute

May 24, 2024

Abstract

In economics and finance, speculative bubbles take the form of locally explosive dynamics that eventually collapse. We propose a test for the presence of speculative bubbles in the context of mixed causal-noncausal autoregressive processes. The test exploits the fact that bubbles are anticipative, that is, they are generated by an extreme shock in the forward-looking dynamics. In particular, the test uses both path level deviations and growth rates to assess the presence of a bubble of given duration and size, at any moment of time. We show that the distribution of the test statistic can be either analytically determined or numerically approximated, depending on the error distribution. Size and power properties of the test are analyzed in controlled Monte Carlo experiments. An empirical application is presented for a monthly oil price index. It demonstrates the ability of the test to detect bubbles and to provide valuable insights in terms of risk assessments in the spirit of Value-at-Risk.

Keywords: noncausality, bubbles, testing, date-stamping, risk assessment.

JEL codes: C22, E31, E37.

*F. Blasques acknowledges the financial support of the Dutch Science Foundation (NWO) under grant Vidi.195.099.

[†]Corresponding author: Gabriele Mingoli, g.mingoli@vu.nl.

1 Introduction

Economic and financial bubbles can have significant disruptive impact for households, companies, governments, and investors. The 2000's housing price bubble in the United States, resulted in the doubling of real-estate prices in a span of just seven years, from 2000 to 2007, and the subsequent collapse triggering the subprime crises of 2008. More recently, the cryptocurrency Bitcoin experienced multiple bubble events. The bubble of 2017-2018 led to a collapse of about 65% in the price of Bitcoin from 6 January to 6 February 2018, leaving many investors with heavy losses. On 26 November 2018, the price of Bitcoin fell by 80%. Similarly, from January 2018 to September 2018, other cryptocurrencies collapsed in value by 80%, making the 2018 cryptocurrency crash worse than the 1995-2000 dot-com bubble collapse of 78%. Early-warning indicators of bubble dynamics and potential bubble collapse size, if available, can play a fundamental role in measuring, understanding and mitigating economic and financial risk.

The recent literature on mixed causal-noncausal autoregressive (MAR) models may offer a systematic way of modeling time-series with local explosive dynamics. As a data generating processes, the MAR model is able to generate short-lived bubbles with characteristics similar to those observed in the data; see e.g., [Gourieroux and Jasiak \(2016\)](#) and [Fries and Zakoïan \(2019\)](#). Recently, research has focused on the forecasting of extreme events using MAR models based on closed-form expressions and approximation methods ([Lanne et al., 2012](#); [Gouriéroux and Zakoïan, 2017](#); [Hecq and Voisin, 2021](#)). However, there is still a lack of formal statistical tests that leverage MAR models to test for the presence of bubbles in the data.

Existing tests for bubbles currently available in the literature extend the research on unit roots to effectively test for the presence of locally explosive dynamics; see e.g. [Phillips and Yu \(2011\)](#) and [Phillips et al. \(2011, 2015\)](#). These papers offer an important set of tools to test for bubbles, but they can be complicated to implement and computationally expensive to run. Their test is based on an autoregressive (AR) model specification $y_t = \rho y_{t-1} + \varepsilon_t$ for which recursive augmented Dickey-Fuller (ADF) tests are used to check for a random walk ($H_0 : \rho = 1$) versus explosive trajectories of the data ($H_1 : \rho > 1$). We propose an alternative method, which is simpler to use, allows for date-stamping a bubble, and is based on the well-studied intuition behind MAR dynamics and the modeling of bubbles. More specifically, we make use of a stationary model in which the linear noncausal component is exclusively responsible for the locally explosive dynamics.




Figure 1: The trajectory of a causal autoregressive process $y_t = \rho y_{t-1} + \varepsilon_t$ (left) and a noncausal autoregressive process $y_t = \rho y_{t+1} + \varepsilon_t$ (right), with $\rho = 0.9$, $\varepsilon_{50} = 1$ and zero otherwise.

The proposed method provides a straightforward way to test for bubbles, which also takes into consideration the size and the date of collapse. Further, it allows us to build confidence bounds for the trajectory of the bubble.

We can illustrate the simple reasoning behind the testing procedure by contrasting a stationary AR process $y_t = \rho y_{t-1} + \varepsilon_t$ with its noncausal counterpart $y_t = \rho y_{t+1} + \varepsilon_t$ where we set $\rho = 0.9$. We consider the simplified scenario where all error terms $\{\varepsilon_t\}_{t=1}^{T=100}$ are zero, except at displacement $t = 50$ where it equals unity. The trajectories of these AR processes are shown in Figure 1. In both scenarios, a single extreme error observation is responsible for the dynamic structure. However, in the causal scenario (left), the shock is introduced to the system at $t = 50$ and its effect fades out at a geometric rate. This means that prior to the shock, for example at $t = 40$, it is impossible to anticipate the shock and its effect on the process. This is different for the noncausal process (right), in which the process grows geometrically towards the extreme observation at $t = 50$, which makes the process anticipative. At $t = 40$, the process already contains information about the shock and this knowledge can be exploited to test whether a bubble emerges in real data by comparing whether its current trajectory is compatible with a noncausal process that has an extreme observation of a certain size h steps ahead (in our illustration of Figure 1, $h = 10$).

Our testing procedure is also valid when the bubble is observed in the presence of stochastic trends, which is often the case in real-life applications. The treatment of non-stationary dynamics and stochastic trends has been until recently one of the main challenges for analysing time series

using MAR models. A common treatment is to detrend the time series prior to modeling; see e.g. [Hecq and Voisin \(2022\)](#) using the Hodrick-Prescott filter, or [Hencic and Gouriéroux \(2015\)](#) using polynomial detrending. More recently, a filtering approach is proposed by [Blasques et al. \(2023\)](#) which is based on a joint model that includes both a random-walk fundamental component and a bubble component that is specified as a MAR model. Both components are treated in the filter updating simultaneously. Hence, the filter also allows the test to be implemented in non-stationary settings. These findings lead to an integrated testing framework for bubbles based on MAR methodology and provide clarity on whether and how to detrend in real-time which is not always unequivocal for alternative testing procedures.

In this study, we develop the test for the presence and location of explosive dynamics and we show that the proposed test has favorable size and power properties in finite samples. For this purpose, we carry out a Monte Carlo simulation study in which we consider different scenarios for the null parameters, the shock size and the horizon, as well as the error distribution. In the empirical study, we highlight the flexibility of the testing procedure and show that it can be further extended to a Value-at-Risk (VaR) application, which provides practitioners with the possibility to make a reasonable risk assessment *a-priori* in the presence of bubble dynamics in the data. Python code to implement the proposed testing method is provided.

This paper is organized as follows. In [Section 2](#) we introduce the mixed causal-noncausal model and an estimation method. [Section 3](#) provides our proposed testing procedure to detect the presence of bubbles and [Section 4](#) investigates its performance in various settings. In [Section 5](#), we illustrate the method in an application on an oil price index and show how it can be used to make a reasonable risk assessment. [Section 6](#) concludes.

2 Mixed causal-noncausal autoregressive model

We specify the stationary mixed causal-noncausal autoregressive (MAR) model in detail and discuss the estimation of its parameters. We further consider a non-stationary extension of the MAR model for the purpose of increasing its empirical relevance.

2.1 Stationary MAR model

Consider a mixed causal-noncausal autoregressive (MAR) process $\{y_t\}_{t \in \mathbb{Z}}$ given by

$$\phi(L)\psi(L^{-1})y_t = \varepsilon_t, \quad (1)$$

where L is the lag operator, $\phi(L) = 1 - \phi_1 L - \dots - \phi_r L^r$ is the lag polynomial of order $r \in \mathbb{N}^+$ with fixed coefficients ϕ_1, \dots, ϕ_r , L^{-1} is the lead operator, $\psi(L^{-1}) = 1 - \psi_1 L^{-1} - \dots - \psi_s L^{-s}$ is the lead polynomial of order $s \in \mathbb{N}^+$ with fixed coefficients ψ_1, \dots, ψ_s , and $\{\varepsilon_t\}_{t \in \mathbb{Z}}$ is a sequence of independent and identically distributed (i.i.d.) non-Gaussian error terms. We denote this process as $\text{MAR}(r, s)$, where the total autoregressive order equals $p = r + s$. To ensure stationarity of the process, we assume that the roots of both polynomials lie strictly outside the unit circle, i.e.,

$$\phi(z) \neq 0 \quad \text{for } |z| \leq 1 \quad \text{and} \quad \psi(z) \neq 0 \quad \text{for } |z| \leq 1.$$

We follow [Lanne and Saikkonen \(2011\)](#) in defining unobserved causal and noncausal components of the MAR process as

$$\begin{aligned} u_t = \psi(L^{-1})y_t &\iff \phi(L)u_t = \varepsilon_t, \\ v_t = \phi(L)y_t &\iff \psi(L^{-1})v_t = \varepsilon_t, \end{aligned} \quad (2)$$

which provide an effective way to disentangle the dynamics due to backward- and forward-looking behavior in the MAR model, as represented by u_t and v_t , respectively. In this paper, the filtered processes $\{u_t\}_{t \in \mathbb{Z}}$ and $\{v_t\}_{t \in \mathbb{Z}}$ play a crucial role in the bubble testing procedure. Whereas the

process $\{y_t\}_{t \in \mathbb{Z}}$ has a two-sided moving average representation¹

$$y_t = \sum_{j=-\infty}^{\infty} \beta_j \varepsilon_{t+j}, \quad (3)$$

where the weights β_j are functions of both the polynomial coefficients ϕ_1, \dots, ϕ_r and ψ_1, \dots, ψ_s , the filtered processes $\{u_t\}_{t \in \mathbb{Z}}$ and $\{v_t\}_{t \in \mathbb{Z}}$ can be expressed as one-sided moving averages given by

$$u_t = \sum_{j=0}^{\infty} \gamma_j \varepsilon_{t-j} \quad \text{and} \quad v_t = \sum_{j=0}^{\infty} \zeta_j \varepsilon_{t+j}, \quad (4)$$

where the weights γ_j and ζ_j are functions of the polynomial coefficients ϕ_1, \dots, ϕ_r and ψ_1, \dots, ψ_s , respectively. In both cases, the weights are exponentially decaying for an increasing j due to the assumption of a stationary MAR. It is well-documented in the literature that bubbles are exclusively generated in the noncausal component $\{v_t\}_{t \in \mathbb{Z}}$, which is often called anticipative. As explained in [Gouriéroux and Zakoïan \(2017\)](#) and [Hecq et al. \(2016\)](#), the development of a bubble is due to the presence of a large shock ε_τ , at some future time point τ relative to current time point $t < \tau$. Equation (4) reveals that the effect of such a future shock is already contained in the noncausal component v_t , while this is not the case for u_t . Given the exponential decay of the coefficient sequence $\{\zeta_j\}_{j \in \mathbb{N}_0}$, the sequence $\{v_t\}_{t \in \mathbb{Z}}$ build up towards this extreme value and reaches its peak at time τ . At time points beyond τ , that is time points $\tau + 1, \tau + 2, \dots$, the large disturbance ε_τ no longer contributes to the formation of v_t (and thus y_t), which causes the process to return to the baseline path. In other words, the bubble crashes at time $\tau + 1$. This insight provides a way to compare whether available data at time t is compatible with a bubble crashing $h + 1$ -step ahead, that is, at $\tau = t + h + 1$, which we exploit in our testing procedure outlined in [Section 3](#).

2.2 Parameter Estimation

The MAR process in (1) requires a non-Gaussian error term to ensure identifiability of the causal and noncausal parts of the model. For this reason, different estimation methods have been proposed, which can be classified as least squares ([Gouriéroux and Zakoïan, 2017](#); [Fries and](#)

¹The process $\{y_t\}_{t \in \mathbb{Z}}$ exists almost surely under quite weak conditions, such as $E|\varepsilon_t|^\kappa < \infty$ for $\kappa > 0$ and $\sum_{j=-\infty}^{\infty} |\beta_j|^\kappa < \infty$ (see e.g., [Brockwell and Davis, 1991](#)) which is automatically satisfied under the assumption on the roots of the polynomials being outside the unit circle.

Zakoian, 2019), spectral methods (Hecq and Velásquez-Gaviria, 2022) and approximate maximum likelihood (Breidt et al., 1991; Lanne and Saikkonen, 2011). The first two methods have the advantage that they do not require the assumption of an error distribution, but they do entail either solving an extensive matching problem (based on extreme residuals clustering) or evaluating relatively non-standard criterion functions. In our study we opt for the approximate maximum likelihood (AML) estimator, as it is well-established in the literature. However, the choice of estimation method does not have a direct influence on the proposed bubble testing procedure set out in Section 3.

We follow the same procedure as Lanne and Saikkonen (2011). More specifically, we assume that the distribution of the non-Gaussian error term ε_t has a (Lebesgue) density $f_\sigma(x; \boldsymbol{\lambda}) = \sigma^{-1}f(\sigma^{-1}x; \boldsymbol{\lambda})$ satisfying the regularity conditions of Andrews et al. (2006), with the $d \times 1$ parameter vector $\boldsymbol{\lambda}$ that contains the scale parameter $\sigma > 0$ and other unknown coefficients in the scaled density function $f(\cdot)$. We have $r \times 1$ vector $\boldsymbol{\phi} = (\phi_1, \dots, \phi_r)'$ and $s \times 1$ vector $\boldsymbol{\psi} = (\psi_1, \dots, \psi_s)'$ containing the causal and noncausal autoregressive parameters, respectively. Their permissible parameter space is defined by the stationarity condition that the roots of both polynomials lie strictly outside the unit circle. We can collect all parameters in the $(r + s + d) \times 1$ parameter vector $\boldsymbol{\theta} = (\boldsymbol{\phi}', \boldsymbol{\psi}', \boldsymbol{\lambda}')'$. The approximate log-likelihood function for $\{y_t\}_{t=1}^T$ is given by

$$l_T(\boldsymbol{\theta}) = \sum_{t=r+1}^{T-s} g_t(\boldsymbol{\theta}) = \sum_{t=r+1}^{T-s} \log f_\sigma(\phi(L)\psi(L^{-1})y_t; \boldsymbol{\lambda}), \quad (5)$$

where g_t is the individual log-likelihood contribution at time t which could also be expressed solely in terms of $\{u_t\}_{t=r+1}^{T-s}$ and $\{v_t\}_{t=r+1}^{T-s}$ using the relations in (2). Maximizing $l_T(\boldsymbol{\theta})$ over permissible values of $\boldsymbol{\theta}$ gives the AML estimator of $\boldsymbol{\theta}$.

2.3 The stochastic trend plus MAR model

To enable the analysis of bubbles for non-stationary time series, we also consider the stochastic trend plus MAR model of Blasques et al. (2023). We refer to this model as the MAR stochastic trend (MARST) model and we formulate the model for an observed non-stationary time series $\{z_t\}_{t \in \mathbb{Z}}$. This observed time series is then decomposed by an unobserved random walk process $\{\mu_t\}_{t \in \mathbb{Z}}$ (stochastic trend) and an unobserved stationary MAR process $\{y_t\}_{t \in \mathbb{Z}}$. Hence, in a stationary setting, $\mu_t = 0$ and $z_t = y_t$ is observed, while in a non-stationary setting $z_t = \mu_t + y_t$

is observed but y_t only is not. The MARST model is given by

$$\begin{aligned} z_t &= \mu_t + y_t, & \mu_t &= \delta + \mu_{t-1} + \alpha \varepsilon_{t-s}, \\ y_t &= [\phi(L)\psi(L^{-1})]^{-1} \sigma \varepsilon_t, \end{aligned} \tag{6}$$

where μ_t represents the stochastic trend with a drift coefficient δ (it can be set to zero) and with scale parameter α that determines the amplitude of the trend update, while the MAR specification remains as before. We explicitly include the scale parameter σ in the specification of the MAR process $\{y_t\}_{t \in \mathbb{Z}}$, in order to differentiate it from the scaling of the trend by α . Hence, the error ε_t has density function $f_\sigma(\varepsilon_t; \boldsymbol{\lambda})$ as introduced in Section 2.2. The trend update is formulated as a random walk with drift δ , where the trend innovation is set to ε_{t-s} where s is the lead order of the noncausal lead polynomial $\psi(L^{-1})$. The process in (6) can be estimated using (5) by augmenting the parameter vector $\boldsymbol{\theta}$ with α and δ and recognizing that in a non-stationary setting, y_t can be regarded as the detrended time series $y_t = z_t - \mu_t(\boldsymbol{\theta})$ where we make the dependence of μ_t on the parameter vector $\boldsymbol{\theta}$ explicit. A filtering procedure for extracting the random walk component $\mu_t(\boldsymbol{\theta})$ based on data available at time t is developed in Blasques et al. (2023). The filtering also enables the actual detrending $y_t = z_t - \mu_t(\boldsymbol{\theta})$ which is repeated whenever $\boldsymbol{\theta}$ changes its value. This notion is particularly relevant during AML parameter estimation which is based on the numerical maximization of the approximate log-likelihood function.

The consideration of the MARST model is motivated by the fact that bubbles in economic and financial data are often observed in the presence of non-stationary features in the time series. In particular, trending time series are often encountered in economics. Various ad-hoc methods have been employed in the MAR literature to remove such non-stationary features, including polynomial regression (Hencic and Gouriéroux, 2015) and the Hodrick-Prescott filter (Hecq and Voisin, 2022). Although these methods may work in practice, their theoretical implications are not documented. The methodology developed by Blasques et al. (2023) for the MARST model ensures strong consistency and asymptotic normality of the AML estimator of $\boldsymbol{\theta}$, for a set of rather flexible assumptions. Moreover, an important feature of the MARST specification is that it allows for detrending in real-time, while the other methods are typically applied *ex-post*. This ability and its practical relevance make it a strong case to proceed in empirical applications, where the bubble is oftentimes observed in time series where non-stationary features are present.

3 A Test for the Presence of Speculative Bubbles

Given the discussion in Section 2.1 and the decomposition implied by equations (2) and (4), we assume that bubbles arise in v_t from the anticipation of a sizeable shock in the near future. If one of the upcoming shocks is a draw far enough in the tails of the error distribution, a relevant bubble can appear in the process y_t . Without loss of generality, we solely focus on extreme positive draws in what follows. To be able to introduce our proposed testing procedure, we formalize this concept as follows. For a random variable X , we consider whether a realization is a draw from the upper quantile of the underlying distribution, i.e. $q_X(\alpha) = \{x : F_X(x) \geq 1 - \alpha\}$, where $F_X(x)$ denotes the cumulative distribution function of a random variable X and α is the chosen significance level. Since we typically require large positive draws in our context, we define the quantile in this alternative way that isolates the right tail of the distribution. We denote by $d_X(\alpha)$ a draw from the quantile $q_X(\alpha)$ and the parameter h represents the horizon at which we expect the bubble process to reach its peak.

The results in this section depend crucially on a correct description of the marginal distribution of the random variable Z_T , as well as the joint distribution of the bivariate random variable (Z_T, Q_T) , which are given by

$$Z_T = \sum_{\substack{j=0 \\ j \neq h}}^{\infty} \zeta_j \varepsilon_{T+j} \quad \text{and} \quad Q_T = \sum_{\substack{j=1 \\ j \neq h}}^{\infty} (\zeta_{j-1} - \zeta_j) \varepsilon_{T+j} - \varepsilon_T,$$

with marginal and joint distributions denoted $p_Z(\boldsymbol{\zeta}, \boldsymbol{\lambda})$, and $p_{ZQ}(\boldsymbol{\zeta}, \boldsymbol{\lambda})$, respectively, where $\boldsymbol{\lambda}$ contains the distributional parameters including the scale parameter σ (and possibly, location, shape, symmetry, etc.) of the error distribution $f_\sigma(\varepsilon_t; \boldsymbol{\lambda})$. In Proposition 1, we define a test statistic which depends on the noncausal component $v_t = \phi(L)y_t$, and parameters α and h which account for the size of a future shock, and the horizon at which the shock ε_{T+h} occurs. Specifically, this proposition is useful in establishing $p_Z(\boldsymbol{\zeta}, \boldsymbol{\lambda})$ as the distribution under the null hypothesis $\varepsilon_{T+h} = d_\varepsilon(\alpha)$ for the infeasible test statistic

$$k_T = v_T - \zeta_h d_\varepsilon(\alpha),$$

which can be used as a finite-sample approximation of the distribution for the feasible counterpart $\hat{k}_T = \hat{v}_T - \hat{\zeta}_h \hat{d}_\varepsilon(\alpha)$ where $\hat{v}_T = \hat{\phi}(L)y_t$ and $\hat{\zeta}_h$ are based on the AML estimates $\hat{\phi}$ and $\hat{\psi}$ respectively

and $\hat{d}_\varepsilon(\alpha)$ depends on the estimated distribution of the residuals.

Proposition 1. *Let the process $\{y_t\}_{t \in \mathbb{Z}}$ be generated by a $MAR(r, s)$ model. Then, given $\zeta_h \in \mathbb{R}$, and under the null hypothesis of a shock corresponding to a certain quantile of size $q_\varepsilon(\alpha)$, occurring h periods ahead, the test statistic k_T satisfies*

$$k_T \sim p_Z(\boldsymbol{\zeta}, \boldsymbol{\lambda}).$$

Proof: Immediate from the fact that $v_T = \sum_{j \neq h}^{\infty} \zeta_j \varepsilon_{T+j} + \zeta_h \varepsilon_{T+h}$, so that $k_T = \sum_{j \neq h}^{\infty} \zeta_j \varepsilon_{T+j} = Z_T$.

We notice that, in practice, the unknown ζ_j 's required to calculate v_t and ζ_h , can be treated as nuisance parameters and substituted by their plug-in estimates $\hat{\zeta}_j$ to obtain a feasible statistic \hat{k}_T . It is in that sense that the distribution stated in Proposition 1 for the infeasible statistic k_T is only a finite sample approximation for the distribution of \hat{k}_T . Naturally, the quality of the approximation will depend on the sample size T as the variance of each $\hat{\zeta}_j$ vanishes with $T \rightarrow \infty$ and we obtain convergence to ζ_j .

Proposition 2. *Let Assumptions (A1) - (A7) in [Lanne and Saikkonen \(2011\)](#) hold. Then the feasible test statistic $\hat{k}_T = \hat{v}_T - \hat{\zeta}_h \hat{d}_\varepsilon(\alpha)$ satisfies $\hat{k}_T \xrightarrow{d} p_Z(\boldsymbol{\zeta}, \boldsymbol{\lambda})$ as $T \rightarrow \infty$.*

Proof: Immediate from Slutsky's theorem given the fact that Assumptions (A1) - (A7) in [Lanne and Saikkonen \(2011\)](#) ensure $(\hat{\boldsymbol{\zeta}}, \hat{\boldsymbol{\lambda}}) \xrightarrow{p} (\boldsymbol{\zeta}, \boldsymbol{\lambda})$ as $T \rightarrow \infty$. The consistency of the AML estimator for the MAR model is proven for a general class of non-Gaussian distributions specified in [Andrews et al. \(2006\)](#). Whereas the error distributions are confined to having finite variance, [Gourieroux and Jasiak \(2016\)](#) provide evidence that AML is also consistent in the context of Cauchy errors.

Naturally, depending on the underlying distribution of the innovations $f_\sigma(\varepsilon_t; \boldsymbol{\lambda})$, the approximate test statistic's distribution $p_Z(\hat{\boldsymbol{\zeta}}, \hat{\boldsymbol{\lambda}})$ will be available in closed form, or may require Monte Carlo approximation. Analytic tractability is available, for example, for the case of Cauchy innovations which is common in the MAR literature.²

²This result can be extended to the general class of α -stable distributions, see [Fries and Zakoian \(2019\)](#) for more details.

Remark 1. Let the innovations be independent identically distributed innovations satisfying $\varepsilon_0 \sim \text{Cauchy}(\alpha, \beta)$. Then,

$$k_T \sim \text{Cauchy}\left(\sum_{j \neq h}^{\infty} \zeta_j, \xi(\beta, \zeta)\right),$$

where ξ is a function of (β, ζ) that defines the scale of the resulting Cauchy.

In practice, suppose that a large shock takes place $h \in \mathbb{N}$ displacements ahead, i.e., $\varepsilon_{T+h} \in q_\varepsilon(\alpha)$, such that we can write

$$v_T = \sum_{j \neq h}^{\infty} \zeta_j \varepsilon_{T+j} + \zeta_h \varepsilon_{T+h}.$$

Under the null that a bubble reaches its peak h steps ahead from now (and consequently crashes $h + 1$ steps ahead) with a size determined by the $q_\varepsilon(\alpha)$ quantile of the underlying distribution, we can now use

$$\hat{k}_T = \hat{v}_T - \hat{\zeta}_h \hat{d}_\varepsilon(\alpha) \sim Z_T,$$

with $Z_T = \sum_{j \neq h}^{\infty} \zeta_j \varepsilon_{T+j}$ to perform testing. The quantity $\Pr(Z_T < \hat{k}_T)$ measures the probability of the level of the data to be compatible with a trajectory including a bubble assumed by the null. In this way, we can test if the data is significantly different from a bubble of a given size, as determined by the quantile $q_\varepsilon(\alpha)$, at a certain horizon h . Depending on the error distribution, the distribution of Z_T can be determined theoretically or approximated by simulations to tabulate appropriate critical values. We refer to this test as the “level test”.

We can further enrich the specification of the test by using growth rates in addition to the levels. For this purpose, we consider temporal changes in v_T as

$$\Delta v_{T+1} = \sum_{\substack{j=1 \\ j \neq h}}^{\infty} (\zeta_{j-1} - \zeta_j) \varepsilon_{T+j} - \varepsilon_T + (\zeta_{h-1} - \zeta_h) \varepsilon_{T+h}.$$

In a similar fashion as for the levels, and under the null hypothesis that $\varepsilon_{T+h} = d_\varepsilon(\alpha)$, we define the test statistic

$$r_T = \Delta v_{T+1} - (\zeta_{h-1} - \zeta_h) d_\varepsilon(\alpha).$$

In Proposition 3, we establish that we can jointly test $\eta_t = (k_t, r_t)$ using their random variables representation $N_t = (Z_t, Q_t)$ under the null hypothesis. We refer to this test as the “level-growth test”.

Proposition 3. *Let the process $\{y_t\}_{t \in \mathbb{Z}}$ be generated by a $MAR(r, s)$ model. Then, given $\zeta_h \in \mathbb{R}$, and under the null hypothesis of a shock corresponding to a certain quantile of size $q_\varepsilon(\alpha)$, occurring h periods ahead, the test statistic N_T satisfies*

$$(k_t, r_t) \sim p_{ZQ}(\boldsymbol{\zeta}, \boldsymbol{\lambda}).$$

Proof: Immediate from the fact that $k_T = Z_T$ and $r_T = \Delta v_{T+1} - (\zeta_{h-1} - \zeta_h)\varepsilon_{T+h} = \sum_{\substack{j=1 \\ j \neq h}}^{\infty} (\zeta_{j-1} - \zeta_j)\varepsilon_{T+j} - \varepsilon_T = Q_T$.

Similar to the first test, the limit distribution will be known in closed form, or may require Monte Carlo approximation, depending on the distribution of the innovations.

Remark 2. *Let the innovations be independent identically distributed innovations satisfying $\varepsilon_0 \sim \text{Cauchy}(\alpha, \beta)$. Then, (k_t, r_t) is multivariate Cauchy.*

As noted before, the unknown ζ_j 's can be substituted by their plug-in estimates $\hat{\zeta}_j$ to obtain a feasible statistic $\hat{\eta}_t = (\hat{k}_t, \hat{r}_t)$. As such, the distribution stated in Proposition 2 for the infeasible statistic (k_t, r_t) is a finite sample approximation for the distribution of (\hat{k}_t, \hat{r}_t) , where

$$\hat{r}_T = \Delta \hat{v}_{T+1} - (\hat{\zeta}_{h-1} - \hat{\zeta}_h)\hat{d}_\varepsilon(\alpha) \sim Q_T.$$

The performance of the test procedure depends again on the choice of the null parameters (α, h) . As such, a grid of tests can be considered to test every (α, h) combination. We can also slice the problem by defining

$$\begin{aligned} \min\{\alpha : P_{\alpha, h}(N_t < \hat{\eta}_t) \leq 0.05\}, \\ \min\{h : P_{\alpha, h}(N_t < \hat{\eta}_t) \geq 0.05\}, \end{aligned} \tag{7}$$

which represents the maximum size of a bubble (at a fixed horizon h) compatible with the data or the earliest burst compatible with the data for a bubble driven by a given quantile $q_\varepsilon(\alpha)$, respectively. Note that we fix the significance level at 5% in (7), but this can easily be adapted if desirable. The definition of the quantile we adopt makes larger significance levels more restrictive in the testing procedure (as will be shown in the empirical application).

4 Monte Carlo Study

In this section, we investigate the performance of the proposed level-growth testing procedure in terms of size and power using an extensive simulation study. We consider a data generating process (DGP) based on (1), which we denote by MAR(1,1), and is given by

$$(1 - \phi_1 L)(1 - \psi_1 L^{-1})y_t = \varepsilon_t,$$

with causal parameter $\phi_1 = 0.7$ and different specifications for the noncausal parameter ψ_1 and error distribution. More specifically, we consider a low ($\psi_1 = 0.5$) and high ($\psi_1 = 0.8$) persistence scenario as they influence the duration of the bubble trajectory. For the error term, we study both the standard Cauchy and Student's t -distribution, as they represent situations in which the distribution of our test statistic is analytically known and requires approximation, respectively. We choose a sample size of $T = 1000$ and run $M = 10,000$ iterations.^{3,4} In addition to the persistency of the noncausal process and the chosen error distribution, the choice of size and occurrence of the shock can influence the power and size of the test. For this reason, we investigate shocks drawn from various upper quantiles $q_\varepsilon(\alpha)$ and different horizons h .

Table 1 displays the size and power of the test for a MAR(1, 1) with $\psi_1 = 0.8$ and standard Cauchy distributed errors for $\alpha \in \{0.01, 0.005, 0.001, 0.0005, 0.0003, 0.0001\}$ and $h \in \{5, 10, 15, 20\}$. The reported results show that the size is very close to the nominal level of 5% for all different combinations of shock size and horizon. At a fixed shock size, we observe a decrease in power as h grows larger, which is in line with our expectations. It highlights that it becomes increasingly difficult to test for a bubble relatively far in the future, as we only have very limited information at our disposal. In contrast, if we have been on the bubble trajectory for multiple consecutive time points, this information is directly exploited by the test. For a fixed horizon h , we also observe that a more extreme shock to the system boosts the power of the test. This can be explained intuitively: whenever the shock is very large, the trajectory deviates a lot from the baseline path and thus exhibits a clear bubble pattern in the data. The large increase in the level and the fast growth rate combined provide valuable information to the test.

³The test relies on the parameter estimates. As a smaller sample size typically means a large sampling variance, the test's performance could be adversely affected. Therefore, we run our simulation study also with a smaller sample size of $T = 300$ and find qualitatively similar results. Results are available upon request.

⁴For the Student's t -distribution innovation, we use $M = 1000$ due to the time consuming procedure to simulate the critical values under the null.

Quantile	Size	Power	Size	Power	Size	Power	Size	Power
	$h = 5$		$h = 10$		$h = 15$		$h = 20$	
$\alpha = 1\%$	0.052	0.402	0.052	0.081	0.054	0.055	0.049	0.055
$\alpha = 0.5\%$	0.054	0.835	0.051	0.154	0.048	0.057	0.050	0.055
$\alpha = 0.1\%$	0.049	0.984	0.052	0.924	0.051	0.446	0.053	0.070
$\alpha = 0.05\%$	0.051	0.990	0.052	0.973	0.051	0.857	0.051	0.168
$\alpha = 0.03\%$	0.049	0.996	0.051	0.984	0.052	0.937	0.049	0.518
$\alpha = 0.01\%$	0.050	0.999	0.052	0.993	0.052	0.983	0.047	0.935

Table 1: Size and power of the test for Cauchy distributed innovation and high persistence

In Table 2, the importance of the persistency in the autoregressive process is revealed. When we perform the same exercise for a process that has $\psi_1 = 0.5$ instead of $\psi_1 = 0.8$, the power is clearly affected. In fact, we present results for the same values of α , but restrict our attention to only $h \in \{2, 4, 6, 8\}$ as testing at longer horizons is not instructive. We find that the power of the test is overall comparable to the ones obtained at larger horizons in Table 1. In the first two columns, we observe a similar increase in power over the different quantiles. However, for the cases $h \in \{6, 8\}$, it generally takes a much larger shock to obtain satisfactory power. This result does not come as a surprise: for the process to really manifest itself as a bubble in the low persistence case, we need a very large disturbance. Only in this way, the effect of the moderate noncausal parameter and its ability to generate a bubble can be identified in the period prior to the crash. This is natural as the information propagates dynamically through the system as ψ_1^h . At horizon $h = 5$, we can see that this quantity is approximately ten times larger for $\psi_1 = 0.8$ compared to $\psi_1 = 0.5$. It has to be mentioned though that empirical applications using MAR methodology typically identify highly persistent noncausal processes. Simulations also show that for a noncausal AR(1), an autoregressive parameter closer to unity is needed to mimic speculative bubbles observed in economic and financial time series. As expected, the size remains stable around 5%, as the test statistic's distribution is analytically known.

Quantile	Size	Power	Size	Power	Size	Power	Size	Power
	$h = 2$		$h = 4$		$h = 6$		$h = 8$	
$\alpha = 1\%$	0.053	0.359	0.053	0.053	0.053	0.044	0.050	0.036
$\alpha = 0.5\%$	0.051	0.912	0.053	0.068	0.050	0.047	0.053	0.039
$\alpha = 0.1\%$	0.055	0.992	0.054	0.939	0.052	0.104	0.050	0.039
$\alpha = 0.05\%$	0.051	0.996	0.054	0.971	0.054	0.650	0.052	0.055
$\alpha = 0.03\%$	0.051	0.998	0.053	0.988	0.054	0.916	0.051	0.082
$\alpha = 0.01\%$	0.053	0.999	0.053	0.997	0.054	0.982	0.054	0.826

Table 2: Size and power of the test for Cauchy distributed innovation and low persistence

To address a scenario in which we have to simulate the critical values of our test statistic's distribution, Table 3 considers the high persistence case where the error term follows a Student's t -distribution with $\nu = 2$ degrees of freedom. We observe a similar pattern of decrease in power for lower quantiles and higher horizons, albeit at a faster pace. This can be attributed to two factors. The approximation of the distribution, but also to the fact that the chosen t -distribution is less leptokurtic than the Cauchy distribution. This means that the same quantiles generally correspond to less pronounced bubbles being generated. Most notably, however, we can see the effect of simulation on the size of the test, which is further away from the nominal value of 5% than in the exact setting. It is important to note, however, that the number of simulations used to compute the critical values are much lower than before ($M = 1000$ instead of $M = 10,000$) due to the time consuming nature of this procedure and this clearly impacts the quality of the approximation.

Quantile	Size	Power	Size	Power	Size	Power
	$h = 5$		$h = 10$		$h = 15$	
$\alpha = 0.1\%$	0.049	0.712	0.047	0.223	0.049	0.059
$\alpha = 0.05\%$	0.047	0.920	0.056	0.230	0.056	0.066
$\alpha = 0.03\%$	0.047	0.990	0.051	0.309	0.054	0.092
$\alpha = 0.01\%$	0.045	0.99	0.049	0.698	0.052	0.171

Table 3: Size and power of the test for Student's t distributed innovation with $\nu = 2$ and high persistence

5 Empirical Study

In this section, we consider real-time testing for the presence of a bubble in the West Texas Intermediate (WTI) monthly oil price index. The data used is the same as in [Hecq and Voisin \(2022\)](#), with the only difference that we shorten the sample to our period of interest.⁵ More specifically, we study the data from March 1995 to October 2014 and focus on the sizeable bubble emerging around 2008. It is followed by a sharp decline in 2009, which is motivated by a strong revision of the short-to-medium-run demand for oil as a result of the unfolding financial crisis. Figure 2a displays the raw data, which appears non-stationary. We follow the convention in the MAR literature to not take the first difference or the growth rate of the series, because such transformations eliminate the bubble pattern of interest. Instead, we consider two types of detrending methods: an ad-hoc procedure using a time polynomial and joint modeling based on the MARST specification. Figure 2b shows the data after MARST detrending, which is the case that we mainly highlight for expository purposes. We observe that the WTI monthly price index now resembles a stationary process with an episode of locally explosive dynamics.




Figure 2: The sample considered in our application, for which our focus mainly lies on the bubble of 2008 and the period displaying spikes afterwards.

Similar to [Hecq and Voisin \(2022\)](#), we fit a $MAR(1, 1)$ to the data obtained after detrending, which yields

$$(1 - 0.28L)(1 - 0.92L^{-1})y_t = e_t,$$

where e_t denotes the residual. For the error term, we assume a non-standardized Student's t -


⁵The data is also freely available from the U.S. Energy Information Administration.

distribution $t(\nu, \sigma)$ and estimate the degrees of freedom parameter ν and scale parameter σ equal to 1.76 and 1.86 respectively. This distribution is often employed in the literature, as it also encompasses the standard Cauchy case (for ν and σ equal to 1). However, as indicated in the previous sections, this choice requires us to tabulate the critical values of the test statistic's distribution. Most importantly, the estimated noncausal parameter $\hat{\psi}_1 = 0.92$ reveals that we are in a high persistence scenario, in which our test performs well according to the simulation study. In this application, we illustrate how the level-growth test can be used for real-time testing, grid testing and risk assessment in the spirit of Value-at-Risk (VaR).


5.1 Real-Time Testing

As discussed in Section 3, our test depends on two dimensions, the horizon h for the bubble to reach its peak and the quantile of the shock driving the bubble $q_\varepsilon(\alpha)$ for a fixed significance level, which we set at 5% throughout this study as in (7). These equations also reveal the exact roles of horizon h and quantile $q_\varepsilon(\alpha)$ in the testing procedure. Let us first fix the horizon $h = 3$ to investigate how the choice of quantiles influences the test. Figure 3 displays the results of the level-growth test for two different quantiles based on $\alpha \in (0.07\%, 0.04\%)$. The shaded area represents the observations that are identified by the test as being compatible to a bubble reaching its extreme point $h = 3$ displacements in the future. In other words, the areas date-stamp the beginning and ending of a bubble according to our test procedure. This has implications in terms of risk as the period associated with the shock that drives the bubble is also the period where the crash is about to happen. By comparing the two panels in Figure 3, it is clear that a higher quantile corresponds to a test result that is more restrictive in highlighting observations compatible with the given shock (and subsequent crash). It is also worth noting that especially at high levels of persistence and with quantiles quite far in the tail, a bubble has a long span (in practice, around 20 time periods). This means that we are already quite far into the bubble when testing at a horizon $h = 3$. However this test still has relevant implications as it gives a sense of the size of the crash that is compatible with the data at a certain horizon.

Figure 4 shows the scenario in which we keep the quantile fixed at $q_\varepsilon = 0.9996$ but vary the horizon $h \in (3, 8)$. We find that the test is more conservative in identifying observations compatible with a bubble for a lower horizon h . This is to be expected, as it is generally more difficult to identify a bubble when we are still relatively far away from its peak. Interestingly, the




(a) $\alpha = 0.07\%$, $h = 3$




(b) $\alpha = 0.04\%$, $h = 3$

Figure 3: The observations compatible with a bubble originating from a shock corresponding to a two different quantiles three periods ahead.

test can be called more conservative in terms of risk assessment for larger h . It might namely identify observations that are congruous with a bubble while they do not truly instigate one. However, this might still be important information, as it provides an early-stage warning to monitor the development in the time series and the corresponding risk involved. Moreover, we compare the data-stamping of bubbles in the WTI oil price index of our test with those obtained by applying the backward ADF (BADF) test (Phillips et al., 2011) and the backward Supremum ADF (BSADF) test (Phillips et al., 2015). We can conclude that the identified bubbles roughly coincide with a quantile based on $\alpha = 0.04\%$ and horizon $h = 5$. The tests are thus able to generate similar results, but in the sequel we want to focus on how the two dimensions (α, h) in our procedure can be exploited to infer more about possible bubble emergence.



(a) $\alpha = 0.04\%$, $h = 3$



(b) $\alpha = 0.04\%$, $h = 8$

Figure 4: The observations compatible with a bubble originated by a shock corresponding to a fixed quantile at two different horizons.

5.2 Grid Testing

The analysis shows that extending the horizon allows for testing at the beginning of an explosive episode. The answer to the question “are we in a bubble?” is however increasingly difficult to answer at such an early stage. Therefore, we propose to exploit the two dimensions of the test simultaneously, as they suggest the existence of a two-dimensional frontier of significance dividing the explosive episodes that are compatible with the data observed today with those that are not. Because the frontier evolves over time, we illustrate its computation by explicitly considering three different points for the 2008 explosive episode. Figure 5 shows the data that we observe in the three considered scenarios: at the start of a bubble, early stage of the bubble and towards the end of the bubble. To calculate the frontier, we perform a grid of tests at these chosen dates.




Figure 5: Observed data for the different scenarios considered.

Table 4 displays the testing results for various combinations of quantiles and horizons, at the three scenarios of interest. The first panel reveals that at the start of the bubble, the null hypothesis is rejected for a large number of performed tests. This occurs due to the fact that we are operating in a portion of the sample without bubble or at best at an early stage of a bubble.⁶ In such a scenario, the observations are not compatible with a large explosive episode happening at short horizon. The table makes the stories of Figure 3 and 4 more explicit: we reject the presence of a bubble when the quantile is too high or the horizon is too low. When the horizon increases, the data start belonging to a period of bubble expansion and less rejections are observed. Interestingly, we have at $h = 10$ that no rejection takes place, not even for the highest quantile. The closer we get to the actual crash of the bubble, the more certain the test becomes and the data becomes compatible with extreme episodes happening at a relatively short horizon. This can be concluded from the second panel, which contains only three rejections for the highest

⁶We approach the problem as real-time testing, i.e., we pretend at this stage that we do not know whether a bubble manifests itself or not.

quantiles at relatively low horizons. Finally, in the third panel, at the final stage of the bubble, the test is vastly certain that an explosive episode is mimicking a bubble that manifests itself in the upcoming period. The only case that is rejected, is a shock from the largest quantile at the lowest horizon (i.e., an exceptionally extreme bubble taking place very soon). However, all other tests do not reject the null, and do this convincingly given the provided p -values.

Start of the bubble				
	$h = 3$	$h = 5$	$h = 7$	$h = 10$
$\alpha = 0.1\%$	0.056	0.256	0.505	0.733
$\alpha = 0.08\%$	0.030	0.174	0.437	0.695
$\alpha = 0.05\%$	0.005	0.066	0.270	0.500
$\alpha = 0.03\%$	0.000	0.012	0.112	0.478
$\alpha = 0.01\%$	0.000	0.000	0.003	0.114
Early stage of the bubble				
	$h = 3$	$h = 5$	$h = 7$	$h = 10$
$\alpha = 0.1\%$	0.434	0.734	0.859	0.900
$\alpha = 0.08\%$	0.306	0.673	0.828	0.896
$\alpha = 0.05\%$	0.086	0.446	0.741	0.879
$\alpha = 0.03\%$	0.008	0.178	0.570	0.839
$\alpha = 0.01\%$	0.000	0.001	0.052	0.576
Final stage of the bubble				
	$h = 3$	$h = 5$	$h = 7$	$h = 10$
$\alpha = 0.1\%$	0.879	0.919	0.937	0.951
$\alpha = 0.08\%$	0.864	0.909	0.929	0.949
$\alpha = 0.05\%$	0.788	0.882	0.921	0.943
$\alpha = 0.03\%$	0.572	0.842	0.898	0.935
$\alpha = 0.01\%$	0.001	0.255	0.743	0.901

Table 4: Test results (p -values) for compatibility of observations with a bubble for different sizes and horizons, tested at different stages of the 2008 bubble. Rejections of the test at a 5% significance level in bold.

5.3 Risk Assessment

The evaluation of various combinations of shock size and horizon suggests an evaluation of risk based on the test presented in (7). For each of the panels considered (i.e., being at different points in the bubble), it is possible to evaluate the maximum bubble compatible with the observed data at the different horizons. This is consistent with the intuition behind risk measures such as Value-at-Risk (VaR). Let us illustrate this by considering the scenario that we are at the start of the bubble. The first panel of Table 4 (first row) reveals that, if the innovation driving the bubble is located $h = 3$ periods ahead, it cannot be larger than an observation corresponding to the quantile $q = 0.9992$. As a consequence, we do not expect a crash of a larger size. If we repeat this exercise for all horizons, we have a confidence bound for the bubble-driving error term at each moment in time. Most importantly, this also provides a bound on the size of the crash, given that the explosive episode indeed ends on the specific date considered.




Figure 6: Confidence bounds for the maximum bubble compatible with the observed data at different horizons, in the spirit of VaR, computed at a significance level of 5% and 10%.

Figure 6 shows the corresponding frontiers that can be generated this way, for significance levels of 5% (orange line) and 10% (green line). We note that the frontier generally becomes closer to the real data trajectory when we are located closer to the peak of the bubble. The confidence bound is quite precise at short horizons and becomes wider as the horizon of the bubble is more distant. It is worth noting that there is no temporal information in this testing procedure about when the bubble actually bursts. However, for risk assessment, it provides an interesting way of doing scenario analysis. It provides an answer to the question “given that the bubble crashes at the preset data, what is the maximum crash size to expect?”. This can be extremely valuable when trying to monitor risk in real-time. It has to be noted that a bubble typically does not crash to zero, but to the baseline path before the emergence of explosive dynamics. If the MAR

specification adequately captures the data, the causal part plays a role in determining the level of the baseline path.

6 Conclusion

We have introduced a novel test for bubbles in time series that are modelled as mixed causal-noncausal autoregressive (MAR) models. The testing procedure is simple to implement and allows for the date-stamping of a bubble. A Monte Carlo study suggests that the test has acceptable power and size properties. In an empirical study of a monthly oil price index, it is shown that the test can produce important insights, allowing for Value-at-Risk calculations which entail a reasonable risk assessment *a-priori* in the presence of bubble dynamics in the data. In particular, the consideration of different combinations of shock size and horizon creates the possibility to perform scenario planning for risk.

The proposed test is a first attempt to leverage the MAR methodology to test for speculative bubbles. Future research could focus on testing for bubbles with different rates of increase. This might be established by considering MAR models with time-varying coefficients or aggregating noncausal processes (see e.g., [Gouriéroux and Zakoïan, 2017](#), Section 4). Moreover, an extension to the multivariate case allowing for possibly cointegrated bubbles can be explored.

References

- Andrews, B., R.A. Davis, and F. Jay Breidt (2006). Maximum likelihood estimation for all-pass time series models. *Journal of Multivariate Analysis* 97(7), 1638–1659.
- Blasques, F., S.J. Koopman, and G. Mingoli (2023). Observation-driven filters for time-series with stochastic trends and mixed causal-noncausal dynamics. *Tinbergen Institute Discussion Paper*.
- Breidt, F., R.A. Davis, K.S. Lii, and M. Rosenblatt (1991). Maximum likelihood estimation for noncausal autoregressive processes. *Journal of Multivariate Analysis* 36(2), 175–198.
- Brockwell, P. and R. Davis (1991). *Time Series: Theory and Methods*. Springer New York.

- Fries, S. and J.M. Zakoïan (2019). Mixed causal-noncausal AR processes and the modelling of explosive bubbles. *Econometric Theory* 35(6), 1234–1270.
- Gourieroux, C. and J. Jasiak (2016). Filtering, prediction and simulation methods for noncausal processes. *Journal of Time Series Analysis* 37(3), 405–430.
- Gouriéroux, C. and J.M. Zakoïan (2017). Local explosion modelling by non-causal process. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79(3), 737–756.
- Hecq, A., L. Lieb, and S. Telg (2016). Identification of mixed causal-noncausal models in finite samples. *Annals of Economics and Statistics* (123/124), 307–331.
- Hecq, A. and D. Velásquez-Gaviria (2022). Spectral estimation for mixed causal-noncausal autoregressive models.
- Hecq, A. and E. Voisin (2021). Forecasting bubbles with mixed causal-noncausal autoregressive models. *Econometrics and Statistics* 20, 29–45.
- Hecq, A. and E. Voisin (2022). Predicting crashes in oil prices during the COVID-19 pandemic with mixed causal-noncausal models. *arXiv working paper*.
- Hencic, A. and C. Gouriéroux (2015). Noncausal autoregressive model in application to Bitcoin/USD exchange rates. In *Econometrics of Risk*, pp. 17–40. New York: Springer.
- Lanne, M., J. Luoto, and P. Saikkonen (2012). Optimal forecasting of noncausal autoregressive time series. *International Journal of Forecasting* 28(3), 623–631.
- Lanne, M. and P. Saikkonen (2011). Noncausal autoregressions for economic time series. *Journal of Time Series Econometrics* 3(3), 1–32.
- Phillips, P.C., S. Shi, and J. Yu (2015). Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. *International Economic Review* 56(4), 1043–1078.
- Phillips, P.C., Y. Wu, and J. Yu (2011). Explosive behavior in the 1990s Nasdaq: When did exuberance escalate asset values? *International Economic Review* 52(1), 201–226.
- Phillips, P.C. and J. Yu (2011). Dating the timeline of financial bubbles during the subprime crisis. *Quantitative Economics* 2(3), 455–491.