Journal of Time Series Analysis

WILEY

NOTES AND COMMENTS

Laurent Series Expansion for MA(∞) Representation of Mixed Causal–Noncausal Autoregressive Processes

Gilles de Truchis[1] | Arthur Thomas[2]

[1]Université d'Orléans, Laboratoire d'Economie d'Orléans (LEO), Orléans, France | [2]Université Paris Dauphine, Université PSL, LEDA, CNRS, IRD, Paris, France

Correspondence: Arthur Thomas (arthur.thomas@dauphine.psl.eu)

Received: 29 September 2025 | **Revised:** 6 March 2026 | **Accepted:** 3 June 2026

Keywords: complex analysis | Laurent series | mixed causal–noncausal autoregression | moving average representation

ABSTRACT

This note develops a rigorous analytical framework for computing exact MA(∞) coefficients of mixed causal–noncausal autoregressive MAR(r, s) processes. While analytical solutions exist only for the MAR(1, 1) specification in the existing literature (Gouriéroux and Jasiak, 2016), general MAR(r, s) processes are typically handled through recursive approximation algorithms that suffer from numerical approximations. Using complex contour integration and the residue theorem, we derive explicit closed-form expressions valid for arbitrary orders (r, s). The derived expressions enable direct simulation algorithms, eliminating the recursive approximation bias while retaining only the standard truncation error inherent to any finite-order approximation, and facilitate implementation of forecasting methodologies. Numerical comparisons with existing recursive methods demonstrate improvements in accuracy, and an experiment with α-stable innovations illustrate the empirical relevance of our results.

MSC2020 Classification: 60G10, 60G52, 62M10

1 | Introduction

Mixed causal–noncausal autoregressive processes have gained significant attention in the econometrics and time series analysis literature. Historically, the increasing attention devoted to these processes has been motivated by their utility in modeling complex temporal patterns in time series data, particularly explosive bubble phenomena in financial asset prices (Andrews et al. 2009; Lanne and Saikkonen 2011; Gouriéroux and Jasiak 2016; Gouriéroux and Zakoian 2017; Velasco and Lobato 2018; Hecq et al. 2020; Gouriéroux and Jasiak 2023; Hecq and Velasquez Gaviria 2025; Blasques et al. 2025). These processes have been investigated as a methodological approach to enhance macroeconometric modeling (Gouriéroux, Monfort, and Renne 2020), as they provide a tractable reduced-form representation for rational expectations models with infinite variance innovations (Gouriéroux, Jasiak, and Monfort 2020). Furthermore, they have proven to be an effective analytical tool for modeling climate variables associated with extreme weather phenomena, including, for example, global temperature anomalies and ocean oscillation indices (Giancaterini et al. 2022; de Truchis et al. 2025). These processes, denoted MAR(r, s), are defined by the equation

$$\Phi(L)\Psi(L^{-1})y_t = \varepsilon_t, \tag{1}$$

where $\Phi(L) = 1 - \phi_1 L - \phi_2 L^2 - \cdots - \phi_r L^r$ is the causal polynomial of order r,$\Psi(L^{-1}) = 1 - \psi_1 L^{-1} - \psi_2 L^{-2} - \cdots - \psi_s L^{-s}$ is the noncausal polynomial of order s, both with roots outside of the unit circle, and ε_t is an independent and identically distributed (i.i.d.) heavy-tailed sequence. This representation is the

© 2026 John Wiley & Sons Ltd.

one that is mostly used in the univariate case for simulation, estimation, and prediction (Andrews et al. 2009; Lanne and Saikkonen 2011; Gouriéroux and Jasiak 2016; Hecq et al. 2020; Gouriéroux and Jasiak 2023). However, under some hypotheses detailed in Assumption 1, there is an equivalent to this multiplicative representation which is the two sided moving average infinity MA(∞) representation

$$y_t = \frac{1}{\Phi(L)\Psi(L^{-1})}\varepsilon_t = \sum_{k\in\mathbb{Z}} \delta_k \varepsilon_{t+k}, \tag{2}$$

where $(\delta_k)_{k\in\mathbb{Z}}$ is a deterministic coefficients sequence. This representation is also used for simulation and forecasting. Fries (2022) demonstrates that the $(\delta_k)_{k\in\mathbb{Z}}$ coefficients are essential for computing predictive conditional moments of mixed causal–noncausal ARMA processes. In an extension of this work, de Truchis et al. (2025) shows that during extreme events and under some conditions (typically, $s > 1$), conditional on observing a segment of the trajectory, the future path is collinear to the $(\delta_k)_{k\in\mathbb{Z}}$ coefficients and hence almost deterministic. However, few papers propose solutions to compute these coefficients. Gouriéroux and Jasiak (2016) suggest an approach based on partial fraction decomposition, as it allows theoretical derivation of an exact representation of the MA coefficients. However, the partial fraction decomposition can be tedious to derive for high-order (r, s) specifications, and the case $r = s = 1$ is essentially the only one considered in the literature. For large values of (r, s), existing approaches are typically based on recursive computation, similar to impulse response calculations. For example, this is the approach used in Lanne and Saikkonen (2011)'s seminal paper, extended in Hecq et al. (2020) and implemented in the so-called MARX package developed by Hecq et al. (2017).

Unfortunately, this recursive approach introduces systematic approximation bias in coefficient computation, as it relies on defining the filtered components through $u_t \equiv \Phi(L)y_t$ and $v_t \equiv \Psi(L^{-1})y_t$ or equivalently $\Psi(L^{-1})u_t = \varepsilon_t$ and $\Phi(L)v_t = \varepsilon_t$ and their associated MA(∞) representations

$$u_t = \sum_{k=0}^{\infty} \delta_k^{(u)} \varepsilon_{t+k} \text{ and } v_t = \sum_{k=0}^{\infty} \delta_k^{(v)} \varepsilon_{t-k},$$

where $(\delta_k^{(u)})_{k\in\mathbb{Z}}$ and $(\delta_k^{(v)})_{k\in\mathbb{Z}}$ are coefficients of $[\Psi(L^{-1})]^{-1}$ and $[\Phi(L)]^{-1}$. It is important to distinguish two sources of error in this recursive approach: (ı) the truncation bias, which is unavoidable and arises from computing only a finite number of coefficients $k \in [-m, m]$ instead of the full sequence $k \in \mathbb{Z}$, and (ıı) the recursive approximation bias which arises from the nested convolution structure where truncation errors propagate and accumulate through successive filtering operations. While truncation bias affects any finite-order approximation method, the recursive approximation bias is specific to this two-step filtering approach and can be eliminated through direct analytical computation of the coefficients. This note extends the MAR(1,1) analytical result by deriving explicit closed-form expressions for the MA(∞) coefficients of any MAR(r, s) processes using complex analytic techniques. The remainder of this note is organized as follows. Section 2 presents the main lemma and associated proof. Section 3 applies this framework to the MAR$(1, 1)$ and the MAR$(1, 2)$ with complex conjugate noncausal roots, with comparisons to recursive approximation methods, and includes a sensitivity analysis quantifying the recursive approximation error as the noncausal root modulus approaches the unit circle. Section 4 presents direct simulation algorithms, based on analytically computed coefficients and shows how these coefficients can be used in a pattern-based forecasting methodology for processes with α-stable innovations. Section 5 concludes.

2 | Theoretical Results

For the analysis of the two sided MA(∞) representation, it is convenient to express the autoregressive polynomials in their factorized forms in terms of their roots. The causal polynomial can be written as

$$\Phi(z) = 1 - \sum_{k=1}^{r} \phi_k z^k = \prod_{i=1}^{r} (1 - \lambda_i z),$$

and the noncausal polynomial as

$$\Psi(z^{-1}) = 1 - \sum_{k=1}^{s} \psi_k z^{-k} = \prod_{j=1}^{s} (1 - \zeta_j z^{-1}) = z^{-s} \prod_{j=1}^{s} (z - \zeta_j),$$

where λ_i and ζ_j represent the complex coefficients after factorization of the polynomials $\Phi(z)$ and $\Psi(z^{-1})$ respectively. In view of application of the residue theorem of Ahlfors (1979), we introduce the following assumption.

Assumption 1. *The process defined in Equation (2) is such that*:

(ı) *the innovations sequence ε_t is i.i.d with $\mathbb{E}(\varepsilon_t) = 0$.*

(ıı) *ε_t is regularly varying with tail index $\alpha > 0$ such that $\mathbb{E}(|\varepsilon_1|) < \infty$ if $\alpha > 1$ and $\mathbb{E}(|\varepsilon_1|^\gamma) < \infty$ for some $\gamma < \alpha$ if $\alpha \leq 1$*

(ıı) *the coefficients sequence $(\delta_k)_{k\in\mathbb{Z}}$ is real-valued and satisfies*

$$\begin{cases} \sum_{k\in\mathbb{Z}} \delta_k^2 < \infty, & \text{if } \alpha > 2, \\ \sum_{k\in\mathbb{Z}} |\delta_k|^\beta < \infty, \quad \beta < \alpha, & \text{if } 0 < \alpha \leq 2, \end{cases} \tag{3}$$

such that the two sided MA(∞) representation of (1) converges a.s. for all $t \in \mathbb{Z}$ (stationarity condition).

Also, under (ıı), for all $i \in \{1, \ldots, r\}$ and $j \in \{1, \ldots, s\}$, we assume:

(ıv) *the polynomials $\Phi(z)$ and $\Psi(z^{-1})$ have distinct roots, such that $\lambda_i \neq \lambda_l$ for $i \neq l$ and $\zeta_j \neq \zeta_m$ for $j \neq m$. Furthermore, let $\mathcal{R}_1 = \max_{j=1,\ldots,s} |\zeta_j| < 1$, $\mathcal{R}_2 = \min_{i=1,\ldots,r} |\lambda_i|^{-1} > 1$ with $\mathcal{R}_1 < \mathcal{R}_2$, and define the annulus $\mathcal{A} = \{z \in \mathbb{C} : \mathcal{R}_1 < |z| < \mathcal{R}_2\}$.*

Assumption 1 provides sufficient conditions for the existence and convergence of the MA(∞) representation (2). The tail conditions in (ıı) align with the non-causal literature requiring non-Gaussian errors for identification of anticipative roots but are unnecessary to apply the residue theorem. The case $\alpha > 2$ is stated for seek of generality. However, as this literature essentially focus on extreme events (e.g., α-stable innovations), the rest of the paper discusses the case $\alpha < 2$. By contrast, (ıv) ensures convergence of the Laurent series in the annulus $\mathcal{A}$, which is essential for the application of the residue theorem and the demonstration of the following Lemma.

Lemma 2.1. *For a MAR(r, s) process defined by Equation (1) with factorized polynomials* $\Phi(z) = \prod_{i=1}^{r}(1-\lambda_i z)$ *and* $\Psi(z^{-1}) = \prod_{j=1}^{s}(1-\zeta_j z^{-1})$ *satisfying Assumption 1, the MA(∞) coefficients of Equation (2) are given by*:

(ı) *For* $k > 0$ *(noncausal coefficients)*:

$$\delta_k = \sum_{j=1}^{s} \frac{\zeta_j^{(s-1)+k}}{\prod_{m\neq j}^{s}(\zeta_j-\zeta_m)\cdot\prod_{i=1}^{r}(\lambda_i\zeta_j-1)}\cdot(-1)^r. \tag{4}$$

(ıı) *For* $k \leq 0$ *(causal coefficients)*:

$$\delta_k = \sum_{i=1}^{r} \frac{\lambda_i^{(r-1)+|k|}}{\prod_{l\neq i}^{r}(\lambda_i-\lambda_l)\cdot\prod_{j=1}^{s}(\lambda_i\zeta_j-1)}\cdot(-1)^s. \tag{5}$$

Proof of Lemma 2.1. The proof uses Laurent series expansion combined with contour integration techniques from complex analysis. The transfer function of the MAR(r, s) process can be written as

$$H(z) = \frac{1}{\Phi(z)\Psi(z^{-1})} = \frac{z^s}{\prod_{i=1}^{r}(1-\lambda_i z)\cdot\prod_{j=1}^{s}(z-\zeta_j)}.$$

The MA(∞) representation seeks coefficients δ_k such that

$$y_t = \sum_{k=-\infty}^{\infty}\delta_k\varepsilon_{t+k} = H(z)\varepsilon_t,$$

where $H(z) = \sum_{k=-\infty}^{\infty}\delta_k z^{-k}$ is the Laurent series expansion. According to Laurent series theory, the coefficients are given by

$$\delta_k = \frac{1}{2\pi i}\oint_C H(z)z^{k-1}dz,$$

where C is an appropriate contour lying in $\mathcal{A} = \{z\in\mathbb{C} : \mathcal{R}_1 < |z| < \mathcal{R}_2\}$. Given our mixed causal–noncausal framework, the Laurent series expansion requires careful consideration of the regions in which each part of the series converges. The function $H(z)$ has two types of poles: poles at $z = 1/\lambda_i$ for $i = 1,\ldots,r$ (all outside the unit circle since $|\lambda_i| < 1$) and poles at $z = \zeta_j$ for $j = 1,\ldots,s$ (all inside the unit circle since $|\zeta_j| < 1$). To compute the coefficients δ_k, we choose contours within the convergence annulus $\mathcal{A}$ depending on the value of k.

(ı) For $k > 0$ (non-causal coefficients), the integrand grows as $|z|\to\infty$, so we use a contour C_+ in the region $\{z\in\mathcal{A} : |z| > \mathcal{R}_1\}$ oriented counterclockwise and enclosing all poles ζ_j. In that case, we denote by $f_+(z)$ the integrand. By the residue theorem (Ahlfors 1979) and under Assumption 1 we have

$$\delta_k = \frac{1}{2\pi i}\oint_{C_+} f_+(z)dz = \sum_{j=1}^{s}\mathrm{Res}(f_+,\zeta_j),$$

where $\mathrm{Res}(f_+,\zeta_j)$ denotes the residue of f_+ at ζ_j. As we assume in Assumption1-(ıv) that the poles are simple, we compute them as

$$\begin{aligned}\mathrm{Res}(f_+,\zeta_j) &= \lim_{z\to\zeta_j}(z-\zeta_j)\cdot\frac{z^{s+k-1}}{\prod_{i=1}^{r}(1-\lambda_i z)\cdot\prod_{m=1}^{s}(z-\zeta_m)}\\ &= \frac{\zeta_j^{s+k-1}}{\prod_{i=1}^{r}(1-\lambda_i\zeta_j)\cdot\prod_{m\neq j}(\zeta_j-\zeta_m)}.\end{aligned}$$

Note that the denominator $\prod_{i=1}^{r}(1-\lambda_i\zeta_j)$ can be rewritten as $(-1)^r\prod_{i=1}^{r}(\lambda_i\zeta_j-1)$, introducing a factor $(-1)^r$. Equation (4) follows. (ıı) Now consider the case $k\leq 0$ (causal coefficients). Since the integrand decays as $|z|\to\infty$ for $k\leq 0$, the integral over a large circle vanishes. Using a contour C_- in the region $\{z\in\mathcal{A} : |z| < \mathcal{R}_2\}$, oriented clockwise when viewed from infinity, the residue theorem yields

$$\delta_k = \frac{1}{2\pi i}\oint_{C_-} f_-(z)dz = -\sum_{i=1}^{r}\mathrm{Res}(f_-,1/\lambda_i), \tag{6}$$

where $f_-(z)$ denotes the integrand in that causal case. To compute the residue at each pole $1/\lambda_i$, we note that $1-\lambda_i z = -\lambda_i(z-1/\lambda_i)$ and thus,

$$\begin{aligned}\mathrm{Res}(f_-,1/\lambda_i) &= \lim_{z\to 1/\lambda_i}(z-1/\lambda_i)\cdot\frac{z^{s+k-1}}{\prod_{l=1}^{r}(1-\lambda_l z)\cdot\prod_{j=1}^{s}(z-\zeta_j)}\\ &= \lim_{z\to 1/\lambda_i}\frac{z^{s+k-1}}{(-\lambda_i)\prod_{l\neq i}(1-\lambda_l z)\cdot\prod_{j=1}^{s}(z-\zeta_j)}.\end{aligned}$$

Substituting $z = 1/\lambda_i$,

$$\mathrm{Res}(f_-,1/\lambda_i) = \frac{(1/\lambda_i)^{s+k-1}}{(-\lambda_i)\prod_{l\neq i}(1-\lambda_l/\lambda_i)\cdot\prod_{j=1}^{s}(1/\lambda_i-\zeta_j)},$$

and simplifying the denominators as

$$\begin{aligned}\prod_{l\neq i}(1-\lambda_l/\lambda_i) &= \prod_{l\neq i}\frac{\lambda_i-\lambda_l}{\lambda_i} = \frac{\prod_{l\neq i}(\lambda_i-\lambda_l)}{\lambda_i^{r-1}},\\ \prod_{j=1}^{s}(1/\lambda_i-\zeta_j) &= \prod_{j=1}^{s}\frac{1-\lambda_i\zeta_j}{\lambda_i} = \frac{\prod_{j=1}^{s}(1-\lambda_i\zeta_j)}{\lambda_i^s},\end{aligned}$$

we obtain

$$\begin{aligned}\mathrm{Res}(f_-,1/\lambda_i) &= \frac{\lambda_i^{-(s+k-1)}}{(-\lambda_i)\cdot\frac{\prod_{l\neq i}(\lambda_i-\lambda_l)}{\lambda_i^{r-1}}\cdot\frac{\prod_{j=1}^{s}(1-\lambda_i\zeta_j)}{\lambda_i^s}}\\ &= \frac{\lambda_i^{r-k-1}}{\prod_{l\neq i}(\lambda_i-\lambda_l)\cdot\prod_{j=1}^{s}(1-\lambda_i\zeta_j)}\cdot(-1).\end{aligned}$$

We can rewrite the denominator $\prod_{j=1}^{s}(1-\lambda_i\zeta_j)$ as $(-1)^s\prod_{j=1}^{s}(\lambda_i\zeta_j-1)$, which gives

$$\mathrm{Res}(f_-,1/\lambda_i) = \frac{(-1)^{1-s}\lambda_i^{r-1-k}}{\prod_{l\neq i}(\lambda_i-\lambda_l)\cdot\prod_{j=1}^{s}(\lambda_i\zeta_j-1)}.$$

Since $k\leq 0$, we have $r-1-k = (r-1)+|k|$, yielding Equation (5) once we accounted for the negative sign in Equation (6). This concludes the proof. □

3 | Numerical Illustrations

We illustrate the formulas of Lemma 2.1 on two model specifications of increasing complexity: a MAR(1, 1) with a single real noncausal root and a MAR(1, 2) with complex conjugate noncausal roots. For each specification we compute the exact MA(∞) coefficients and compare them with those obtained by the standard recursive two-step filtering method used, for example, in

TABLE 1 | Recursive approximation error for the MAR(1, 2) specification with complex conjugate noncausal roots ($m = 20$). The RMSE is defined in (7).

Length factor ℓ	Total length $L = \ell \times m$	RMSE (δ^e, δ^r)	max $\lvert\delta^e - \delta^r\rvert$
4	80	1.78×10^{-4}	8.27×10^{-4}
6	120	2.14×10^{-7}	9.76×10^{-7}
8	160	2.46×10^{-10}	1.19×10^{-9}
10	200	2.42×10^{-13}	1.32×10^{-12}

Lanne and Saikkonen (2011) and Hecq et al. (2020). As we shall see, the two methods are numerically equivalent in the low-order MAR(1, 1) case, but the recursive approximation bias becomes significant for the MAR(1, 2) specification with complex conjugate roots (Table 1). A numerical experiment in Section 3.3 then reveals that this bias can grow by several orders of magnitude as the noncausal root modulus approaches the unit circle (Figures 1 and 2).

The recursive approach computes the MA(∞) coefficients by passing a unit impulse successively through the truncated noncausal filter $[\Psi(L^{-1})]^{-1}$ (backward pass) and the truncated causal filter $[\Phi(L)]^{-1}$ (forward pass), over a finite window of total length $L = \ell \times m$, where m is the number of desired MA coefficients on each side and ℓ denotes the length factor. Higher values of ℓ reduce boundary effects but increase computational cost. In all the experiments below, we set $m = 20$ and vary the length factor $\ell \in \{4, 6, 8, 10\}$. For every pair of coefficient vectors $\boldsymbol{\delta}^e = (\delta^e_{-m}, \ldots, \delta^e_m)$ and $\boldsymbol{\delta}^r = (\delta^r_{-m}, \ldots, \delta^r_m)$, denoting respectively the exact and recursive MA(∞) coefficients, the approximation quality is measured by the root mean squared error

$$\mathrm{RMSE}\left(\boldsymbol{\delta}^e, \boldsymbol{\delta}^r\right) = \sqrt{\frac{1}{2m}\sum_{k=-m}^{m-1}\left(\delta^e_k - \delta^r_k\right)^2}. \tag{7}$$

3.1 | MAR(1, 1) Case

We consider the simple MAR(1, 1) process $(1 - 0.6L)(1 - 0.7L^{-1})y_t = \varepsilon_t$, with $\lambda_1 = 0.6$ and $\zeta_1 = 0.7$, which corresponds to the known case from Gouriéroux and Jasiak (2016). Applying Lemma 2.1, we get

$$\delta_k = \frac{\zeta_1^k}{1 - \lambda_1\zeta_1} = \frac{(0.7)^k}{0.58}, \quad k > 0,$$
$$\delta_0 = \frac{1}{1 - \lambda_1\zeta_1} = \frac{1}{0.58} \approx 1.724,$$
$$\delta_k = \frac{\lambda_1^{|k|}}{1 - \lambda_1\zeta_1} = \frac{(0.6)^{|k|}}{0.58}, \quad k < 0.$$

The coefficients display the characteristic asymmetric geometric decay of MAR(1, 1) processes: past influences ($k < 0$) decay at rate $\lambda_1 = 0.6$, while future influences ($k > 0$) decay more slowly at rate $\zeta_1 = 0.7$. In this particular low-order specification, the impulse response of each one-sided filter decays rapidly enough for the truncation window $L = \ell \times m$ to capture it fully, even at $\ell = 4$. As a result, the recursive method achieves highly accurate results and the two approaches are numerically equivalent.

3.2 | MAR(1,2) Case With Complex Conjugate Noncausal Roots

We consider a MAR(1, 2) process $(1 - 0.6L)(1 - 0.8L^{-1} + 0.5L^{-2})y_t = \varepsilon_t$, with $\lambda_1 = 0.6$ and complex conjugate noncausal roots $\zeta_1 = 0.4 + 0.583i$ and $\zeta_2 = 0.4 - 0.583i$ (modulus $|\zeta_{1,2}| = \sqrt{0.5} \approx 0.707$). From Lemma 2.1, the backward coefficients simplify to

$$\delta_k = \frac{(0.6)^{|k|}}{(\lambda_1\zeta_1 - 1)(\lambda_1\zeta_2 - 1)} = \frac{(0.6)^{|k|}}{0.7}, \quad k \le 0,$$

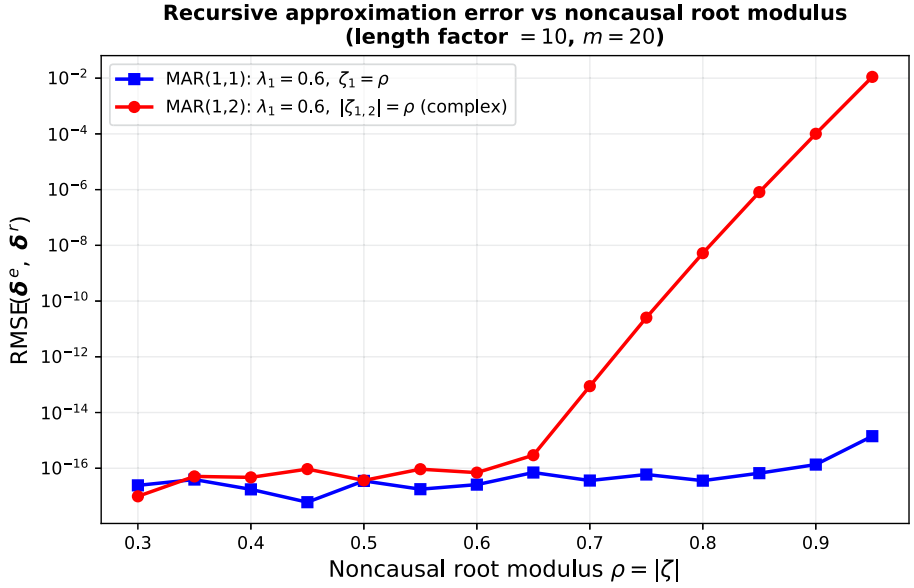


FIGURE 1 | RMSE of our exact coefficients and the recursive approximation as a function of the noncausal root modulus $|\zeta|$, for two MAR specifications (length factor $\ell = 10$, $m = 20$).

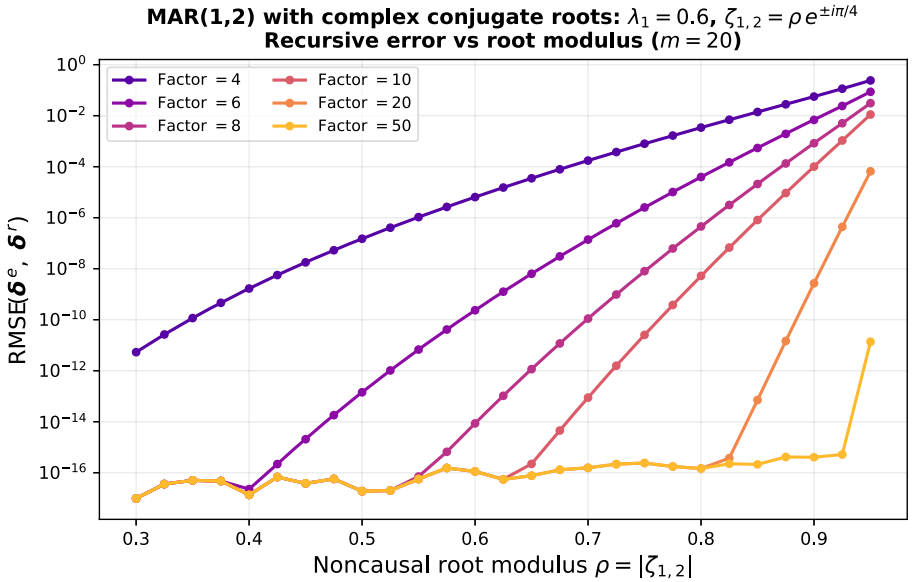


FIGURE 2 | RMSE of exact and recursive MA(∞) coefficients for the MAR(1,2) specification with complex conjugate noncausal roots $\zeta_{1,2} = \rho\, e^{\pm i\pi/4}$ and $\lambda_1 = 0.6$, as a function of $\rho = |\zeta_{1,2}|$ for different length factors $\ell \in \{4,\ 6,\ 8,\ 10,\ 20,\ 50\}$ ($m = 20$).

while the forward coefficients ($k > 0$) are given by the partial fraction sum

$$\delta_k = \sum_{j=1}^{2} \frac{\zeta_j^{1+k}}{(\zeta_j - \zeta_{3-j})(1 - \lambda_1 \zeta_j)},$$

which yields real-valued oscillatory decay with envelope rate $\sqrt{0.5} \approx 0.707$.

Unlike the MAR(1, 1) case, the recursive method now exhibits a non-negligible approximation bias. Table 1 reports the RMSE and maximum absolute error for different length factors. At $\ell = 4$ (i.e., $L = 80$), the RMSE is of order 10^{-4}, which may be significant in applications involving iterated filtering or forecasting over long horizons. At $\ell = 10$ (i.e., $L = 200$), the RMSE remains at the order of 10^{-13}, whereas our analytical formulas yield practically exact results ($\approx 10^{-16}$).[1]

3.3 | Sensitivity Analysis to Root Modulus

The previous examples use fixed parameter values with moderate noncausal root moduli ($|\zeta_1| = 0.7$ for MAR(1, 1) and $|\zeta_{1,2}| \approx 0.707$ for MAR(1, 2)). In practice, however, estimated noncausal roots can be much closer to the unit circle, especially in financial time series exhibiting pronounced bubble dynamics. To quantify the advantage of our closed-form expressions in such close to the unit-root configurations, we conduct a sensitivity analysis in which the noncausal root modulus $\rho = |\zeta|$ varies from 0.3 to 0.95 while we keep the length factor fixed at $\ell = 10$.

Figure 1 reports the RMSE defined in (7) as a function of ρ for two model specifications:

(ı) MAR(1, 1): $\lambda_1 = 0.6$ and $\zeta_1 = \rho$ (single real noncausal root),

(ıı) MAR(1, 2): $\lambda_1 = 0.6$ and complex conjugate noncausal roots $\zeta_{1,2} = \rho\, e^{\pm i\pi/4}$.

In both cases, the recursive approximation error increases sharply as ρ approaches 1, spanning more than ten orders of magnitude over the range considered. This behavior is expected: as the noncausal roots approach the unit circle, the impulse response of $[\Psi(L^{-1})]^{-1}$ decays ever more slowly, and the finite truncation window $L = \ell \times m$ becomes increasingly insufficient to capture the tail of the response. By contrast, our analytical expressions remain exact regardless of ρ, since the residue computation does not involve any truncation of the impulse response.

A natural question is whether increasing the length factor ℓ can compensate for the loss of accuracy near the unit circle. Figure 2 addresses this question for the MAR(1, 2) specification with complex conjugate roots $\zeta_{1,2} = \rho\, e^{\pm i\pi/4}$ and $\lambda_1 = 0.6$. Even with $\ell = 50$ (i.e., $L = 1000$ for $m = 20$), the recursive RMSE at $\rho = 0.95$ remains at the order of 10^{-3}, while our closed-form expressions yield highly precise results. Moreover, the computational cost of the recursive method scales linearly in L, so that large length factors entail a substantial overhead. These results demonstrate that the analytical approach of Lemma 2.1 is not merely a theoretical convenience but provides a concrete numerical advantage, especially when roots are close to the unit circle or when high-precision coefficients are required for downstream tasks such as forecasting.

4 | Simulation and Forecasting Algorithms

Based on the closed-form expressions derived in Lemma 2.1, we present a direct simulation algorithm for MAR(r, s) processes. This method explicitly computes the MA(∞) coefficients to generate process realizations with controlled truncation error. Accordingly, the algorithm proceeds through the following steps:

(ı) *Coefficient computation*: Using Equations (4) and (5), we compute the MA(∞) coefficients δ_k for a sufficiently large

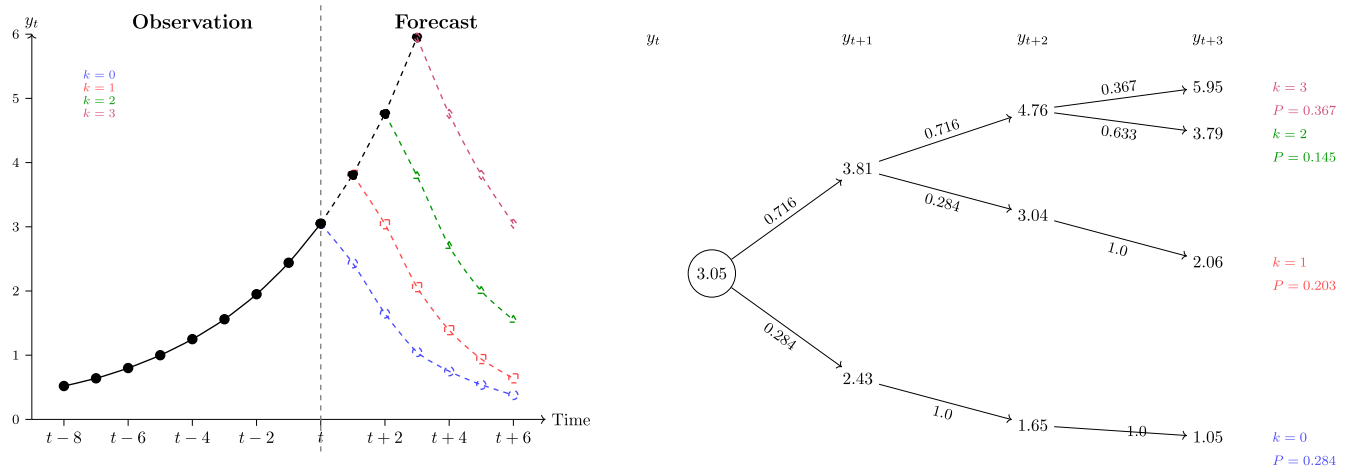


FIGURE 3 | Pattern-based forecasting illustration for MAR(3, 1) process with exponential growth observation and probabilistic forecast trajectories. Node values computed as $y_{t+h} = y_t \times \mathbf{D}_k^{(h)}$ where $\mathbf{D}_k$ are normalized patterns from (9). Probabilities follow $P(k) = \zeta_1^{\alpha k}(1 - \zeta_1^{\alpha})$ for $k \in \{0, 1, 2\}$ and $P(k = 3) = \zeta_1^{\alpha h}$ as in (10).

range $k \in [-m, m]$, where m is chosen such that $|\delta_k| < \epsilon$ for some predetermined tolerance $\epsilon > 0$ and $|k| > m$.

(*ii*) *Heavy-tailed innovation generation*: Generate i.i.d. innovations ε_t following an α-stable distribution $S(\alpha, \beta, \sigma, \mu)$. The parameters control tail behavior ($\alpha \in (0, 2]$), asymmetry ($\beta \in [-1, 1]$), scale ($\sigma > 0$), and location ($\mu \in \mathbb{R}$). Other heavy-tailed distributions may be substituted as needed (Student t distribution)

(*iii*) *Process simulation*: For each time point $t \in \{m + 1, \ldots, T - m\}$, compute the exact process values:

$$y_t = \sum_{k=-m}^{m} \delta_k \varepsilon_{t+k}. \tag{8}$$

The explicit computation of the MA(∞) coefficients δ_k provides a theoretical foundation for advanced forecasting methods based on pattern recognition during extreme events (see e.g., Fries 2022; de Truchis et al. 2025). As an illustration, consider a MAR(3, 1) process defined as the strictly stationary solution of

$$(1 - 0.6L)(1 - 0.3L)(1 - 0.1L)(1 - 0.8L^{-1})y_t = \varepsilon_t,$$

where $\zeta_1 = 0.8$ and $\lambda_1 = 0.6$, $\lambda_2 = 0.3$, $\lambda_3 = 0.1$. Here, (ε_t) is an i.i.d. sequence of regularly varying errors with tail index $\alpha = 1.5$. The process (y_t) admits the MA(∞) representation $y_t = \sum_{k\in\mathbb{Z}} \delta_k \varepsilon_{t+k}$ with coefficients computed via Lemma 2.1. Following Fries (2022), denote $\mathbf{y}_t := (y_{t-m}, \ldots, y_t)$, $\mathbf{y}_{t+h} := (y_{t+1}, \ldots, y_{t+h})$, where $m \geq 1$, and $h \geq 1$ the forecasting horizon. We define the forecasting patterns,

$$\mathbf{D}_k := \frac{(\delta_{k-1}, \delta_{k-2}, \ldots, \delta_{k-h})}{|\delta_k|} \quad \text{for } k \in \{0, 1, \ldots, h\}.$$

The index $k \in \{0, 1, \ldots, h\}$ represents the number of periods until the bubble peak: $k = 0$ corresponds to the peak occurring at time t (immediate collapse), while $k = h$ means the peak has not yet occurred within the forecasting horizon. Applying Lemma 2.1 yields $\delta_0 = 2.750$, $\delta_1 = 2.200$, $\delta_2 = 1.760$, $\delta_3 = 1.408$ for $k \geq 0$, and $\delta_{-1} = 2.188$, $\delta_{-2} = 1.485$, $\delta_{-3} = 0.944$ for $k < 0$. For a forecasting horizon $h = 3$, the normalized patterns are:

$$\begin{aligned}
\mathbf{D}_0 &= (0.796, 0.540, 0.343), \\
\mathbf{D}_1 &= (1.25, 0.995, 0.675), \\
\mathbf{D}_2 &= (1.25, 1.562, 1.243), \\
\mathbf{D}_3 &= (1.25, 1.562, 1.953).
\end{aligned} \tag{9}$$

Under appropriate regularity conditions, if during an extreme event the trajectory $\mathbf{y}_t$ is approximately collinear to $(\zeta_1^{-m}, \zeta_1^{-(m-1)}, \ldots, \zeta_1^{-1}, 1)$ with exponential growth rate $\zeta_1^{-1} = 1.25$, then the forecasting probabilities are:

$$\begin{aligned}
&\mathbb{P}\left(\left\|\frac{\mathbf{y}_{t+h}}{|\mathbf{y}_t|} - \vartheta \boldsymbol{D}_k\right\| < \nu \middle| |y_t| > y, \left\|\frac{\mathbf{y}_t}{|\mathbf{y}_t|} - \vartheta\boldsymbol{\xi}\right\| < \eta\right) \\
&\qquad \xrightarrow[y\to+\infty]{} \begin{cases} \zeta_1^{\alpha k}(1 - \zeta_1^{\alpha}) & \text{if } k \in \{0, 1, \ldots, h-1\} \\ \zeta_1^{\alpha h} & \text{if } k = h \end{cases},
\end{aligned} \tag{10}$$

where $\boldsymbol{\zeta} := (\zeta_1^{-m}, \zeta_1^{-(m-1)}, \ldots, \zeta_1^{-1}, 1)$, $\vartheta \in \{-1, +1\}$ indicates the sign of the extreme event, and $\nu, \eta > 0$ are small tolerance parameters controlling the proximity to the theoretical patterns). Given an observed trajectory (y_{t-2}, y_{t-1}, y_t) approximately collinear to $\boldsymbol{\zeta} = (0.64, 0.8, 1)$ and tail parameter $\alpha = 1.5$, the conditional forecasting probabilities converge to approximately 0.284, 0.203, 0.145, and 0.367 for $k = 0, 1, 2, 3$, respectively. This conditioning set differs from the standard conditional extreme value (CEV) framework as developed in Gourieroux et al. (2025), who condition solely on $|y_t| > y$. In their approach, the limiting distribution of the tail process (y_{t+h}) depends on a discrete random variable N representing the stochastic time lag to the "single big jump" (SBJ), that is, the location of the dominating innovation. Our conditioning in Equation (10) combines two facts: (i) that $|y_t|$ exceeds a large threshold (standard CEV), and (ii) that, as a consequence of conditioning on a large event when the process has at least one anticipative component, the normalized past trajectory $\mathbf{y}_t/|y_t|$ is approximately collinear to the geometric pattern $\boldsymbol{\zeta}$, which characterizes the exponential growth phase of a bubble (Fries 2022; de Truchis et al. 2025). This trajectory-based

conditioning effectively restricts attention to scenarios where the bubble is still inflating, providing richer initial information that potentially avoids the need for sequential updating as in Gourieroux et al. (2025). The limiting probabilities in (10) admit a natural interpretation: k represents the number of periods until the bubble peak, and the geometric distribution $\zeta_1^{\alpha k}(1-\zeta_1^{\alpha})$ emerges as the conditional survival probability of the bubble. This result follows from Fries (2022, Proposition 4.2), which establishes that for processes satisfying $\delta_k = C \cdot \zeta_1^k$ for $k \geq 0$, where C is the normalization constant that depends on the causal and non-causal roots (as verified above for our MAR(3, 1)), the crash odds are geometric with parameter ζ_1^{α}, regardless of the causal component structure. This memory-less property implies that the forecaster only needs to estimate ζ_1 and α, and can remain agnostic about the post-peak collapse dynamics governed by the causal roots $\lambda_1, \lambda_2, \lambda_3$.

Figure 3 demonstrates the implementation of the theoretical framework where the left panel shows the transition from observed exponential growth to multiple forecasting scenarios, while the right panel presents the probability tree structure starting from $y_t = 3.05$ with computed conditional probabilities for different bubble peak timings ($k = 0, 1, 2, 3$). The visualization illustrates how the pattern-matching algorithm leverages the explicit MA(∞) coefficients to assign forecast probabilities based on the observed trajectory's alignment with the theoretical patterns $\mathbf{D}_k$, demonstrating the practical utility of the coefficient formulas from Lemma 2.1 for systematic bubble forecasting through pattern recognition. During extreme events, the process dynamics simplify to geometric patterns governed by the noncausal root ζ_1, enabling tractable probabilistic forecasting of bubble peaks where each branch represents a distinct collapse scenario with its associated probability.

5 | Conclusion

This note derives closed-form expressions for the MA(∞) coefficients of mixed causal–noncausal autoregressive MAR(r, s) processes through Laurent series expansion and residue calculus. Our results extend the partial fraction decomposition approach from the MAR(1, 1) case to arbitrary MAR(r, s) specifications. The derived expressions enable direct simulation algorithms that avoid the approximation bias inherent in recursive methods, as demonstrated through numerical comparisons across different model specifications. Furthermore, the explicit coefficient structure facilitates the development of pattern-based forecasting methodologies for extreme events in processes with α-stable innovations.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The figures in this article are generated using Python code; no real data is used, and all figures along with the Python code implementing the paper's formulas are available in the replication code package. Data sharing not applicable to this article as no datasets were generated or analyzed during the current study.

Endnote

[1]This residual difference simply reflects standard numerical rounding errors, indicating that the analytical formula is exact up to machine precision.

References

Ahlfors, L. V. 1979. *Complex Analysis*. 3rd ed. McGraw-Hill.

Andrews, B., M. Calder, and R. Davis. 2009. "Maximum Likelihood Estimation for α-Stable Autoregressive Process." *Annals of Statistics* 37: 1946–1982.

Blasques, F., S. J. Koopman, G. Mingoli, and S. Telg. 2025. "A Novel Test for the Presence of Local Explosive Dynamics." *Journal of Time Series Analysis* 46, no. 5: 966–980.

de Truchis, G., S. Fries, and A. Thomas. 2025. "Forecasting Extreme Trajectories Using Seminorm Representations". Working Paper 2025–06, Chaire Économie du Climat, Paris.

Fries, S. 2022. "Conditional Moments of Noncausal Alpha-Stable Processes and the Prediction of Bubble Crash Odds." *Journal of Business & Economic Statistics* 40, no. 4: 1–21.

Giancaterini, F., A. Hecq, and C. Morana. 2022. "Is Climate Change Time-Reversible?" *Econometrics* 10, no. 4: 36.

Gourieroux, C., Y. Lu, and C. Y. Robert. 2025. "The Causal–Noncausal Tail Processes: An Introduction". arXiv Preprint arXiv:2506.04046.

Gouriéroux, C., and J. Jasiak. 2016. "Filtering, Prediction and Simulation Methods for Noncausal Processes." *Journal of Time Series Analysis* 37: 405–430.

Gouriéroux, C., and J. Jasiak. 2023. "Generalized Covariance Estimator." *Journal of Business & Economic Statistics* 41, no. 4: 1315–1327.

Gouriéroux, C., J. Jasiak, and A. Monfort. 2020. "Stationary Bubble Equilibria in Rational Expectation Models." *Journal of Econometrics* 218, no. 2: 714–735.

Gouriéroux, C., A. Monfort, and J.-P. Renne. 2020. "Identification and Estimation in Non-Fundamental Structural VARMA Models." *Review of Economic Studies* 87, no. 4: 1915–1953.

Gouriéroux, C., and J.-M. Zakoian. 2017. "Local Explosion Modelling by Non-Causal Process." *Journal of the Royal Statistical Society, Series B: Statistical Methodology* 79: 737–756.

Hecq, A., J. V. Issler, and S. Telg. 2020. "Mixed Causal–Noncausal Autoregressions With Exogenous Regressors." *Journal of Applied Econometrics* 35, no. 3: 328–343.

Hecq, A., S. Telg, and L. Lieb. 2017. "Simulation, Estimation and Selection of Mixed Causal-Noncausal Autoregressive Models: The MARX Package." *SSRN*. https://ssrn.com/abstract=3015797.

Hecq, A., and D. Velasquez Gaviria. 2025. "Non-Causal and Non-Invertible ARMA Models: Identification, Estimation and Application in Equity Portfolios." *Journal of Time Series Analysis* 46, no. 2: 325–352.

Lanne, M., and P. Saikkonen. 2011. "Noncausal Autogressions for Economic Time Series." *Journal of Time Series Econometrics* 3.

Velasco, C., and I. N. Lobato. 2018. "Frequency Domain Minimum Distance Inference for Possibly Noninvertible and Noncausal ARMA Models." *Annals of Statistics* 46, no. 2: 555–579.