

TI 2023-054/III
Tinbergen Institute Discussion Paper

Pseudo-variance quasi-maximum likelihood estimation of semi- parametric time series models

*Mirko Armillotta*¹
*Paolo Gorgi*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

Pseudo-variance quasi-maximum likelihood estimation of semi-parametric time series models*

Mirko Armillotta^{1,2} and Paolo Gorgi^{1,2}

¹Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

²Tinbergen Institute, Amsterdam, The Netherlands.

September 6, 2023

Abstract

We propose a novel estimation approach for a general class of semi-parametric time series models where the conditional expectation is modeled through a parametric function. The proposed class of estimators is based on a Gaussian quasi-likelihood function and it relies on the specification of a parametric pseudo-variance that can contain parametric restrictions with respect to the conditional expectation. The specification of the pseudo-variance and the parametric restrictions follow naturally in observation-driven models with bounds in the support of the observable process, such as count processes and double-bounded time series. We derive the asymptotic properties of the estimators and a validity test for the parameter restrictions. We show that the results remain valid irrespective of the correct specification of the pseudo-variance. The key advantage of the restricted estimators is that they can achieve higher efficiency compared to alternative quasi-likelihood methods that are available in the literature. Furthermore, the testing approach can be used to build specification tests for parametric time series models. We illustrate the practical use of the methodology in a simulation study and two empirical applications featuring integer-valued autoregressive processes, where assumptions on the dispersion of the thinning operator are formally tested, and autoregressions for double-bounded data with application to a realized correlation time series.

Keywords: Double-bounded time series, integer-valued autoregressions, quasi-maximum likelihood.

JEL codes: C32, C52, C58.

*Corresponding author: Paolo Gorgi. Email address: p.gorgi@vu.nl.

1 Introduction

A wide range of time series models have been proposed in the literature to model the conditional mean of time series data. Their specification often depends on the nature of the time series variable of interest. For example, AutoRegressive Moving Average (ARMA) models (Box et al., 1970) are typically employed for time series variables that are continuous and take values on the real line. INteger-valued AutoRegressive (INAR) models (Al-Osh and Alzaid, 1987; McKenzie, 1988) and INteger-valued GARCH models (INGARCH) (Heinen, 2003; Ferland et al., 2006) are designed to account for the discrete and non-negative nature of count processes. Autoregressive Conditional Duration (ACD) models (Engle and Russell, 1998) are used for modeling non-negative continuous processes. Beta autoregressive models (Rocha and Cribari-Neto, 2009) are employed for modeling double-bounded time series data lying in a specified interval domain. The estimation of such models can be carried out by the Maximum Likelihood Estimator (MLE), which constitutes the gold standard approach for the estimation of unknown parameters in parametric models. However, the MLE requires parametric assumptions on the entire distribution of the time series process. This feature is not appealing when the interest of the study is only on modeling the conditional mean instead of the entire distribution. Furthermore, the likelihood function can sometimes present a complex form and the implementation of the MLE can become unfeasible. For instance, exact likelihood inference of INAR models is well-known to be cumbersome and numerically difficult, especially when the order of the model is larger than one (Bu et al., 2008; Drost et al., 2009; Pedeli et al., 2015). In such situations, the use of quasi-likelihood methods becomes attractive.

The Quasi-MLE (QMLE), introduced by Wedderburn (1974), is a likelihood-based estimator where there is a quasi-likelihood that is not necessarily the true distribution of the data. Quasi-likelihoods are typically a member of the one-parameter exponential family. Gouriéroux et al. (1984) show that the QMLE is consistent for the true unknown parameters of the model. Nevertheless, QMLEs can be inefficient because, given a parametric definition for the conditional mean of the process, the conditional variance is implicitly constrained to be a function of the conditional mean as determined by the exponential family of distributions that is considered. In order to improve the estimation efficiency for the parameters of the conditional mean in time series models, Aknouche and Francq (2021) propose a two-stage Weighted Least Squares Estimator (WLSE) where in the first step the conditional variance of the process is estimated and it is then used in the second step as weighting sequence for the solution of the weighted least squares problem. It is shown that this WLSE leads to improved efficiency with respect to QMLE if the variance function is correctly specified. A similar estimator has been more recently proposed in the context of estimating functions approach leading to the same type of efficiency improvement (Francq and Zakoian, 2023).

In this paper, we propose a novel class of QMLEs for the estimation of the conditional expectation of semi-parametric time series models. The estimators are based on a Gaussian quasi-likelihood and a pseudo-variance specification, which can contain restrictions with the parameters of the conditional expectation. The Pseudo-Variance QMLEs (PVQMLEs) only require parametric assumptions on the conditional expectation as the pseudo-variance function does not need to be correctly specified. We establish strong consistency and asymptotic normality of the PVQMLEs under very general conditions. The case in which the pseudo-variance formulation corresponds to the true conditional variance of the process is obtained as a special case. We show that when no restrictions are imposed between the mean and pseudo-variance, the resulting unrestricted PVQMLE has the same asymptotic efficiency of a particular WLSE. Furthermore, if the pseudo-variance is correctly specified it achieves the same asymptotic efficiency as the efficient WLSE. On the other hand, when parameter restrictions are considered, the resulting restricted PVQMLEs can achieve higher efficiency compared to the efficient WLSE and alternative QMLEs. This result is theoretically shown in some special cases and empirically verified for INAR models through an extensive numerical exercise. We discuss how the specification of the pseudo-variance and the parameter restrictions naturally arise for time series processes with bounded support. We obtain that the restricted PVQMLEs retain the desired asymptotic properties when the imposed restrictions are valid with respect to the true parameter of the mean and a pseudo-true parameter of the conditional variance. The validity of such restrictions can be tested without requiring correct specification of the conditional variance. We derive a test for this purpose that can be used as a consistency test for restricted PVQMLEs. When the evidence-based parameter constraints are identified and validated, they constitute a restriction set where an higher-efficiency restricted PVQMLE can be obtained. Furthermore, under correct specification of the pseudo-variance, the test can be used as a specification test on the underlying process generating the data.

Finally, the practical usefulness of PVQMLE approach is illustrated by means of two real data applications. One is concerned with INAR models and one with a Beta autoregression for double-bounded data. INAR processes depend on the distribution assumed for the innovation and the thinning specification (Lu, 2021). Our test allows us to test for the degree of dispersion in the thinning operator as well as the error term. There exists a vast literature of INAR models in testing innovations and marginal distributions dispersion (Schweer and Weiß, 2014; Aleksandrov and Weiß, 2020), testing for serial dependence (Sun and McCabe, 2013), and general goodness of fit tests (Weiß, 2018). However, to the best of our knowledge, specification tests are not available for the thinning dispersion. The thinning operator is typically assumed to be binomial, which implies underdispersion in the thinning. Once appropriate thinning and innovation restrictions are identified

through the specification test, the corresponding PVQMLE is used to estimate the parameters of the INAR model. The second application concerns the analysis of daily realized correlations between a pair of stock returns, which forms a double-bounded time series as the realized correlation takes values between minus one and one. We consider a pseudo-variance specification based on the implied variance from Beta-distributed variables for the definition of PVQMLEs. We then test the validity of parametric restrictions between the mean and pseudo-variance to validate the use of restricted PVQMLEs.

The remainder of the paper is organized as follows. Section 2 introduces the general mean and pseudo-variance framework and the PVQMLEs, together with some examples. Section 3 presents the main theoretical results of the paper and a comparison between the PVQMLE and alternative quasi-likelihood methods. Section 4 introduces the specification test for the validity of the constraints with an extensive simulation study in the case of INAR models. Section 5 presents empirical applications.

2 Specification and estimation

2.1 PVQML estimators

Consider a stationary and ergodic time series process $\{Y_t\}_{t \in \mathbb{Z}}$ with elements taking values in the sample space $\mathcal{Y} \subseteq \mathbb{R}$ and with conditional mean given by

$$E(Y_t | \mathcal{F}_{t-1}) = \lambda(Y_{t-1}, Y_{t-2}, \dots; \psi_0) = \lambda_t(\psi_0), \quad t \in \mathbb{Z}, \quad (1)$$

where \mathcal{F}_t denotes the σ -field generated by $\{Y_s, s \leq t\}$, $\lambda : \mathbb{R}^\infty \times \Psi \rightarrow \mathbb{R}$ is a known measurable function, and $\psi_0 \in \Psi \subset \mathbb{R}^p$ is the true unknown p -dimensional parameter vector. We denote with ν_t the conditional variance of the process, i.e. $V(Y_t | \mathcal{F}_{t-1}) = \nu_t$, which is considered to have an unknown specification. The model is a semi-parametric model as the quantity of interest is the parameter vector of the conditional mean ψ_0 and other distributional properties are left unspecified and treated as an infinite dimensional nuisance parameter. The general specification of the model in (1) includes a wide range of time series models as special case. For instance, it includes linear and non-linear ARMA models when $\mathcal{Y} = \mathbb{R}$, INGARCH and INAR models when $\mathcal{Y} = \mathbb{N}$, ACD models when $\mathcal{Y} = (0, \infty)$, and Beta autoregressive models for bounded data when $\mathcal{Y} = (0, 1)$.

The main objective is to estimate the parameter vector ψ_0 of the conditional expectation. For this purpose, we consider the specification of a pseudo-variance

$$\nu_t^*(\gamma) = \nu^*(Y_{t-1}, Y_{t-2}, \dots; \gamma), \quad t \in \mathbb{Z}, \quad (2)$$

where $\nu^* : \mathbb{R}^\infty \times \Gamma \rightarrow [0, +\infty)$ is a known function that is indexed by the k -dimensional parameter $\gamma \in \Gamma \subset \mathbb{R}^k$. We refer to this as a pseudo-variance as it is not necessarily correctly specified,

i.e. there may be no value $\gamma \in \Gamma$ such that $\nu_t^*(\gamma) = \nu_t$. The idea is to use the pseudo-variance $\nu_t^*(\gamma)$ to enhance the efficiency of the estimation of ψ_0 by means of a Gaussian QMLE. We denote the whole parameter vector that contains both the parameter of the mean and pseudo-variance with $\theta = (\psi', \gamma)'$ and $\theta \in \Theta = \Psi \times \Gamma \subset \mathbb{R}^m$, $m = p + k$.

We introduce the class of PVQMLES that relies on a Gaussian quasi-likelihood for the mean equation with the pseudo-variance as scale of the Gaussian density. We consider estimators based on both unrestricted and restricted quasi-likelihood functions. Assume that we have an observed sample of size T from the process defined in (1), given by $\{Y_t\}_{t=1}^T$. Since $\lambda_t(\psi)$ and $\nu_t^*(\gamma)$ can depend on the infinite past of Y_t , we define their approximations of $\tilde{\lambda}_t(\psi)$ and $\tilde{\nu}_t^*(\gamma)$ based on the available finite sample $\{Y_t\}_{t=1}^T$,

$$\tilde{\lambda}_t(\psi) = \lambda(Y_{t-1}, \dots, Y_1, \tilde{Y}_0, \tilde{Y}_{-1}, \dots; \psi), \quad \tilde{\nu}_t^*(\gamma) = \nu^*(Y_{t-1}, \dots, Y_1, \tilde{Y}_0, \tilde{Y}_{-1}, \dots; \gamma), \quad (3)$$

where $\tilde{Y}_0, \tilde{Y}_{-1}, \dots$ are given initial values. The Gaussian quasi-likelihood for ψ with the pseudo-variance scaling is defined as

$$\tilde{L}_T(\theta) = \frac{1}{T} \sum_{t=1}^T \tilde{l}_t(\theta), \quad \tilde{l}_t(\theta) = -\frac{1}{2} \log \tilde{\nu}_t^*(\gamma) - \frac{[Y_t - \tilde{\lambda}_t(\psi)]^2}{2\tilde{\nu}_t^*(\gamma)}. \quad (4)$$

Based on the quasi-likelihood function in (4), we define the unrestricted and restricted PVQMLE. The unrestricted PVQMLE is based on the unconstrained maximization of the pseudo-likelihood without imposing any constraints between ψ and γ . The unrestricted PVQMLE $\hat{\theta}$ is defined as

$$\hat{\theta} = \arg \max_{\theta \in \Theta} \tilde{L}_T(\theta), \quad (5)$$

where $\hat{\theta} = (\hat{\psi}', \hat{\gamma})'$ and $\hat{\psi}$ is the unrestricted PVQMLE of ψ_0 . In Section 3, we shall see that the unrestricted PVQMLE $\hat{\psi}$ is a consistent estimator of ψ_0 and, in fact, it is asymptotically equivalent to a specific WLSE. If the pseudo-variance is correctly specified, i.e. there is $\gamma_0 \in \Gamma$ such that $\nu^*(\gamma_0) = \nu_t$, then $\hat{\psi}$ is asymptotically equivalent to the efficient WLSE.

In models where the sample space \mathcal{Y} is bounded, such as count-time series models, there can be a natural relationship between the conditional mean and variance of the process. For example, in a count time series process we have that if the mean goes to zero, then also the variance goes to zero as, in fact, the limit case is the mean being exactly zero. Such relationship between mean and variance, as given by parametric models, provide a natural way to introduce restrictions between the mean and pseudo-variance parameters ψ and γ . Several examples are presented at the end of this section.

To specify the restricted PVQMLE, we consider the constrained parameter set Θ_R that imposes r restrictions on the pseudo-variance parameters

$$\Theta_R = \{\theta \in \Theta : S\gamma = g(\psi)\},$$

where S is a $r \times k$ selection matrix and $g : \Psi \rightarrow \mathbb{R}^r$. The estimator derived from the maximization of (4) over the set Θ_R is the restricted PVQMLE,

$$\hat{\theta}_R = \arg \max_{\theta \in \Theta_R} \tilde{L}_T(\theta) \quad (6)$$

where $\hat{\theta}_R = (\hat{\psi}'_R, \hat{\gamma}'_R)'$ and $\hat{\psi}_R$ is the restricted PVQMLE of ψ_0 . In Section 3, we shall see that the restricted PVQMLE $\hat{\psi}_R$ is a consistent estimator of ψ_0 if the constraints in Θ_R hold with respect to a pseudo-true parameter γ^* . The advantage of the restricted PVQMLE $\hat{\psi}_R$ is that it can achieve higher efficiency than the unrestricted one. Furthermore, as it shall be presented in Section 4, the validity of the restrictions can be tested under both misspecification and correct specification of the pseudo-variance. The test can be interpreted as a consistency test for the restricted estimator when the pseudo-variance is misspecified. Instead, it can be employed as a specification test if we assume correct specification of the pseudo-variance. For instance, it shall be employed to test for underdispersion, equidispersion or overdispersion in the thinning operator of INAR models.

2.2 Examples

The model specification in (1) is very general and it covers a wide range of semi-parametric observation-driven time series model. The unrestricted and restricted QMLE based on the pseudo-variance in (2) can be employed for such general class of models. However, PVQMLEs are particularly suited for time series processes where the support of the conditional mean is bounded and a natural relationship with the conditional variance can be assumed. In Section 3.2 it will be shown that in models where conditional mean and pseudo-variance share some parameter restrictions, a more efficient estimator may be obtained with respect to alternative estimation approaches available in the literature. The specification of the pseudo-variance and the parameter restrictions with the conditional mean can be based on well known model specifications. The validity of such restrictions is testable and the asymptotic properties do not require correct specification of the pseudo-variance. This means that no assumptions on the true conditional variance are needed and the consistency of the restricted PVQMLE can also be tested without relying on such assumptions. Below we present some examples of models that are encompassed in the framework defined in equations (1) and (2), and provide a general way to specify the pseudo-variance and the parameter restrictions with the conditional mean.

Example 1 (INAR models). INAR models are widely used in the literature to model count time series. The INAR(1) model is given by

$$Y_t = a \circ Y_{t-1} + \varepsilon_t, \quad t \in \mathbb{Z}, \quad (7)$$

where $\{\varepsilon_t\}_{t \in \mathbb{Z}}$ is an iid sequence of non-negative integer-valued random variables with mean $\omega_1 > 0$ and variance $\omega_2 > 0$, and ‘ \circ ’ is the thinning operator of [Steutel and Van Harn \(1979\)](#). For a given $N \in \mathbb{N}$ and $a \in (0, 1)$, the most general formulation of the thinning operator $a \circ N$ is defined to be a count random variable with mean aN . The most common formulation ([Steutel and Van Harn, 1979](#)) is the Bernoulli thinning where $a \circ N$ is a binomial random variable with N trials and success probability a . The conditional mean of the INAR(1) is

$$\lambda_t = aY_{t-1} + \omega_1,$$

and the pseudo-variance can be specified as

$$\nu_t^* = bY_{t-1} + \omega_2.$$

As discussed in Section 5, several restrictions can be considered for the PVQMLE. For instance, the restriction $b = a(1 - a)$ is implied by a binomial thinning and $\omega_1 = \omega_2$ is implied by a Poisson error.

Example 2 (INGARCH models). Another popular model for time series of counts is the INGARCH model. The conditional mean of the INGARCH(1,1) model takes the form

$$\lambda_t = \omega_1 + \alpha_1 Y_{t-1} + \beta_1 \lambda_{t-1}, \tag{8}$$

where $\omega_1, \alpha_1, \beta_1 \geq 0$. The pseudo-variance can be specified as

$$\nu_t^* = \omega_2 + \alpha_2 Y_{t-1} + \beta_2 \lambda_{t-1}.$$

Also in this case, several restrictions can be considered for the PVQMLE. For instance, the restrictions $\omega_2 = \omega_1$, $\alpha_2 = \alpha_1$ and $\beta_2 = \beta_1$ are implied by an equidispersion assumption $\nu_t^* = \lambda_t$, which follows assuming a conditional Poisson distribution for example. Alternatively, the restrictions $\omega_2 = c\omega_1$, $\alpha_2 = c\alpha_1$ and $\beta_2 = c\beta_1$ with $c > 0$ are implied by a proportional variance assumption $\nu_t^* = c\lambda_t$.

Example 3 (ACD models). ACD models are typically used to model non-negative continuous time series variables, like durations or volumes. These models take the form $Y_t = \lambda_t \varepsilon_t$ where ε_t is a sequence of positive variables with mean equal to 1. The conditional expectation λ_t may take the form as in equation (8). The pseudo-variance can be specified in several ways and restrictions can be imposed. For instance, the restriction $\nu_t^* = \lambda_t^2$ follows by assuming an exponential error distribution. An alternative restriction is given by $\nu_t^* = c\lambda_t^2$, $c > 0$.

Example 4 (double-bounded autoregressions). For double-bounded time series data the conditional mean λ_t can be specified as in equation (8), see [Gorgi and Koopman \(2021\)](#) for instance. Several specifications and restrictions for the pseudo-variance can be considered. For instance, the restriction $\nu_t^* = \lambda_t(1 - \lambda_t)/(1 + \phi)$ is implied by a beta conditional distribution with dispersion parameter ϕ . Intermediate restrictions on the pseudo-variance are discussed in the corresponding application in Section 5.

We note that the example presented in this section are focused on a linear mean equation for simplicity of exposition. Several other non-linear model specifications are encompassed in the general framework in (1) and (2), see for example [Creal et al. \(2013\)](#) and [Christou and Fokianos \(2015\)](#).

3 Asymptotic theory

In this section, the asymptotic properties of the PVQMLEs in (5) and (6) are formally derived. Although asymptotic results related to quasi-maximum likelihood estimators of observation-driven models are well-established in the literature, the associated theory for PVQMLEs differs as it relies on simultaneous estimation of mean and pseudo-variance parameters, where the latter can be misspecified and present parameter restrictions with the mean. Since the pseudo-variance can be misspecified, the estimator of the pseudo-variance parameter $\hat{\gamma}$ will be consistent with respect to a pseudo-true value γ^* , which is given by

$$\gamma^* = \arg \max_{\gamma \in \Gamma} -\frac{1}{2} \mathbb{E} \left(\log \nu_t^*(\gamma) + \frac{[Y_t - \lambda_t(\psi_0)]^2}{\nu_t^*(\gamma)} \right). \quad (9)$$

We define the vector $\theta_0 = (\psi_0', \gamma^{*'})'$ that contains both true and pseudo-true parameters. The estimator of the mean parameters preserves the consistency and asymptotic normality results to the true parameter vector ψ_0 . Moreover, we will show that such result holds for both unrestricted (5) and restricted (6) estimators.

We start by showing consistency and asymptotic normality of the unrestricted PVQMLE in (5). We first obtain the score function related to (4)

$$\tilde{S}_T(\theta) = \frac{1}{T} \sum_{t=1}^T \tilde{s}_t(\theta), \quad \tilde{s}_t(\theta) = \frac{Y_t - \tilde{\lambda}_t(\psi)}{\tilde{\nu}_t^*(\gamma)} \frac{\partial \tilde{\lambda}_t(\psi)}{\partial \theta} + \frac{[Y_t - \tilde{\lambda}_t(\psi)]^2 - \tilde{\nu}_t^*(\gamma)}{2\tilde{\nu}_t^{*2}(\gamma)} \frac{\partial \tilde{\nu}_t^*(\gamma)}{\partial \theta}. \quad (10)$$

Furthermore, we define $L_T(\theta)$, $l_t(\theta)$, $S_T(\theta)$ and $s_t(\theta)$ as the random functions obtained from $\tilde{L}_T(\theta)$, $\tilde{l}_t(\theta)$, $\tilde{S}_T(\theta)$ and $\tilde{s}_t(\theta)$ by substituting $\tilde{\lambda}_t(\psi)$ and $\tilde{\nu}_t^*(\gamma)$ with $\lambda_t(\psi)$ and $\nu_t^*(\gamma)$, respectively. We consider the following assumptions.

A1 The process $\{Y_t, t \in Z\}$ is strictly stationary and ergodic.

A2 $\lambda_t(\cdot)$ is continuous in Ψ , $\nu_t^*(\cdot)$ is continuous in Γ and the set Θ is compact. Moreover,

$$\mathbb{E} \sup_{\gamma \in \Gamma} |\log \nu_t^*(\gamma)| < \infty, \quad \mathbb{E} \sup_{\theta \in \Theta} \frac{[Y_t - \lambda_t(\psi)]^2}{\nu_t^*(\gamma)} < \infty.$$

A3 $\lambda_t(\psi) = \lambda_t(\psi_0)$ a.s. if and only if $\psi = \psi_0$.

A4 There is a constant $\underline{\nu}^* > 0$ such that $\nu_t^*(\gamma), \tilde{\nu}_t^*(\gamma) \geq \underline{\nu}^*$ for any $t \geq 1$ and any $\gamma \in \Gamma$.

A5 Define $a_t = \sup_{\psi \in \Psi} |\tilde{\lambda}_t(\psi) - \lambda_t(\psi)|$ and $b_t = \sup_{\gamma \in \Gamma} |\tilde{\nu}_t^*(\gamma) - \nu_t^*(\gamma)|$, it holds that

$$\lim_{t \rightarrow \infty} \left(1 + |Y_t| + \sup_{\psi \in \Psi} |\lambda_t(\psi)|\right) a_t = 0, \quad \lim_{t \rightarrow \infty} \left(1 + Y_t^2 + \sup_{\psi \in \Psi} \lambda_t^2(\psi)\right) b_t = 0 \quad a.s.$$

A6 The pseudo-true parameter $\gamma^* \in \Gamma$ defined in (9) is unique.

A7 Define $c_t = \sup_{\theta \in \Theta} \|\partial \tilde{\lambda}_t(\psi)/\partial \theta - \partial \lambda_t(\psi)/\partial \theta\|$, $d_t = \sup_{\theta \in \Theta} \|\partial \tilde{\nu}_t^*(\gamma)/\partial \theta - \partial \nu_t^*(\gamma)/\partial \theta\|$. The following quantities are of order $\mathcal{O}(t^{-\delta})$ a.s. for some $\delta > 1/2$

$$\begin{aligned} & \sup_{\theta \in \Theta} \left\| \frac{\partial \lambda_t(\psi)}{\partial \theta} \right\| a_t, \quad \sup_{\theta \in \Theta} \left\| \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \right\| \left(1 + |Y_t| + \sup_{\psi \in \Psi} |\lambda_t(\psi)|\right) a_t, \\ & \sup_{\theta \in \Theta} \left\| \frac{\partial \lambda_t(\psi)}{\partial \theta} \right\| \left(|Y_t| + \sup_{\psi \in \Psi} |\lambda_t(\psi)|\right) b_t, \quad \sup_{\theta \in \Theta} \left\| \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \right\| \left(1 + Y_t^2 + \sup_{\psi \in \Psi} \lambda_t^2(\psi)\right) b_t, \\ & (1 + |Y_t| + \sup_{\psi \in \Psi} |\lambda_t(\psi)|) c_t, \quad (1 + Y_t^2 + \sup_{\psi \in \Psi} \lambda_t^2(\psi)) d_t. \end{aligned}$$

A8 $\lambda_t(\cdot)$ and $\nu_t^*(\cdot)$ have continuous second-order derivatives in their spaces. Moreover,

$$\begin{aligned} & \mathbb{E} \sup_{\theta \in \Theta} \frac{[Y_t - \lambda_t(\psi)]^4}{\nu_t^{*2}(\gamma)} < \infty, \quad \mathbb{E} \sup_{\theta \in \Theta} \left\| \frac{1}{\sqrt{\nu_t^*(\gamma)}} \frac{\partial^2 \lambda_t(\psi)}{\partial \theta \partial \theta'} \right\|^2 < \infty, \\ & \mathbb{E} \sup_{\theta \in \Theta} \left\| \frac{1}{\nu_t^*(\gamma)} \frac{\partial \lambda_t(\psi)}{\partial \theta} \frac{\partial \lambda_t(\psi)}{\partial \theta'} \right\| < \infty, \quad \mathbb{E} \sup_{\theta \in \Theta} \left\| \frac{1}{\nu_t^*(\gamma)} \frac{\partial \lambda_t(\psi)}{\partial \theta} \frac{\partial \nu_t^*(\gamma)}{\partial \theta'} \right\|^2 < \infty, \\ & \mathbb{E} \sup_{\theta \in \Theta} \left\| \frac{1}{\nu_t^{*2}(\gamma)} \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \frac{\partial \nu_t^*(\gamma)}{\partial \theta'} \right\|^2 < \infty, \quad \mathbb{E} \sup_{\theta \in \Theta} \left\| \frac{1}{\nu_t^*(\gamma)} \frac{\partial^2 \nu_t^*(\gamma)}{\partial \theta \partial \theta'} \right\|^2 < \infty. \end{aligned}$$

A9 The matrices $H(\theta_0) = \mathbb{E}[-\partial^2 l_t(\theta_0)/\partial \theta \partial \theta']$, $I(\theta_0) = \mathbb{E}[s_t(\theta_0)s_t'(\theta_0)]$ exist with $H(\theta_0)$ invertible.

A10 $\theta_0 \in \dot{\Theta}$, where $\dot{\Theta}$ is the interior of Θ .

A11 The sequence $\sqrt{T}S_T(\theta_0)$ obeys the central limit theorem.

The strict stationarity and ergodicity in assumption **A1** depends upon the model formulation in (1) and (2) and it can be established by means of different probabilistic approaches, see for instance [Straumann and Mikosch \(2006\)](#) and [Debaly and Truquet \(2021\)](#). Assumption **A2** is a standard moment condition. Assumption **A3** is required for the identification of the true parameter ψ_0 . Assumptions **A5** and **A7** are needed to guarantee that the initialization of filters in (3) is asymptotically irrelevant. Assumption **A6** imposes the uniqueness of the pseudo-true parameter for the variance equation. In Corollary 3 below, we show that this assumption can be dropped if the researcher is not interested in the asymptotic normality of the estimator but only in the consistency. Assumptions **A8** and **A9** impose moments on the second derivatives of the log-quasi-likelihood that are required for asymptotic normality to apply. Assumption **A10** is the standard condition for asymptotic normality that the true parameter value is in the interior of the parameter set. Finally, assumption **A11** is an high-level condition that a central limit theorem applies to the score. This condition is left high-level for generality purposes since the score function $s_t(\theta_0)$ is not a martingale difference sequence, see equation (10). There are several alternative CLTs for non-martingale sequences and the choice of the most appropriate one is strongly dependent on the specific mean-variance model formulation. For example, CLTs appealing the concept of mixing processes or mixingales are widely available, see the surveys in [Doukhan \(1994\)](#), [Bradley \(2005\)](#) and [White \(1994\)](#). In case of correct conditional variance specification then assumption **A11** can be dropped, see Corollary 2. Theorem 1 delivers the consistency and asymptotic normality of the unrestricted PVQMLE of the true parameter ψ_0 .

Theorem 1. *Consider the unrestricted PVQMLE in (5). Under conditions **A1-A6***

$$\hat{\psi} \rightarrow \psi_0, \quad a.s. \quad T \rightarrow \infty. \quad (11)$$

Moreover, if also **A7-A11** hold, as $T \rightarrow \infty$

$$\sqrt{T} \left(\hat{\psi} - \psi_0 \right) \xrightarrow{d} N(0, \Sigma_\psi), \quad \Sigma_\psi = H_\psi^{-1}(\theta_0) I_\psi(\theta_0) H_\psi^{-1}(\theta_0), \quad (12)$$

where

$$H_\psi(\theta_0) = E \left[\frac{1}{\nu_t^*(\gamma^*)} \frac{\partial \lambda_t(\psi_0)}{\partial \psi} \frac{\partial \lambda_t(\psi_0)}{\partial \psi'} \right], \quad I_\psi(\theta_0) = E \left[\frac{\nu_t}{\nu_t^{*2}(\gamma^*)} \frac{\partial \lambda_t(\psi_0)}{\partial \psi} \frac{\partial \lambda_t(\psi_0)}{\partial \psi'} \right]. \quad (13)$$

The asymptotic properties of the estimator of the variance parameter γ are obtained in Corollary 1 below. Let $s_t(\theta_0)' = [s_t^{(\psi)}(\theta_0)', s_t^{(\gamma)}(\theta_0)']'$ be the partition of the score with respect to the mean and (pseudo-)variance parameters. Let $H_\gamma(\theta_0) = E[-\partial^2 l_t(\theta_0)/\partial \gamma \partial \gamma']$ and $I_\gamma(\theta_0) = E[s_t^{(\gamma)}(\theta_0) s_t^{(\gamma)}(\theta_0)']$.

Corollary 1. *Under the assumptions of Theorem 1 we have that as $T \rightarrow \infty$, a.s. $\hat{\gamma} \rightarrow \gamma^*$ and $\sqrt{T}(\hat{\gamma} - \gamma^*) \xrightarrow{d} N(0, \Sigma_\gamma)$, where $\Sigma_\gamma = H_\gamma^{-1}(\theta_0) I_\gamma(\theta_0) H_\gamma^{-1}(\theta_0)$.*

Theorem 1 determines the asymptotic distribution of the unrestricted PVQMLE of ψ_0 without requiring correct specification of the pseudo-variance. The following result shows that in the special case in which the variance is well-specified then the estimator $\hat{\psi}$ gains in efficiency.

Corollary 2. *Consider the assumptions of Theorem 1. If, in addition, the variance (2) is correctly specified, i.e. $\nu_t^*(\gamma^*) = \nu_t$, then **A1-A10** entail (11) and*

$$\sqrt{T} \left(\hat{\psi} - \psi_0 \right) \xrightarrow{d} N(0, I_\psi^{-1}), \quad I_\psi = \mathbf{E} \left[\frac{1}{\nu_t} \frac{\partial \lambda_t(\psi_0)}{\partial \psi} \frac{\partial \lambda_t(\psi_0)}{\partial \psi'} \right], \quad (14)$$

where $\Sigma_\psi - I_\psi^{-1}$ is positive semi-definite.

We also note that in Corollary 2 the uniqueness of the variance parameter in assumption **A6** is implied by the condition $\nu_t^*(\gamma) = \nu_t^*(\gamma^*)$ a.s. if and only if $\gamma = \gamma^*$. This follows immediately from the correct specification of the pseudo-variance. Corollary 3 below shows that even if the pseudo-true parameter γ^* is not unique, i.e. assumption **A6** does not hold, the consistency of the unrestricted estimator $\hat{\psi}$ is retained without any additional assumption. The overall estimator $\hat{\theta}$ will instead be set consistent over the set of values that maximize the limit of the quasi-likelihood, Θ_0 , since the pseudo-true parameter γ^* is not uniquely identified.

Corollary 3. *Consider the unrestricted PVQMLE (5) and assume conditions **A1-A5** hold. Then, as $T \rightarrow \infty$, $\inf_{\theta_0 \in \Theta_0} \|\hat{\theta} - \theta_0\| \rightarrow 0$ a.s. and $\hat{\psi} \rightarrow \psi_0$ a.s.*

We now treat the case in which the conditional mean and pseudo-variance parameters are constrained. We study the asymptotic properties of the restricted PVQMLE $\hat{\psi}_R$ defined in (6).

A12 The equality $S\gamma^* = g(\psi_0)$ holds and $g(\cdot)$ is continuous.

Assumption **A12** is required to ensure that $\theta_0 \in \Theta_R$, i.e. the imposed restrictions are valid with respect to the true parameter ψ_0 and the pseudo-true parameter γ^* . The continuity of $g(\cdot)$ guarantees that Θ_R remains compact. Define $\gamma = (\gamma'_1, \gamma'_2)'$ where $\gamma_1 = S\gamma = g(\psi)$ is the sub-vector of pseudo-variance parameters that are restricted to mean parameters and γ_2 constitutes the sub-vector of remaining free parameters. For $\theta \in \Theta_R$, with some abuse of notation, we have $\theta = (\psi', \gamma'_1, \gamma'_2)' = (\psi', g(\psi)', \gamma'_2)' = (\psi', \gamma'_2)'$. Recall that $H_x(\theta_0) = \mathbf{E}[-\partial^2 l_t(\theta_0)/\partial x \partial x']$ and $I_x(\theta_0) = \mathbf{E}[s_t^{(x)}(\theta_0) s_t^{(x)'}(\theta_0)']$. Moreover, define $H_{x,z}(\theta_0) = \mathbf{E}[-\partial^2 l_t(\theta_0)/\partial x \partial y']$, $I_{x,z}(\theta_0) = \mathbf{E}[s_t^{(x)}(\theta_0) s_t^{(z)'}(\theta_0)']$ and $I_{z,x}(\theta_0) = I'_{x,z}(\theta_0)$. Analogously, set $D(\theta_0) = H^{-1}(\theta_0)$ and $D_{x,y}(\theta_0)$ being the corresponding partition related to rows x and columns y of $D(\theta_0)$. Theorem 2 delivers the asymptotic distribution of the restricted PVQMLE.

Theorem 2. Consider the restricted PVQMLE in (6). Under conditions **A1-A6** and **A12**

$$\hat{\psi}_R \rightarrow \psi_0, \quad a.s. \quad T \rightarrow \infty. \quad (15)$$

Moreover, if also **A7-A11** hold, as $T \rightarrow \infty$

$$\sqrt{T} \left(\hat{\psi}_R - \psi_0 \right) \xrightarrow{d} N(0, \Sigma_R), \quad (16)$$

where

$$\begin{aligned} \Sigma_R = & D_\psi(\theta_0) I_\psi(\theta_0) D_\psi(\theta_0) + D_{\psi, \gamma_2}(\theta_0) I_{\gamma_2, \psi}(\theta_0) D_\psi(\theta_0) \\ & + D_\psi(\theta_0) I_{\psi, \gamma_2}(\theta_0) D_{\gamma_2, \psi}(\theta_0) + D_{\psi, \gamma_2}(\theta_0) I_{\gamma_2}(\theta_0) D_{\gamma_2, \psi}(\theta_0). \end{aligned} \quad (17)$$

We note that Corollaries 1-3 can easily be adapted to hold also for $\hat{\theta}_R$. In Section 3.2 below, we shall see that the restricted PVQMLE can lead to substantial gains in efficiency with respect to the unrestricted PVQMLE. The consistency of the restricted PVQMLE requires the additional assumption L. However, as discussed in Section 4, this assumption can be tested and the correct specification of the pseudo-variance is not required. Clearly, when ψ and γ do not have parameter restrictions, i.e. $\hat{\psi}_R = \hat{\psi}$, it can be noted that Theorem 1 is equivalent to Theorem 2 with $\Sigma_R = \Sigma_\psi$, since $H_{\psi, \gamma_2}(\theta_0) = 0$, $H(\theta_0)$ becomes block diagonal, its inverse has block elements $D_x(\theta_0) = H_x^{-1}(\theta_0)$ and $D_{x, y}(\theta_0) = D_{y, x}(\theta_0) = 0$, implying that $\Sigma_R = \Sigma_\psi$.

3.1 Comparison to alternative estimators

In this section, we show that the unrestricted PVQMLE achieves the same asymptotic variance of existing estimators. Consider the unrestricted PVQMLE depicted in Theorem 1. The partition of the score related to the mean parameter ψ is

$$\tilde{s}_t^{(\psi)}(\theta) = \frac{Y_t - \tilde{\lambda}_t(\psi)}{\tilde{v}_t^*(\gamma)} \frac{\partial \tilde{\lambda}_t(\psi)}{\partial \psi}. \quad (18)$$

We compare (18) with some alternative semi-parametric estimators presented in the literature.

Consider the two-stage Weighted Least Squares (WLSE) of [Aknouche and Francq \(2021\)](#) defined as

$$\hat{\psi}_W = \arg \max_{\psi \in \Psi} \frac{1}{T} \sum_{t=1}^T \tilde{l}_{s_t}(\psi, \hat{w}_t), \quad \tilde{l}_{s_t}(\psi, \hat{w}_t) = -\frac{[Y_t - \tilde{\lambda}_t(\psi)]^2}{\hat{w}_t},$$

where \hat{w}_t is a first-step estimator of the set of weights w_t . The resulting score of the WLSE is

$$\tilde{s}_t(\psi, \hat{w}_t) = \frac{Y_t - \tilde{\lambda}_t(\psi)}{\hat{w}_t} \frac{\partial \tilde{\lambda}_t(\psi)}{\partial \psi}. \quad (19)$$

Since it is well-known that the conditional variance is the optimal weight for the WLSE, the same authors set $w_t = \nu_t^*(\xi) = \nu^*(Y_{t-1}, Y_{t-2}, \dots; \xi)$ by defining a functional form for a pseudo-variance, where the parameters ξ^* may also contain ψ_0 or parts of it. The corresponding first-step estimated weights are $\hat{w}_t = \tilde{\nu}_t^*(\hat{\xi})$, where $\hat{\xi}$ represents the first-step estimate of the parameter ξ .

Consider the general QMLE of [Wedderburn \(1974\)](#) and [Gourieroux et al. \(1984\)](#) based on the exponential family of quasi-likelihoods defined as

$$\hat{\psi}_Q = \arg \max_{\psi \in \Psi} \tilde{l}_T(\psi),$$

where the log-quasi-likelihood $\tilde{l}_T(\psi)$ is a member of the one-parameter exponential family with respect to $\tilde{\lambda}_t(\psi)$. The corresponding score is given by

$$\tilde{s}_t(\psi) = \frac{Y_t - \tilde{\lambda}_t(\psi)}{\tilde{\nu}_t(\psi)} \frac{\partial \tilde{\lambda}_t(\psi)}{\partial \psi}. \quad (20)$$

where the conditional variance $\tilde{\nu}_t(\psi)$ is typically a function of the mean, i.e. $\tilde{\nu}_t(\psi) = h(\tilde{\lambda}_t(\psi))$ for some function $h(\cdot)$. For example, selecting the Poisson quasi-likelihood yields $\tilde{\nu}_t(\psi) = \tilde{\lambda}_t(\psi)$ ([Ahmad and Francq, 2016](#)), see [Aknouche and Francq \(2021, Sec. 2.2\)](#) for other examples.

The expression of the scores in (18)-(20) highlight how the unrestricted PVQMLE is closely related to WLSE and the QMLE based on the exponential family. The main difference between the unrestricted PVQMLE and the QMLE with score in (20) is that the QMLE only considers the specification of the conditional mean and the conditional variance is a function of the conditional mean that is implied by the selected distribution in the exponential family. On the the other hand, the unrestricted PVQMLE differs from the WLSE as the parameters are estimated jointly instead of a multi-step estimation. The unrestricted PVQMLE, the QMLE and the WLSE enjoy the same consistency property for the mean parameters ψ_0 irrespective of the correct specification of the conditional variance. Furthermore, when they have the same specification of the conditional variance, these estimators are asymptotically equivalent.

Corollary 4. *Assume Theorem 1 holds. Moreover, suppose the WLSE (19) with $w_t = \nu_t^*(\gamma^*)$ is consistent and asymptotically normal with limiting variance Σ_W . Then the unrestricted PVQMLE in (5) is asymptotically as efficient as the WLSE, meaning that $\Sigma_\psi = \Sigma_W$. In addition, if $\nu_t^*(\cdot) = \nu_t(\cdot)$, then $\Sigma_W = \Sigma_\psi = I_\psi^{-1}$.*

The result in Corollary 4 follows immediately from Theorem 1 and Corollary 2. We also note that if Corollary 4 holds then also Corollaries 2.1-2.3 in [Aknouche and Francq \(2021\)](#) hold for the unrestricted PVQMLE. This has two direct consequences: (i) if the variance is well-specified, the unrestricted PVQMLE is asymptotically more efficient than the QMLE of ψ_0 , if the variance implied

by the exponential family is not the true one, and (ii) if the conditional distribution of Y_t comes from the exponential family, then the well-specified PVQMLE is asymptotically as efficient as the MLE of ψ_0 .

We note that the comparison discussed so far only concerns the unrestricted PVQMLE. This equivalence of the PVQMLE with respect to the WLSE and the QMLE does not hold for the restricted PVQMLE. This can be noted from the form of the score function given in equation (10) and the fact that the partial derivative of $\tilde{\nu}_t^*(\gamma)$ with respect to ψ is no longer equal to zero. Below we discuss how the restricted PVQMLE can achieve higher efficiency compared to the unrestricted PVQMLE.

3.2 Efficiency of PVQMLE

Given that the PVQMLE with distinct parameters on mean and pseudo-variance is asymptotically equivalent to the WLSE for the mean parameters ψ_0 (Corollary 4), it may be expected that if the mean and pseudo-variance equations share common parameters in θ , i.e. ψ_0 and γ^* are not completely distinct so that $\theta_0 \in \Theta_R$, then the restricted PVQMLE in (6) could show improved efficiency over the unrestricted PVQMLE and the WLSE. This result cannot be proved in general but for the following special cases it is verified.

A13 One of the following conditions holds:

A13.a $Y_t | \mathcal{F}_{t-1} \sim q(\lambda_t, \nu_t)$ where $q(\cdot)$ is Gaussian.

A13.b $m = 1$, $q(\cdot)$ is symmetric and is meso- or platy-kurtic.

A13.c $m = 1$, the first derivatives of the functions $\lambda_t(\psi_0)$ and $\nu_t(\gamma_0)$ have the (opposite) same sign, and $q(\cdot)$ is (positive) negative skewed and meso- or platy-kurtic.

Proposition 1. *Assume that Assumptions **A1-A13** hold with $\nu_t^*(\gamma^*) = \nu_t$. Moreover, suppose that the WLSE in (19) with $w_t = \nu_t$ is consistent and asymptotically normal with asymptotic variance I_ψ^{-1} . Then, the restricted PVQMLE in (6) is asymptotically more efficient than the unrestricted PVQMLE and the WLSE, i.e. $I_\psi^{-1} - \Sigma_R$ is positive semi-definite.*

The conditions stated in Assumption M can be somewhat restrictive, however, we note that they are only sufficient conditions. In general, it is not straightforward to derive sharper theoretical conditions under which the restricted PVQMLE is more efficient than the unrestricted PVQMLE. However, for specific models, we can appeal to numerical methods to obtain the asymptotic covariance matrix of the two estimators and evaluate their relative efficiency.

We consider the INAR(1) model in (7) with binomial thinning and Poisson error distribution as an example. The unrestricted PVQMLE $\hat{\psi}$ is based on the following conditional mean and pseudo-variance equations

$$\lambda_t(\psi) = aY_{t-1} + \omega_1, \quad \nu_t^*(\gamma) = bY_{t-1} + \omega_2, \quad (21)$$

where $\psi' = (a, \omega_1)$ and $\gamma' = (b, \omega_2)$. Instead, the restricted PVQMLE $\hat{\psi}_R$ imposes the restrictions $b = a(1 - a)$ and $\omega_2 = \omega_1 = \omega$.

We focus on the analysis of the asymptotic variances of these estimators. To this aim, we simulate a long time series ($T = 10,000$) from the INAR(1) process (binomial thinning and Poisson errors) for different values of the parameters a and ω_1 over a grid. The asymptotic covariance matrices of the two estimators are computed by approximating their expectations with the corresponding sample means. Figure 1 reports an heatmap plot of the ratio (in \log_{10} scale) between the asymptotic variance of the unrestricted and the restricted PVQMLEs for the parameter estimates of a and ω_1 . The regions of the parameter set where the \log_{10} -variance ratio is greater than zero, i.e. variance ratio is greater than one, indicate the parameter values for which the restricted estimator is more efficient of the unrestricted one, and vice versa. The pictures suggest that the restricted estimator $\hat{\psi}_R$ is more efficient than the unrestricted estimator $\hat{\psi}$ in most cases, except when a and ω_1 are close to zero. Furthermore, the lack of efficiency of the restricted PVQMLE in the green areas is showed to be minimal. For example, a \log_{10} -variance ratio around -0.05 indicates a variance ratio around 0.9. Therefore, for small values of a and ω_1 the two estimators are essentially equivalent. Instead, for larger values of a and ω_1 , the variance ratio gets substantially larger with the unrestricted PVQMLE estimator having up to 30 times larger variance of the restricted one. This is further illustrated in Figure 2, which displays a graph of cross-section of the \log_{10} -variance ratio for some fixed values of a and ω_1 .

Another way to grasp the intuition behind the improved efficiency of the restricted PVQMLE comes from the literature on saddlepoint approximations (Daniels, 1954). Saddlepoint approximations are used to approximate a density function with a function that is based on the cumulant generating function of the data, which is typically called saddlepoint density. Pedeli et al. (2015) show that the conditional saddlepoint density can approximate the conditional density of the INAR(p) model in (7) to a certain degree of accuracy. It is not hard to see that the conditional saddlepoint density is approximately equal to the pseudo-variance quasi-likelihood in (4) with correctly specified variance (Pedeli et al., 2015, Sec. 3.4). Therefore, when the variance is correctly specified, the restricted PVQMLE of the INAR(p) model is close to the maximizer of the log-likelihood obtained by the saddlepoint density, which in turn is expected to get closer to the MLE as $\lambda_t \rightarrow \infty$. This is confirmed empirically from the results in figures 1 and 2, where the efficiency of restricted PVQMLE

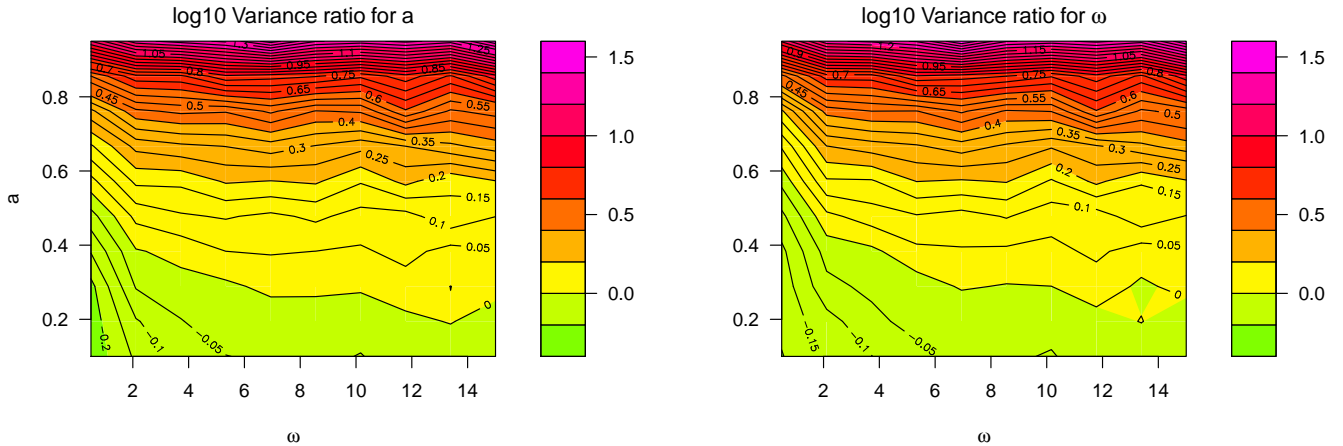


Figure 1: Contour plots of log₁₀-variance ratios for the INAR coefficients. Left: ratio $\log_{10}[Var(\hat{a})/Var(\hat{a}_R)]$ plotted for several values of a and ω . Right: ratio $\log_{10}[Var(\hat{\omega})/Var(\hat{\omega}_R)]$ plotted for several values of a and ω . The green area indicates a variance ratio smaller than one.




Figure 2: Log₁₀-variance ratios plots for the INAR coefficients. Dashed red line: y -axis=0. Left: ratio $\log_{10}[Var(\hat{a})/Var(\hat{a}_R)]$ plotted for several values of a and $\omega = 3$. Right: ratio $\log_{10}[Var(\hat{\omega})/Var(\hat{\omega}_R)]$ plotted for several values of ω and $a = 0.85$.

over the unrestricted PVQMLE grows as $a, w \rightarrow \infty$ i.e. where restricted PVQMLE approximates more accurately the MLE. We conjecture that similar results may apply also to other models. For case of independent observations, [Goodman \(2022\)](#) has recently shown that the approximation error in using saddlepoint approximation is negligible compared to the inferential uncertainty inherent in the MLE. Although the literature is still under development, these arguments provide reliable evidence on the higher asymptotic performance of restricted PVQMLEs compared to the unrestricted one and other quasi-likelihood methods presented in [Section 3.1](#).

Finally, we consider a simulation study to assess the small sample properties of PVQMLEs in comparison with several other alternative estimators. The study consists of 1,000 Monte Carlo replications where we generate data from the Poisson INAR(1) process and estimate the mean parameter vector ψ . We consider several PVQMLEs based on different restrictions of the variance parameter vector γ . The unrestricted PVQMLE $\hat{\psi}$ is based on the mean and pseudo-variance equations in (21). The first restricted PVQMLE $\hat{\psi}_{R_1}$ imposes the restriction $R_1 : b = a(1 - a)$, the second restricted PVQMLE $\hat{\psi}_{R_2}$ imposes the restriction $R_2 : \omega_2 = \omega_1$, and the third restricted PVQMLE $\hat{\psi}_{R_3}$ imposes the restriction $R_3 : b = a(1 - a), \omega_2 = \omega_1$. Furthermore, we consider the QMLE based on the Poisson quasi-likelihood $\hat{\psi}_Q$, the conditional least squares estimator (CLSE) $\hat{\psi}_{LS}$, the WLSE in (19) with weights $\hat{w}_t = \hat{a}_{LS}(1 - \hat{a}_{LS})Y_{t-1} + \hat{\omega}_{LS}$ where $(\hat{\omega}_{LS}, \hat{a}_{LS})' = \hat{\psi}_{LS}$ are first-step estimates obtained from the CLSE, and the unfeasible WLSE $\hat{\psi}_{WUN}$ with weights given by the true conditional variance. The results of the simulation study are reported in Table 1.

Table 1: Bias and RMSE of estimators of the mean parameters when the data generating process is an INAR(1) with $a = 0.85$ and $\omega = 3$, and sample sizes $T = 500$ and $T = 2000$.

Est.	$T = 500$				$T = 2000$			
	ω_1		a		ω_1		a	
	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE
$\hat{\psi}_Q$	0.1190	0.5080	-0.0063	0.0255	0.0323	0.2425	-0.0016	0.0121
$\hat{\psi}_{LS}$	0.1206	0.5003	-0.0064	0.0251	0.0290	0.2401	-0.0014	0.0119
$\hat{\psi}_W$	0.1175	0.5001	-0.0062	0.0251	0.0301	0.2392	-0.0015	0.0119
$\hat{\psi}_{WUN}$	0.1174	0.5001	-0.0062	0.0251	0.0300	0.2393	-0.0015	0.0119
$\hat{\psi}$	0.1159	0.5021	-0.0061	0.0252	0.0295	0.2388	-0.0015	0.0119
$\hat{\psi}_{R_1}$	0.1103	0.4819	-0.0058	0.0241	0.0305	0.2314	-0.0015	0.0115
$\hat{\psi}_{R_2}$	0.1109	0.4911	-0.0059	0.0246	0.0246	0.2332	-0.0012	0.0115
$\hat{\psi}_{R_3}$	0.0052	0.2028	-0.0006	0.0098	0.0027	0.1011	-0.0001	0.0049

Since the PVQMLE without constraints on the first two moments is asymptotically equivalent to the WLSE, it can be expected that the restricted PVQMLE where suitable constraints corresponding to the true model are imposed should show improved performances over the other estimators. Indeed, from Table 1 it can be seen that QMLE, CLSE, WLSE and unrestricted PVQMLE of model (21) share similar performances both in terms of bias and RMSE. Instead, a partial specification of the true constraints underlying the model in $\hat{\psi}_{R_1}$ and $\hat{\psi}_{R_2}$ already leads to an improvement with respect to the other estimation techniques; such improvement becomes substantial in $\hat{\psi}_{R_3}$ where all the correct constraints are considered.

4 Testing restrictions

In the previous section, we have seen that correctly identified constraints on mean and pseudo-variance equations can deliver a restricted PVQMLE with improved efficiency. In this section, we develop a test based on the unrestricted estimator in (5) which allows us to test the validity of the restriction $S\gamma = g(\psi)$. We define $r(\theta) = S\gamma - g(\psi)$ and we denote with $\Sigma(\theta_0) = H^{-1}(\theta_0)I(\theta_0)H^{-1}(\theta_0)$ the asymptotic covariance matrix of the entire unrestricted estimator vector $\hat{\theta}$. Moreover, consider the following plug-in estimators of $H(\theta_0)$ and $I(\theta_0)$ given by $\tilde{H}_T(\hat{\theta}) = T^{-1} \sum_{t=1}^T -\partial^2 \tilde{l}_t(\hat{\theta}) / \partial \theta \partial \theta'$ and $\tilde{I}_T(\hat{\theta}) = T^{-1} \sum_{t=1}^T \tilde{s}_t(\hat{\theta}) \tilde{s}_t'(\hat{\theta})$, respectively. The following result holds.

Proposition 2. *Assume that the assumptions of Theorem 1 hold. Consider the test $H_0 : r(\theta_0) = 0$ versus $H_1 : r(\theta_0) \neq 0$ where the function $r(\cdot)$ is continuously differentiable. Let $R(\theta) = \partial r(\theta) / \partial \theta'$. Then, under H_0 , as $T \rightarrow \infty$*

$$W_T = Tr'(\hat{\theta}) [R(\hat{\theta}) \tilde{\Sigma}_T(\hat{\theta}) R'(\hat{\theta})]^{-1} r(\hat{\theta}) \xrightarrow{d} \chi_r^2,$$

where $\tilde{\Sigma}_T(\hat{\theta}) = \tilde{H}_T^{-1}(\hat{\theta}) \tilde{I}_T(\hat{\theta}) \tilde{H}_T^{-1}(\hat{\theta})$.

The result follows immediately by the multivariate delta method, the continuous mapping theorem and standard asymptotic convergence arguments. Proposition 2 provides us a testing procedure for $H_0 : \theta_0 \in \Theta_R$ versus $H_1 : \theta_0 \notin \Theta_R$. It is worth nothing that the hypothesis test depicted in Proposition 2 does not require the variance of the model to be correctly specified. In the special case in which the pseudo-variance is correctly specified, then the test can be interpreted as a test of correct specification.

For example, consider the INAR(1) model in (7) with conditional mean and pseudo-variance equations as defined in equation (21). We may consider the following test

$$H_0 : b = a(1 - a) \quad \text{vs} \quad H_1 : b \neq a(1 - a), \quad (22)$$

which is a test for the assumption of a binomial thinning operator ‘ \circ ’. This follows from the definition of the INAR model in (7) as the autoregressive coefficient of the variance takes the form $b = a(1 - a)$ under the assumption of binomial thinning. Alternative thinning specifications can be tested leading to a different form of the autoregressive variance parameter b , see Latour (1998) for the properties of INAR models with a general thinning specification. For instance, if we have a Poisson distribution for the thinning operator we have the restriction $b = a$. The corresponding test is

$$H_0 : b = a \quad \text{vs} \quad H_1 : b \neq a, \quad (23)$$

Table 2: Empirical size for test in (23). The model considered under H_0 is an INAR(1) model with Poisson thinning as well as Poisson error with parameter values $a = 0.5$ and $\omega = 2$.

Nominal size	T			
	250	500	1000	2000
0.1000	0.1224	0.1170	0.1058	0.1094
0.0500	0.0690	0.0662	0.0542	0.0582
0.0100	0.0218	0.0182	0.0110	0.0138




Figure 3: Empirical power for test in (23). The true parameter values of the INAR(1) model with negative binomial thinning and Poisson error are $a = 0.5$ and $\omega = 2$. The value of the dispersion parameter v changes as indicated in the horizontal axis through the % of overdispersion: $1 - a/(a + a^2/v)$.

which assesses the validity of the assumption of equidispersion in the thinning operator versus either overdispersion or underdispersion.

We carry out a simulation study with 5000 Monte Carlo replications to assess the empirical size and power of the test of the parameter restrictions for the INAR(1) model. We consider the hypothesis in (23). To assess the size of the test we simulate under H_0 from a model with Poisson thinning operator and a Poisson distribution of the error term. Table 2 reports the results on the empirical size of the test. We can see that the test is slightly oversized for the smallest sample size, though still close to the nominal level, and it quickly becomes correctly sized as the sample size increases. Next, we evaluate the power of the test by simulating under the alternative. We consider a negative binomial thinning specification such that $a \circ X$ has a negative binomial distribution with mean aX and variance bX , $b = a + a^2/v$, where v is the dispersion parameter of the negative binomial. We note that this generates overdispersion in the thinning as $b = a + a^2/v > a$ and

the smaller the parameter v the more the overdispersion. Figure 3 shows the power of the test in (23) to reject the null hypothesis. As expected, we see that the power increases as the relative overdispersion $1 - a/b$ increases (v decreases) and as the sample size increases. Overall, the results show how the test has appropriate size and it has power against alternative hypotheses.

5 Real data applications

In this section, we present two empirical applications where we employ PVQMLEs. We consider the test described in Section 4 to select appropriate parameter restrictions and compare different PVQMLEs. The first application concerns a dataset of crime counts, where the INAR model is considered for the specification of the conditional mean and the pseudo-variance. The second application concerns the realized correlation between two financial assets that forms a double-bounded time series, where we consider a beta autoregression for the specification of the conditional mean and the pseudo-variance.

5.1 INAR model for crime counts

We consider an empirical application to the monthly number of offensive conduct reports in the city of Blacktown, Australia, from January 1995 to December 2014. This dataset has been employed in several articles featuring the INAR(1) model (Gorgi, 2018; Leisen et al., 2019). The time series is displayed in Figure 4. In the literature, the distributional structure of the INAR innovation term ε_t is typically allowed to be flexible or left unspecified but the thinning operator is typically considered to be binomial. We consider the test proposed in the previous section to formally test the validity of binomial thinning assumption as well as the dispersion of the error term. We obtain the unrestricted PVQMLE for the INAR conditional mean and pseudo-variance equations in (21) and test several restrictions based on the test in Proposition 2. We test for equidispersion in the error $H_0 : \omega_1 = \omega_2$, binomial thinning $H_0 : b = a(1 - a)$, Poisson thinning $H_0 : a = b$ and geometric thinning $H_0 : b = a + a^2$.

Table 3: p -values of the restriction tests for the INAR(1) model.

$\omega_1 = \omega_2$	Thinning		
	binomial	Poisson	geometric
0.372	0.005	0.043	0.229




Figure 4: Monthly number of offensive conduct reports in Blacktown, Australia, from January 1995 to December 2014.

Table 4: PVQMLEs of the INAR(1) model for the crime time series dataset. Standard errors in brackets.

	$\hat{\omega}_1$	$\hat{\omega}_2$	\hat{a}	\hat{b}
Unrestricted	4.559 (0.520)	6.644 (2.374)	0.509 (0.058)	1.170 (0.330)
binomial thinning	6.280 (0.434)	-	0.371 (0.040)	-
Poisson thinning	4.820 (0.523)	-	0.524 (0.058)	-
geometric thinning	4.129 (0.500)	-	0.592 (0.059)	-

The results of the tests are summarized in Table 3. We can see that the test does not reject the hypothesis of equidispersion in the error $\omega_1 = \omega_2$. As it concerns the tests on the thinning, the binomial and Poisson thinning are rejected at 5% significance level, instead, the geometric thinning is not rejected. This indicates that there is overdispersion in the thinning and the geometric one may be appropriate to describe the degree of overdispersion. Table 4 reports the estimation results for several PVQMLEs that are based on the different restrictions on the thinning operator. We can see that restricting to a binomial thinning leads to substantially biased estimates with respect to the unrestricted PVQMLE. Instead, from the geometric thinning we do not have such bias and the estimator can be expected to have an higher efficiency.

5.2 Double-bounded autoregression for realized correlation

The second application we present concerns the modelling of daily realized correlations between Boeing and Honeywell stocks as considered in [Gorgi and Koopman \(2021\)](#). Figure 5 reports the plot of the time series. The sample size is $T = 2515$. Realized correlation measures take values in the interval $[-1, 1]$ and the transformation $Y_t/2 + 1/2$ is applied to rescale the realized correlation in the unit interval $[0, 1]$. We refer to [Gorgi and Koopman \(2021\)](#) for a discussion on how models on the unit interval can be extended to a general interval with known bounds.




Figure 5: Daily time series of realized correlations between Boeing (BA) and Honeywell (HON) asset returns, from January 2001 to December 2010.

We consider the following specification for the conditional mean and pseudo-variance

$$\begin{aligned} \lambda_t &= \delta_1 + \alpha_1 Y_{t-1} + \beta_1 \lambda_{t-1}, \\ \nu_t^* &= \frac{\mu_t(1 - \mu_t)}{1 + \phi}, \quad \mu_t = \delta_2 + \alpha_2 Y_{t-1} + \beta_2 \mu_{t-1}, \end{aligned}$$

where the double-bounded nature of the data requires $\delta_i + \alpha_i + \beta_i < 1$ for $i = 1, 2$ and $\phi > 0$. Besides the unrestricted PVQMLE, we consider a restricted PVQMLE with $\delta_2 = \delta_1$, $\alpha_2 = \alpha_1$, $\beta_1 = \beta_2$, which implies $\mu_t = \lambda_t$. These restrictions impose that the pseudo-variance is equal to the conditional variance implied by a beta distribution with mean parameter λ_t and precision parameter ϕ (see [Example 4](#)). In this way, we can also test the adequacy of the beta autoregression for modelling the analyzed data through the specification test on the restriction. [Table 5](#) reports the estimation results together with the restriction tests. We can see that the specification test rejects the null hypothesis of equality for the estimated α coefficients. For the same reason also the null assumption of the combined joint test is rejected. However, the null hypothesis is instead not

rejected for δ and β coefficients at 1% level. This leans in favour of the restricted PVQMLE. We also notice that the estimated coefficients and the corresponding standard errors of the restricted PVQMLE are fairly close to the ones obtained from the beta autoregression reported in Table 1 of [Gorgi and Koopman \(2021\)](#).

Table 5: Estimation results for the realized correlation series. Standard errors in brackets. The bottom of the table reports the p -values of the tests on the restrictions.

	$\hat{\delta}_1$	$\hat{\alpha}_1$	$\hat{\beta}_1$	$\hat{\phi}$	$\hat{\delta}_2$	$\hat{\alpha}_2$	$\hat{\beta}_2$
Unrestricted	0.01 (0.003)	0.163 (0.013)	0.822 (0.015)	22.226 (2.745)	0.055 (0.019)	0.045 (0.007)	0.898 (0.022)
Restricted	0.01 (0.003)	0.161 (0.013)	0.826 (0.015)	36.963 (1.073)	-	-	-
H_0	$\delta_1 = \delta_2$	$\alpha_1 = \alpha_2$	$\beta_1 = \beta_2$	joint test			
p -value	0.02	<0.001	0.01	<0.001			

Appendix A: Proofs of Results

A.1 Proofs

Proof of Theorem 1: Let $L(\theta) = \mathbb{E}[l_t(\theta)]$ be the limit log-quasi-likelihood. In what follows we show the following intermediate results.

- (i) Uniform convergence: $\sup_{\theta \in \Theta} |\tilde{L}_T(\theta) - L(\theta)| \rightarrow 0$ almost surely, as $T \rightarrow \infty$.
- (ii) Identifiability: the true parameter value θ_0 is the unique maximizer of $L(\theta)$, i.e. $\mathbb{E}[l_t(\theta)] < \mathbb{E}[l_t(\theta_0)]$ for all $\theta \in \Theta, \theta \neq \theta_0$.

In order to prove (i) the uniform convergence of the two summands of (A.1) should be shown.

$$|\tilde{L}_T(\theta) - L(\theta)| \leq |\tilde{L}_T(\theta) - L_T(\theta)| + |L_T(\theta) - L(\theta)|. \quad (\text{A.1})$$

The first term converges uniformly by Lemma 1, under [A4-A5](#), implying that the starting value of the process is asymptotically unimportant for the quasi-likelihood contribution. By assumption [A1](#) the log-quasi-likelihood contribution $l_t(\theta)$ is stationary and ergodic. Moreover, it is uniformly bounded

$$\mathbb{E} \sup_{\theta \in \Theta} |l_t(\theta)| \leq \frac{1}{2} \mathbb{E} \sup_{\gamma \in \Gamma} |\log \nu_t^*(\gamma)| + \frac{1}{2} \mathbb{E} \sup_{\theta \in \Theta} \left(\frac{[Y_t - \lambda_t(\psi)]^2}{\nu_t^*(\gamma)} \right) < \infty$$

by assumption **A2**. For the continuity of the quasi-likelihood and the compactness of Θ , [Straumann and Mikosch \(2006, Thm. 2.7\)](#) applies providing the uniform convergence of the second term in [\(A.1\)](#); in symbols $\sup_{\theta \in \Theta} |L_T(\theta) - L(\theta)| \rightarrow 0$ almost surely, as $T \rightarrow \infty$. This concludes the proof of [\(i\)](#).

We now prove [\(ii\)](#). First note that by the uniform limit theorem $L(\theta) = E[l_t(\theta)]$ is a continuous function and it attains at least a maximum in Θ since Θ is compact. We now prove that such maximum is unique so that it can be univocally identified. Recall that $\theta = (\psi', \gamma)'$, assumption **A2** provides $E \sup_{\psi \in \Psi} |l_t(\psi, \gamma)| < \infty$ and $E \sup_{\gamma \in \Gamma} |l_t(\psi_0, \gamma)| < \infty$ so also the function $l_t(\psi, \gamma)$ has at least a maximum for $\psi \in \Psi$, and $l_t(\psi_0, \gamma)$ has at least a maximum for $\gamma \in \Gamma$. Consider now $E \{l_t(\theta) - l_t(\theta_0)\} = E \{l_t(\theta) - l_t(\psi_0, \gamma)\} + E \{l_t(\psi_0, \gamma) - l_t(\theta_0)\}$.

$$\begin{aligned} E \{l_t(\theta) - l_t(\psi_0, \gamma)\} &= E \left\{ -\frac{E[(Y_t - \lambda_t(\psi))^2 | \mathcal{F}_{t-1}]}{2\nu_t^*(\gamma)} + \frac{\nu_t}{2\nu_t^*(\gamma)} \right\} \\ &\leq E \left\{ -\frac{\nu_t}{2\nu_t^*(\gamma)} + \frac{\nu_t}{2\nu_t^*(\gamma)} \right\} = 0 \end{aligned}$$

with equality if and only if $\psi = \psi_0$ by assumption **A3**. Moreover, $E \{l_t(\psi_0, \gamma) - l_t(\theta_0)\} = E [l_t(\psi_0, \gamma)] - E [l_t(\psi_0, \gamma^*)] \leq 0$ by assumption **A6**. This concludes the proof of [\(ii\)](#). The consistency of the whole estimator $\hat{\theta}$ follows from [\(i\)](#), [\(ii\)](#) and the compactness of Θ by [Pötscher and Prucha \(1997, Lemma 3.1\)](#). This implies [\(11\)](#).

To prove the asymptotic normality of the estimator we establish additional intermediate results.

- (a) $\sqrt{T} \sup_{\theta \in \Theta} \|S_T(\theta) - \tilde{S}_T(\theta)\| \rightarrow 0$ almost surely, as $T \rightarrow \infty$.
- (b) Define $V_T(\theta) = T^{-1} \sum_{t=1}^T -\partial^2 l_t(\theta) / \partial \theta \partial \theta'$. $V_T(\theta) \rightarrow H(\theta_0)$ almost surely uniformly over $\theta \in \Theta$, as $T \rightarrow \infty$.
- (c) $E[s_t(\theta_0)] = 0$.

The condition [\(a\)](#) is satisfied by [Lemma 2](#), under [A4-A5](#) and [A7](#) implying that initial values of the process do not affect the asymptotic distribution of the PVQMLE.

Consider the second derivative of the log-quasi-likelihood contribution.

$$\begin{aligned} \frac{\partial^2 l_t(\theta)}{\partial \theta \partial \theta'} &= \left(\frac{1}{2\nu_t^{*2}(\gamma)} - \frac{[Y_t - \lambda_t(\psi)]^2}{\nu_t^{*3}(\gamma)} \right) \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \frac{\partial \nu_t^*(\gamma)}{\partial \theta'} \\ &\quad - \frac{Y_t - \lambda_t(\psi)}{\nu_t^{*2}(\gamma)} \left(\frac{\partial \lambda_t(\psi)}{\partial \theta} \frac{\partial \nu_t^*(\gamma)}{\partial \theta'} - \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \frac{\partial \lambda_t(\psi)}{\partial \theta'} \right) \\ &\quad - \frac{1}{\nu_t^*(\gamma)} \frac{\partial \lambda_t(\psi)}{\partial \theta} \frac{\partial \lambda_t(\psi)}{\partial \theta'} + \frac{Y_t - \lambda_t(\psi)}{\nu_t^*(\gamma)} \frac{\partial^2 \lambda_t(\psi)}{\partial \theta \partial \theta'} \\ &\quad + \left(\frac{[Y_t - \lambda_t(\psi)]^2}{2\nu_t^{*2}(\gamma)} - \frac{1}{2\nu_t^*(\gamma)} \right) \frac{\partial^2 \nu_t^*(\gamma)}{\partial \theta \partial \theta'}. \end{aligned}$$

Assumption **A8** and the Cauchy-Schwarz inequality yield $E \sup_{\theta \in \Theta} |\partial^2 l_t(\theta) / \partial \theta_i \partial \theta_j| < \infty$ for all $i, j = 1, \dots, m$. Furthermore, the second derivative is a continuous, stationary and ergodic sequence. Then, an application of [Straumann and Mikosch \(2006, Thm. 2.7\)](#) provides the condition **(b)**. Note that since in this case $\partial \lambda_t(\psi) / \partial \gamma = \partial \nu_t^*(\gamma) / \partial \psi = 0$ the matrix $H(\theta_0)$ is block diagonal with diagonal block matrices $H_\psi(\theta_0) = E[-\partial^2 l_t(\theta_0) / \partial \psi \partial \psi']$ and $H_\gamma(\theta_0) = E[-\partial^2 l_t(\theta_0) / \partial \gamma \partial \gamma']$. The former is defined in [\(13\)](#).

For establishing the asymptotic normality of the estimator $\hat{\theta}$ the proof of **(c)** is needed. Let $s_t(\theta_0) = [s_t^{(\psi)}(\theta_0)', s_t^{(\gamma)}(\theta_0)']'$ be the partition of the score between mean and pseudo-variance parameters. Observe that $E(s_t^{(\psi)}(\theta_0) | \mathcal{F}_{t-1}) = 0$ but $E(s_t(\theta_0) | \mathcal{F}_{t-1}) \neq 0$. Note that $\sup_{\theta \in \Theta} |\partial l_t(\theta) / \partial \theta_i| \leq 2 [\sup_{\theta \in \Theta} |l_t(\theta)|]^{1/2} [\sup_{\theta \in \Theta} |\partial^2 l_t(\theta) / \partial \theta_i \partial \theta_i|]^{1/2}$, by [Rudin \(1976, p. 115\)](#). Moreover, $E \sup_{\theta \in \Theta} |l_t(\theta)| < \infty$, and $E \sup_{\theta \in \Theta} |\partial^2 l_t(\theta) / \partial \theta_i \partial \theta_j| < \infty$. Then an application of Cauchy-Schwarz inequality entails $E \sup_{\theta \in \Theta} |\partial l_t(\theta) / \partial \theta_i| < \infty$. Finally, $\|\partial l_t(\theta) / \partial \theta\| \leq \sup_{\theta \in \Theta} \|\partial l_t(\theta) / \partial \theta\|$ and an application of the dominated convergence theorem leads to $E[\partial l_t(\theta) / \partial \theta] = \partial E[l_t(\theta)] / \partial \theta$. By noting that θ_0 is the unique maximizer of $E[l_t(\theta)]$ the result **(c)** follows.

For T large enough $\hat{\theta} \in \dot{\Theta}$ by [A10](#), so the following derivatives exist almost surely

$$0 = \sqrt{T} \tilde{S}_T(\hat{\theta}) = \sqrt{T} S_T(\hat{\theta}) + o(1) = \sqrt{T} S_T(\theta_0) - V_T(\bar{\theta}) \sqrt{T}(\hat{\theta} - \theta_0) + o(1),$$

where the first equality comes from the definition [\(4\)](#), the second equality holds by [\(a\)](#), and the third equality is obtained by Taylor expansion at θ_0 with $\bar{\theta}$ lying between $\hat{\theta}$ and θ_0 . By assumption [A11](#) and **(c)** we have $\sqrt{T} S_T(\theta_0) \xrightarrow{d} N(0, I(\theta_0))$ with $I(\theta_0) = E[s_t(\theta_0) s_t(\theta_0)']$. This fact and **(b)** establish the asymptotic normality of the estimator $\hat{\theta}$ with covariance matrix $\Sigma(\theta_0) = H^{-1}(\theta_0) I(\theta_0) H^{-1}(\theta_0)$ by assumption [A9](#), where

$$H(\theta_0) = \begin{pmatrix} H_\psi(\theta_0) & 0 \\ 0 & H_\gamma(\theta_0) \end{pmatrix}, \quad I(\theta_0) = \begin{pmatrix} I_\psi(\theta_0) & I_{\psi,\gamma}(\theta_0) \\ I_{\psi,\gamma}(\theta_0)' & I_\gamma(\theta_0) \end{pmatrix}, \quad (\text{A.2})$$

with $H_x(\theta_0) = E[-\partial^2 l_t(\theta_0) / \partial x \partial x']$, $I_x(\theta_0) = E[s_t^{(x)}(\theta_0) s_t^{(x)}(\theta_0)']$ and $I_{x,z}(\theta_0) = E[s_t^{(x)}(\theta_0) s_t^{(z)}(\theta_0)']$. In particular, standard algebra shows that $I_\psi(\theta_0)$ equals [\(13\)](#). See also equation [\(18\)](#). A suitable block matrix multiplication of [\(A.2\)](#) provides

$$\Sigma(\theta_0) = \begin{pmatrix} \Sigma_\psi(\theta_0) & \Sigma_{\psi,\gamma}(\theta_0) \\ \Sigma_{\psi,\gamma}(\theta_0)' & \Sigma_\gamma(\theta_0) \end{pmatrix},$$

where $\Sigma_\psi(\theta_0)$ takes the form defined in [\(12\)](#). Finally, note that for the marginal property of the multivariate Gaussian distribution result [\(12\)](#) holds with covariance matrix Σ_ψ being the partition of $\Sigma(\theta_0)$ for the mean parameters ψ . \square

Proof of Corollary 2: Condition [A11](#) is not required since in this case is easily showed by [\(10\)](#) that $E(s_t(\theta_0)|\mathcal{F}_{t-1}) = 0$. Recall that $\sqrt{T}s_T(\theta_0) = T^{-1/2} \sum_{t=1}^T U_t$ where $U_t = s_t(\theta_0)$. Note that $\{U_t, \mathcal{F}_t\}$ is a stationary martingale difference, and due to [A9](#) it has finite second moments. Then [A11](#) follows by the central limit theorem for martingales ([Billingsley, 1961](#)) and the Cramér-Wold device. The consistency and asymptotic normality of $\hat{\theta}$ follow as above. Finally, in view of [\(18\)](#) and $E(s_t^{(\psi)}(\theta_0)|\mathcal{F}_{t-1}) = 0$

$$\text{Var}[H_\psi^{-1}(\theta_0)s_t^{(\psi)}(\theta_0) - I_\psi^{-1}(\theta_0)s_t^{(\psi)}(\theta_0)] = \Sigma_\psi - I_\psi$$

being necessarily positive semi-definite. □

Proof of Corollary 3: Analogously to the proof of [Theorem 1](#), [A1-A5](#) guarantee that $L_t(\theta)$ is continuous and a.s. uniformly convergent to $E[l_t(\theta)]$. By recalling that Θ is compact the result follows by [Pötscher and Prucha \(1997, Lemma 4.2\)](#). □

Proof of Theorem 2: The consistency of $\hat{\theta}_R$ follows from the fact that by the proof of [Theorem 1](#) we have that $E[l_t(\psi, \gamma)] \leq E[l_t(\psi_0, \gamma^*)]$ for any $\theta \in \Theta$ with equality holding only if $\theta = (\psi_0', \gamma^{*'})'$, and assumption L ensures that $(\psi_0', \gamma^{*'})' \in \Theta_R$ with $\Theta_R \subseteq \Theta$. The consistency in [\(15\)](#) follows. The asymptotic normality of the estimator $\hat{\theta}_R$ follows as in the proof of [Theorem 1](#) with covariance matrix $\Sigma(\theta_0) = H^{-1}(\theta_0)I(\theta_0)H^{-1}(\theta_0)$. In this case Hessian and Fisher information matrices can be written in the following block matrix form

$$H(\theta_0) = \begin{pmatrix} H_\psi(\theta_0) & H_{\psi, \gamma_2}(\theta_0) \\ H_{\psi, \gamma_2}(\theta_0)' & H_{\gamma_2}(\theta_0) \end{pmatrix}, \quad I(\theta_0) = \begin{pmatrix} I_\psi(\theta_0) & I_{\psi, \gamma_2}(\theta_0) \\ I_{\psi, \gamma_2}(\theta_0)' & I_{\gamma_2}(\theta_0) \end{pmatrix}. \quad (\text{A.3})$$

Moreover, recall that

$$H^{-1}(\theta_0) = D(\theta_0) = \begin{pmatrix} D_\psi(\theta_0) & D_{\psi, \gamma_2}(\theta_0) \\ D_{\psi, \gamma_2}(\theta_0)' & D_{\gamma_2}(\theta_0) \end{pmatrix}. \quad (\text{A.4})$$

By computing $\Sigma(\theta_0)$ using the block matrix multiplication as defined in [\(A.3\)](#) and [\(A.4\)](#) the partition of $\Sigma(\theta_0)$ for the mean parameters ψ equals Σ_R . This entails [\(16\)](#). □

Proof of Proposition 1: It is not hard to show that under the conditions of [Proposition 1](#) the Hessian and information matrices of [\(4\)](#) take the form defined in [\(A.5\)](#) and [\(A.6\)](#), respectively.

$$H(\theta_0) = E \left[\frac{1}{\nu_t(\gamma_0)} \frac{\partial \lambda_t(\psi_0)}{\partial \theta} \frac{\partial \lambda_t(\psi_0)}{\partial \theta'} + \frac{1}{2\nu_t^2(\gamma_0)} \frac{\partial \nu_t(\gamma_0)}{\partial \theta} \frac{\partial \nu_t(\gamma_0)}{\partial \theta'} \right], \quad (\text{A.5})$$

$$\begin{aligned}
I(\theta_0) = & \mathbb{E} \left[\frac{1}{\nu_t(\gamma_0)} \frac{\partial \lambda_t(\psi_0)}{\partial \theta} \frac{\partial \lambda_t(\psi_0)}{\partial \theta'} + \frac{h_t}{2\nu_t^3(\gamma_0)} \left(\frac{\partial \lambda_t(\psi_0)}{\partial \theta} \frac{\partial \nu_t(\gamma_0)}{\partial \theta'} + \frac{\partial \nu_t(\gamma_0)}{\partial \theta} \frac{\partial \lambda_t(\psi_0)}{\partial \theta'} \right) \right] \\
& + \mathbb{E} \left[\left(\frac{k_t}{\nu_t^2(\gamma_0)} - 1 \right) \frac{1}{4\nu_t^2(\gamma_0)} \frac{\partial \nu_t(\gamma_0)}{\partial \theta} \frac{\partial \nu_t(\gamma_0)}{\partial \theta'} \right], \tag{A.6}
\end{aligned}$$

where $h_t = \mathbb{E}[(Y_t - \lambda_t(\psi_0))^3 | \mathcal{F}_{t-1}]$ and $k_t = \mathbb{E}[(Y_t - \lambda_t(\psi_0))^4 | \mathcal{F}_{t-1}]$. From (A.5) and (A.6), we can see that when the variance is correctly specified the PVQMLE can be considered as a Gaussian QMLE. In case the data are normally distributed the PVQMLE eventually becomes MLE since $h_t = 0$ and $k_t = 3\nu_t^2$ implying $I(\theta_0) = H(\theta_0)$. So the Gaussian case in assumption A13 is trivial. For the other cases we have $m = 1$, so $\psi = \theta$ and $I(\theta_0) \leq H(\theta_0)$, then $\Sigma_R = I(\theta_0)/H(\theta_0)^2 \leq 1/I(\theta_0)$. \square

A.2 Technical lemmas

Lemma 1. *Consider the PVQMLE in (5) with log-quasi-likelihood (4). Under conditions A4-A5, almost surely as $T \rightarrow \infty$, $\sup_{\theta \in \Theta} |\tilde{L}_T(\theta) - L_T(\theta)| \rightarrow 0$.*

Proof of Lemma 1: From assumption A4, we have that

$$\begin{aligned}
& \sup_{\theta \in \Theta} |l_t(\theta) - \tilde{l}_t(\theta)| \\
\leq & \sup_{\theta \in \Theta} \left| \frac{[\tilde{\lambda}_t(\psi) - \lambda_t(\psi)][\tilde{\lambda}_t(\psi) + \lambda_t(\psi) - 2Y_t]}{2\tilde{\nu}_t^*(\gamma)} + \frac{[\nu_t^*(\gamma) - \tilde{\nu}_t^*(\gamma)][Y_t - \lambda_t(\psi)]^2}{2\nu_t^*(\gamma)\tilde{\nu}_t^*(\gamma)} \right| \\
& + \frac{1}{2} \sup_{\gamma \in \Gamma} \left| \log \frac{\tilde{\nu}_t^*(\gamma)}{\nu_t^*(\gamma)} \right| \\
\leq & \frac{1}{\underline{\nu}^*} a_t \left(a_t + |Y_t| + \sup_{\psi \in \Psi} |\lambda_t(\psi)| \right) + \frac{1}{\underline{\nu}^{*2}} b_t \left(Y_t^2 + \sup_{\psi \in \Psi} \lambda_t^2(\psi) \right) \\
& + \frac{1}{2} \sup_{\gamma \in \Gamma} \left| \log \left(1 + \frac{\tilde{\nu}_t^*(\gamma) - \nu_t^*(\gamma)}{\nu_t^*(\gamma)} \right) \right| \\
\leq & \frac{1}{\underline{\nu}^*} a_t \left(1 + |Y_t| + \sup_{\psi \in \Psi} |\lambda_t(\psi)| \right) + \frac{1}{\underline{\nu}^{*2}} b_t \left(Y_t^2 + \sup_{\psi \in \Psi} \lambda_t^2(\psi) \right) + \frac{1}{2\underline{\nu}^*} b_t,
\end{aligned}$$

for t large enough since, by assumption A5, a.s. $a_t \rightarrow 0$ as $t \rightarrow \infty$ and by noticing that $\log(1+x) \leq x$ for $x > -1$. Assumption A5 and an application of Cesaro's lemma lead to

$$\sup_{\theta \in \Theta} |\tilde{L}_T(\theta) - L_T(\theta)| \leq T^{-1} \sum_{t=1}^T \sup_{\theta \in \Theta} |\tilde{l}_t(\theta) - l_t(\theta)| \rightarrow 0, \quad a.s.$$

as $T \rightarrow \infty$. \square

Lemma 2. Consider the PVQMLE in (5) with score (10). Under conditions A4-A5 and A7, almost surely as $T \rightarrow \infty$, $\sqrt{T} \sup_{\theta \in \Theta} \|\tilde{S}_T(\theta) - S_T(\theta)\| \rightarrow 0$.

Proof of Lemma 2: We obtain that

$$\begin{aligned} \sup_{\theta \in \Theta} \|s_t(\theta) - \tilde{s}_t(\theta)\| &\leq \sup_{\theta \in \Theta} \left\| \frac{1}{2\tilde{\nu}_t^*(\gamma)} \frac{\partial \tilde{\nu}_t^*(\gamma)}{\partial \theta} - \frac{1}{2\nu_t^*(\gamma)} \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \right\| \\ &\quad + \sup_{\theta \in \Theta} \left\| \frac{Y_t - \tilde{\lambda}_t(\psi)}{\tilde{\nu}_t^*(\gamma)} \frac{\partial \tilde{\lambda}_t(\psi)}{\partial \theta} - \frac{Y_t - \lambda_t(\psi)}{\nu_t^*(\gamma)} \frac{\partial \lambda_t(\psi)}{\partial \theta} \right\| \\ &\quad + \sup_{\theta \in \Theta} \left\| \frac{[Y_t - \tilde{\lambda}_t(\psi)]^2}{2\tilde{\nu}_t^{*2}(\gamma)} \frac{\partial \tilde{\nu}_t^*(\gamma)}{\partial \theta} - \frac{[Y_t - \lambda_t(\psi)]^2}{2\nu_t^{*2}(\gamma)} \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \right\| = \delta_t^1 + \delta_t^2 + \delta_t^3, \end{aligned}$$

with obvious notation. We now bound the single terms individually. In what follows the notation $o(1)$ almost surely, as $t \rightarrow \infty$, will be abbreviated to $o(1)$.

$$\begin{aligned} \delta_t^1 &\leq \sup_{\theta \in \Theta} \left\| \frac{1}{2\tilde{\nu}_t^*(\gamma)} \left(\frac{\partial \tilde{\nu}_t^*(\gamma)}{\partial \theta} - \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \right) + \frac{[\nu_t^*(\gamma) - \tilde{\nu}_t^*(\gamma)]}{2\tilde{\nu}_t^*(\gamma)\nu_t^*(\gamma)} \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \right\| \\ &\leq \frac{d_t}{2\underline{\nu}^*} + \frac{b_t}{2\underline{\nu}^{*2}} \sup_{\theta \in \Theta} \left\| \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \right\|. \end{aligned}$$

Similarly,

$$\begin{aligned} \delta_t^2 &\leq \sup_{\theta \in \Theta} \left\| \frac{Y_t - \tilde{\lambda}_t(\psi)}{\tilde{\nu}_t^*(\gamma)} \left(\frac{\partial \tilde{\lambda}_t(\psi)}{\partial \theta} - \frac{\partial \lambda_t(\psi)}{\partial \theta} \right) \right\| \\ &\quad + \sup_{\theta \in \Theta} \left\| \frac{\partial \lambda_t(\psi)}{\partial \theta} \left(\frac{\lambda_t(\psi) - \tilde{\lambda}_t(\psi)}{\tilde{\nu}_t^*(\gamma)} + \frac{Y_t - \lambda_t(\psi)}{\tilde{\nu}_t^*(\gamma)} - \frac{Y_t - \lambda_t(\psi)}{\nu_t^*(\gamma)} \right) \right\| \\ &\leq \frac{c_t}{\underline{\nu}^*} \left(|Y_t| + \sup_{\psi \in \Psi} |\lambda_t(\psi)| + a_t \right) \\ &\quad + \sup_{\theta \in \Theta} \left\| \frac{\partial \lambda_t(\psi)}{\partial \theta} \right\| \left(\frac{a_t}{\underline{\nu}^*} + \sup_{\theta \in \Theta} \left| \frac{[\nu_t^*(\gamma) - \tilde{\nu}_t^*(\gamma)] [Y_t - \lambda_t(\psi)]}{\tilde{\nu}_t^*(\gamma)\nu_t^*(\gamma)} \right| \right) \\ &\leq \frac{c_t}{\underline{\nu}^*} \left(|Y_t| + \sup_{\psi \in \Psi} |\lambda_t(\psi)| + o(1) \right) + \sup_{\theta \in \Theta} \left\| \frac{\partial \lambda_t(\psi)}{\partial \theta} \right\| \left(\frac{a_t}{\underline{\nu}^*} + \frac{b_t}{\underline{\nu}^{*2}} \left(|Y_t| + \sup_{\psi \in \Psi} |\lambda_t(\psi)| \right) \right). \end{aligned}$$

Using similar arguments for δ_t^3 and assumption **A5** leads to

$$\begin{aligned}
\delta_t^3 &\leq \frac{d_t}{\underline{\nu}^{*2}} \sup_{\theta \in \Theta} (Y_t^2 + \lambda_t^2(\psi) + a_t^2 + 2a_t \lambda_t(\psi)) \\
&\quad + \sup_{\theta \in \Theta} \left\| \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \right\| \sup_{\theta \in \Theta} \left| \frac{[\tilde{\lambda}_t(\psi) - \lambda_t(\psi)][\tilde{\lambda}_t(\psi) + \lambda_t(\psi) - 2Y_t]}{2\tilde{\nu}_t^{*2}(\gamma)} \right| \\
&\quad + \sup_{\theta \in \Theta} \left\| \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \right\| \sup_{\theta \in \Theta} \left| \frac{[\nu_t^*(\gamma) - \tilde{\nu}_t^*(\gamma)][\nu_t^*(\gamma) + \tilde{\nu}_t^*(\gamma)][Y_t - \lambda_t(\psi)]^2}{2\nu_t^{*2}(\gamma)\tilde{\nu}_t^{*2}(\gamma)} \right| \\
&\leq \frac{d_t}{\underline{\nu}^{*2}} \left(Y_t^2 + \sup_{\psi \in \Psi} \lambda_t^2(\psi) + o(1) \right) + \sup_{\theta \in \Theta} \left\| \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \right\| \frac{a_t}{\underline{\nu}^{*2}} \left(|Y_t| + \sup_{\psi \in \Psi} |\lambda_t(\psi)| + o(1) \right) \\
&\quad + \sup_{\theta \in \Theta} \left\| \frac{\partial \nu_t^*(\gamma)}{\partial \theta} \right\| \frac{2b_t}{\underline{\nu}^{*3}} \left(Y_t^2 + \sup_{\psi \in \Psi} \lambda_t^2(\psi) \right).
\end{aligned}$$

By assumption **A7**, $\delta_t^j = \mathcal{O}(t^{-\delta})$, for $\delta > 1/2$ and $j = 1, 2, 3$. Therefore $\sqrt{T} \sup_{\theta \in \Theta} \|S_T(\theta) - \tilde{S}_T(\theta)\| \leq T^{-1/2} \sum_{t=1}^T \mathcal{O}(t^{-\delta})$ converges to 0 almost surely as $T \rightarrow \infty$. \square

Funding

Mirko Armillotta acknowledges financial support from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No.101108797.

References

- Ahmad, A. and Francq, C. (2016). Poisson QMLE of count time series models. *Journal of Time Series Analysis*, 37:291–314.
- Aknouche, A. and Francq, C. (2021). Two-stage weighted least squares estimator of the conditional mean of observation-driven time series models. *Journal of Econometrics*, To appear:To appear.
- Al-Osh, M. and Alzaid, A. A. (1987). First-order integer-valued autoregressive (INAR (1)) process. *Journal of Time Series Analysis*, 8:261–275.
- Aleksandrov, B. and Weiß, C. H. (2020). Testing the dispersion structure of count time series using pearson residuals. *ASTA Advances in Statistical Analysis*, 104(3):325–361.
- Billingsley, P. (1961). The Lindeberg-Lévy theorem for martingales. *Proceedings of the American Mathematical Society*, 12:788–792.

- Box, G. E., Jenkins, G. M., and Reinsel, G. (1970). Time series analysis: forecasting and control. *Holden-day San Francisco*.
- Bradley, R. C. (2005). Basic properties of strong mixing conditions. A survey and some open questions. *Probability surveys*, 2:107–144.
- Bu, R., McCabe, B., and Hadri, K. (2008). Maximum likelihood estimation of higher-order integer-valued autoregressive processes. *Journal of Time Series Analysis*, 29:973–994.
- Christou, V. and Fokianos, K. (2015). Estimation and testing linearity for non-linear mixed Poisson autoregressions. *Electron. J. Stat.*, 9:1357–1377.
- Creal, D., Koopman, S. J., and Lucas, A. (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, 28(5):777–795.
- Daniels, H. E. (1954). Saddlepoint approximations in statistics. *The Annals of Mathematical Statistics*, 25:631–650.
- Debaly, Z. M. and Truquet, L. (2021). A note on the stability of multivariate non-linear time series with an application to time series of counts. *Statistics & Probability Letters*, 179:1–7.
- Doukhan, P. (1994). *Mixing*, volume 85 of *Lecture Notes in Statistics*. Springer-Verlag, New York.
- Drost, F. C., Akker, R. v. d., and Werker, B. J. (2009). Efficient estimation of auto-regression parameters and innovation distributions for semiparametric integer-valued ar (p) models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(2):467–485.
- Engle, R. F. and Russell, J. R. (1998). Autoregressive conditional duration: a new model for irregularly spaced transaction data. *Econometrica*, 66:1127–1162.
- Ferland, R., Latour, A., and Oraichi, D. (2006). Integer-valued GARCH process. *Journal of Time Series Analysis*, 27:923–942.
- Francq, C. and Zakoian, J.-M. (2023). Optimal estimating function for weak location-scale dynamic models. *Journal of Time Series Analysis*, To appear.
- Goodman, J. (2022). Asymptotic accuracy of the saddlepoint approximation for maximum likelihood estimation. *The Annals of Statistics*, 50(4):2021–2046.
- Gorgi, P. (2018). Integer-valued autoregressive models with survival probability driven by a stochastic recurrence equation. *Journal of Time Series Analysis*, 39(2):150–171.

- Gorgi, P. and Koopman, S. J. (2021). Beta observation-driven models with exogenous regressors: A joint analysis of realized correlation and leverage effects. *Journal of Econometrics*, To appear.
- Gourieroux, C., Monfort, A., and Trognon, A. (1984). Pseudo maximum likelihood methods: Theory. *Econometrica*, pages 681–700.
- Heinen, A. (2003). Modelling time series count data: an autoregressive conditional Poisson model. Technical Report MPRA Paper 8113, University Library of Munich, Germany. available at <http://mpra.ub.uni-muenchen.de/8113/>.
- Latour, A. (1998). Existence and stochastic structure of a non-negative integer-valued autoregressive process. *Journal of Time Series Analysis*, 19(4):439–455.
- Leisen, F., Mena, R. H., Palma, F., and Rossini, L. (2019). On a flexible construction of a negative binomial model. *Statistics & Probability Letters*, 152:1–8.
- Lu, Y. (2021). The predictive distributions of thinning-based count processes. *Scandinavian Journal of Statistics*, 48:42–67.
- McKenzie, E. (1988). Some ARMA models for dependent sequences of Poisson counts. *Advances in Applied Probability*, 20:822–835.
- Pedeli, X., Davison, A. C., and Fokianos, K. (2015). Likelihood estimation for the INAR(p) model by saddlepoint approximation. *Journal of the American Statistical Association*, 110(511):1229–1238.
- Pötscher, B. M. and Prucha, I. R. (1997). *Dynamic Nonlinear Econometric Models*. Springer-Verlag, Berlin. Asymptotic theory.
- Rocha, A. V. and Cribari-Neto, F. (2009). Beta autoregressive moving average models. *Test*, 18(3):529.
- Rudin, W. (1976). *Principles of Mathematical Analysis*, volume 3. McGraw-hill New York.
- Schweer, S. and Weiß, C. H. (2014). Compound poisson inar (1) processes: Stochastic properties and testing for overdispersion. *Computational Statistics & Data Analysis*, 77:267–284.
- Steutel, F. and Van Harn, K. (1979). Discrete analogues of self-decomposability and stability. *The Annals of Probability*, 7(5):893–899.

- Straumann, D. and Mikosch, T. (2006). Quasi-maximum-likelihood estimation in conditionally heteroscedastic time series: A stochastic recurrence equations approach. *The Annals of Statistics*, 34(5):2449–2495.
- Sun, J. and McCabe, B. P. (2013). Score statistics for testing serial dependence in count data. *Journal of Time Series Analysis*, 34(3):315–329.
- Wedderburn, R. W. (1974). Quasi-likelihood functions, generalized linear models, and the gauss-newton method. *Biometrika*, 61:439–447.
- Wei, C. H. (2018). Goodness-of-fit testing of a count time series’ marginal distribution. *Metrika*, 81(6):619–651.
- White, H. (1994). *Estimation, Inference and Specification Analysis*. Cambridge university press.