

Noise contrastive estimation: Asymptotic properties, formal comparison with MC-MLE

Lionel Riou-Durand

e-mail: lionel.riou-durand@ensae.fr

and

Nicolas Chopin

e-mail: nicolas.chopin@ensae.fr

ENSAE
route de Saclay
91128 PALAISEAU Cedex
FRANCE

Abstract: A statistical model is said to be un-normalised when its likelihood function involves an intractable normalising constant. Two popular methods for parameter inference for these models are MC-MLE (Monte Carlo maximum likelihood estimation), and NCE (noise contrastive estimation); both methods rely on simulating artificial data-points to approximate the normalising constant. While the asymptotics of MC-MLE have been established under general hypotheses (Geyer, 1994), this is not so for NCE. We establish consistency and asymptotic normality of NCE estimators under mild assumptions. We compare NCE and MC-MLE under several asymptotic regimes. In particular, we show that, when $m \rightarrow \infty$ while n is fixed (m and n being respectively the number of artificial data-points, and actual data-points), the two estimators are asymptotically equivalent. Conversely, we prove that, when the artificial data-points are IID, and when $n \rightarrow \infty$ while m/n converges to a positive constant, the asymptotic variance of a NCE estimator is always smaller than the asymptotic variance of the corresponding MC-MLE estimator. We illustrate the variance reduction brought by NCE through a numerical study.

Received February 2018.

Contents

1	Introduction	3474
2	Set-up and notations	3476
3	Asymptotics of the Monte Carlo error	3477
	3.1 Consistency	3478
	3.2 Asymptotic normality, comparison with MC-MLE	3478
4	Asymptotics of the overall error	3480
	4.1 Consistency	3481
	4.2 Asymptotic normality	3481

4.3	Comparison of asymptotic variances	3484
5	Numerical example	3485
6	Conclusion	3487
	Acknowledgements	3487
A	Main Proofs	3488
A.1	Technical lemmas	3488
A.2	Proof of Theorem 1	3490
A.3	Proof of Theorem 2	3492
A.4	Proof of Theorem 3	3495
A.4.1	NCE consistency	3495
A.5	Proof of Theorem 4	3499
A.5.1	NCE asymptotic normality	3499
A.6	Proof of Theorem 5	3503
B	Supplementary Proofs	3506
B.1	Proofs of technical lemmas	3506
B.1.1	Proof of Lemma 1	3506
B.1.2	Proof of Lemma 2	3506
B.1.3	Proof of Lemma 3	3507
B.1.4	Proof of Lemma 4	3507
B.1.5	Proof of Lemma 5	3509
B.1.6	Proof of Lemma 6	3510
B.2	Proofs of the remaining lemmas	3511
B.2.1	Proof of Lemma 7	3511
B.2.2	Proof of Lemma 9	3511
B.2.3	Proof of Lemma 10	3512
B.3	Proofs of MC-MLE consistency and asymptotic normality	3513
B.3.1	MC-MLE consistency	3513
B.3.2	MC-MLE asymptotic normality	3514
B.4	Proofs related to exponential families	3516
	References	3517

1. Introduction

Consider a set of probability densities $\{f_\theta : \theta \in \Theta\}$ with respect to some measure μ , defined on a space \mathcal{X} , such that:

$$f_\theta(x) = \frac{h_\theta(x)}{\mathcal{Z}(\theta)}$$

where h_θ is non-negative, and $\mathcal{Z}(\theta)$ is a normalising constant, $\mathcal{Z}(\theta) = \int_{\mathcal{X}} h_\theta(x) \mu(dx)$. A model based on such a family of densities is said to be un-normalised if function h_θ may be computed point-wise, but $\mathcal{Z}(\theta)$ is not available (i.e. it may not be computed in a reasonable CPU time).

Un-normalised models arise in several areas of machine learning and Statistics, such as deep learning (Salakhutdinov and Hinton, 2009), computer vision (Wang et al., 2013), image segmentation (Gu and Zhu, 2001), social network

modelling (Caimo and Friel, 2011), directional data modelling (Walker, 2011), among others. In most applications, data-points are assumed to be IID (independent and identically distributed); see however e.g. Mnih and Teh (2012) or Barthelmé and Chopin (2015) for applications of non-IID un-normalised models. In that spirit, we consider an un-normalised model of IID variables Y_1, \dots, Y_n , with log-likelihood (divided by n):

$$\ell_n(\theta) = \frac{1}{n} \sum_{i=1}^n \log h_\theta(y_i) - \log \mathcal{Z}(\theta). \tag{1}$$

The fact that $\mathcal{Z}(\theta)$ is intractable precludes standard maximum likelihood estimation.

Geyer (1994) wrote a seminal paper on un-normalised models, in which he proposed to estimate θ by maximising function

$$\ell_{n,m}^{\text{IS}}(\theta) = \frac{1}{n} \sum_{i=1}^n \log \frac{h_\theta(y_i)}{h_\psi(y_i)} - \log \left\{ \frac{1}{m} \sum_{j=1}^m \frac{h_\theta(x_j)}{h_\psi(x_j)} \right\} \tag{2}$$

where the x_j 's are m artificial data-points generated from a user-chosen distribution \mathbb{P}_ψ with density $f_\psi(x) = h_\psi(x)/\mathcal{Z}(\psi)$. The empirical average inside the second log is a consistent (as $m \rightarrow \infty$) importance sampling estimate of $\mathcal{Z}(\theta)/\mathcal{Z}(\psi)$. Function $\ell_{n,m}^{\text{IS}}$ is thus an approximation of the log-likelihood ratio $\ell_n(\theta) - \ell_n(\psi)$, whose maximiser is the MLE.

In many applications, the easiest way to sample from \mathbb{P}_ψ is to use MCMC (Markov chain Monte Carlo). Geyer (1994) established the asymptotic properties of the MC-MLE estimates under general conditions; in particular that the x_j 's are realisations of an ergodic process. This is remarkable, given that most of the theory on M-estimation (i.e. estimation obtained by maximising functions) is restricted to IID data.

More recently, Gutmann and Hyvärinen (2012) proposed an alternative approach to parameter estimation of un-normalised models, called noise contrastive estimation (NCE). It also relies on simulating artificial data-points x_1, \dots, x_m from distribution \mathbb{P}_ψ . The method consists in maximising the likelihood of a logistic classifier, where actual (resp. artificial) data-points are assigned label 1 (resp. 0). With symbols, the log-likelihood divided by n rewrites:

$$\ell_{n,m}^{\text{NCE}}(\theta, \nu) = \frac{1}{n} \sum_{i=1}^n \log q_{\theta, \nu}(y_i) + \frac{1}{m} \sum_{i=1}^m \log \left\{ (1 - q_{\theta, \nu}(x_i))^{m/n} \right\} \tag{3}$$

where $q_{\theta, \nu}(x)$, the probability of label 1 for a value x , is defined through odd-ratio function:

$$\log \left\{ \frac{q_{\theta, \nu}(x)}{1 - q_{\theta, \nu}(x)} \right\} = \log \left\{ \frac{h_\theta(x)}{h_\psi(x)} \right\} + \nu + \log \left(\frac{n}{m} \right).$$

The NCE estimator of θ is obtained by maximising function $\ell_{n,m}^{\text{NCE}}(\theta, \nu)$ with respect to both $\theta \in \Theta$ and $\nu \in \mathbb{R}$. In particular, when the considered model

is exponential, i.e. when $h_\theta(x) = \exp\{\theta^T S(x)\}$, for some statistic S , $\ell_{n,m}^{\text{NCE}}$ is the log-likelihood of a standard logistic regression, with covariate $S(x)$. In that case, implementing NCE is particularly straightforward.

This paper has two objectives: first, to establish the asymptotic properties of NCE when the artificial data-points are generated from an ergodic process (typically a MCMC sampler) in order to show that NCE is as widely applicable as MC-MLE; second, to compare the statistical efficiency of both methods.

As a preliminary step, we replace the original log-likelihood by a function defined on the extended space $\Theta \times \mathbb{R}$, called Poisson transform:

$$\ell_n(\theta, \nu) = \frac{1}{n} \sum_{i=1}^n \log \left\{ \frac{h_\theta(y_i)}{h_\psi(y_i)} \right\} + \nu - e^\nu \times \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)}. \quad (4)$$

This function is so called as it corresponds to the log-likelihood (up to a linear transformation) of a Poisson process with intensity $h_\theta(y) + \nu$, see Barthelmé and Chopin (2015) for details. The main property of this transformation is that it produces exactly the same MLE as the original likelihood: $(\hat{\theta}_n, \hat{\nu}_n)$ maximises (4) if and only if $\hat{\theta}_n$ maximises (1) and $\hat{\nu}_n = \log \left\{ \mathcal{Z}(\psi) / \mathcal{Z}(\hat{\theta}_n) \right\}$.

In the same way, we replace the MC-MLE log-likelihood by function

$$\ell_{n,m}^{\text{IS}}(\theta, \nu) = \frac{1}{n} \sum_{i=1}^n \log \left\{ \frac{h_\theta(y_i)}{h_\psi(y_i)} \right\} + \nu - \frac{e^\nu}{m} \sum_{j=1}^m \frac{h_\theta(x_j)}{h_\psi(x_j)} \quad (5)$$

which has the same maximiser (with respect to θ) as function (2). We thus obtain three objective functions defined with respect to the same parameter space, $\Theta \times \mathbb{R}$. This will greatly facilitate our analysis.

The paper is organised as follows. In Section 2, we introduce the set up and notations. In Section 3, we study the behaviour of the NCE estimator as $m \rightarrow \infty$ (while n is kept fixed). We prove that the NCE estimator converges to the MLE at the same $m^{-1/2}$ rate as the MC-MLE estimator, and the difference between the two estimators converges faster, at rate m^{-1} . In Section 4, we let both m and n go to infinity while $m/n \rightarrow \tau > 0$. We obtain asymptotic variances for both estimators, which admit a simple and interpretable decomposition. Using this decomposition, we are able to establish that when the artificial data-points are IID, the asymptotic variance of NCE is always smaller than the asymptotic variance of MC-MLE (for the same computational budget). Section 5 assesses this variance reduction in a numerical example. Section 6 discusses the practical implications of our results. All proofs are delegated to the appendix.

2. Set-up and notations

Unless explicitly stated, we will consider Θ to be an open subset of \mathbb{R}^d , with natural topology associated to the Euclidean norm. We consider a parametric statistical model $\{\mathbb{P}_\theta^{\otimes n} : \theta \in \Theta\}$, corresponding to n IID data-points lying in space $\mathcal{X} \subset \mathbb{R}^k$, associated with the corresponding Borel σ -Field. We assume

that the model is identifiable, and equipped with some dominating measure μ , inducing the log-likelihood (1). From now on, we work directly with the “extended” version of approximate and exact log-likelihoods, i.e. functions (3), (4) and (5), which are functions of extended parameter $\xi = (\theta, \nu)$, with $\xi \in \Xi = \Theta \times \mathbb{R}$. When convenient, we also write $\ell_n(\xi)$ for $\ell_n(\theta, \nu)$ and so on. An open ball in Ξ , centered on ξ and of radius ϵ , is denoted $B(\xi, \epsilon)$. We may also use this notation for balls in Θ .

The point of this paper is to study and compare point estimates $\hat{\xi}_{n,m}^{\text{IS}}$ and $\hat{\xi}_{n,m}^{\text{NCE}}$, which maximise functions (5) and (3). For the sake of generality, we allow these estimators to be approximate maximisers; i.e. we will refer to $\hat{\xi}_{n,m}^{\text{IS}}$ as an approximate MC-MLE if

$$\ell_{n,m}^{\text{IS}}(\hat{\xi}_{n,m}^{\text{IS}}) \geq \sup_{\xi \in \Xi} \ell_{n,m}^{\text{IS}}(\xi) - o(1) \quad \text{a.s.} \tag{6}$$

and with a similar definition for $\hat{\xi}_{n,m}^{\text{NCE}}$. The meaning of symbol $o(1)$ in (6) depends on the asymptotic regime: in Section 3, n is kept fixed, while $m \rightarrow \infty$, hence $o(1)$ means “converges to zero as $m \rightarrow \infty$ ”. In Section 4, both m and n go to infinity, and the meaning of $o(1)$ must be adapted accordingly.

In both asymptotic regimes, the main assumption regarding the sampling process is as follows.

(X1) The artificial data-points are realisations of a \mathbb{P}_ψ -ergodic process $(X_j)_{j \geq 1}$.

By \mathbb{P}_ψ -ergodicity, we mean that the following law of large number holds:

$$\frac{1}{m} \sum_{j=1}^m \varphi(X_j) \xrightarrow{m \rightarrow \infty} \mathbb{E}_\psi [\varphi(X)] = \int_{\mathcal{X}} \varphi(x) f_\psi(x) \mu(dx) \quad \text{a.s.}$$

for any measurable, real-valued function φ such that $\mathbb{E}_\psi [|\varphi(X)|] < +\infty$.

Assumption (X1) is mild. For instance, if the X_j 's are generated by a MCMC algorithm, this is equivalent to assuming that the simulated chain is aperiodic and ϕ -irreducible, which is true for all practical MCMC samplers; see e.g. Roberts and Rosenthal (2004).

Finally, note that, although notation \mathbb{P}_ψ suggests that the distribution of the artificial data-points belongs to the considered parametric model, this is not compulsory. All our results hold provided that the model is dominated by \mathbb{P}_ψ (i.e. $\mathbb{P}_\theta \ll \mathbb{P}_\psi$ for every $\theta \in \Theta$).

3. Asymptotics of the Monte Carlo error

In this section, the analysis is conditional on the observed data: n and y_1, \dots, y_n are fixed. The only source of randomness is the Monte Carlo error, and the quantity we seek to estimate is the (intractable) MLE. This regime was first studied for MC-MLE by Geyer (1994). For convenience, we suppose that the MLE exists and is unique; or equivalently that $\hat{\xi}_n = (\hat{\theta}_n, \hat{\nu}_n)$ is the unique maximiser of ℓ_n .

3.1. Consistency

We are able to prove NCE consistency (towards the MLE) using the same approach as Geyer (1994) for MC-MLE. Our consistency result relies on the following assumptions:

(C1) The random sequence $(\hat{\xi}_{n,m}^{\text{NCE}})_{m \geq 1}$ is an approximate NCE estimator, which belongs to a compact set almost surely.

(H1) The maps $\theta \mapsto h_\theta(x)$ are:

1. lower semi-continuous at each $\theta \in \Theta$, except for x in a \mathbb{P}_ψ -null set that may depend on θ ;
2. upper semi-continuous, for any x not in a \mathbb{P}_ψ -null set (that does not depend on θ), and for all $x = y_i$, $i = 1, \dots, n$.

Theorem 1. *Under assumptions (X1), (C1) and (H1), almost surely:*
 $\hat{\xi}_{n,m}^{\text{NCE}} \xrightarrow{m \rightarrow \infty} \hat{\xi}_n$.

This result is strongly linked to Theorems 1 and 4 of Geyer (1994), which state that $\hat{\theta}_{n,m}^{\text{IS}} \rightarrow \hat{\theta}_n$ as $m \rightarrow \infty$ under essentially the same assumptions. These assumptions are very mild: they basically require continuity of the maps $\theta \mapsto h_\theta(x)$, without any integrability condition.

Remark 1. *As noticed by Geyer (1994), the proof does not require Θ to be a subset of \mathbb{R}^d , consistency of MC-MLE as well as Theorem 1 hold more generally as soon as Θ is a separable metric space.*

3.2. Asymptotic normality, comparison with MC-MLE

In order to compare the Monte Carlo error of MC-MLE and NCE estimators, we make the following extra assumptions:

(H2) The maps $\theta \mapsto h_\theta(x)$ are twice continuously differentiable in a neighborhood of $\hat{\theta}_n$ for \mathbb{P}_ψ -almost every x , and for $x = y_i$, $i = 1, \dots, n$. The Hessian matrix $\mathbf{H} = \nabla^2 \ell_n(\hat{\theta}_n)$ is invertible. Moreover, for some $\varepsilon > 0$

$$\int_{\mathcal{X}} a_\varepsilon(x) \sup_{\theta \in B(\hat{\theta}_n, \varepsilon)} h_\theta(x) \mu(dx) < +\infty$$

$$\text{where } a_\varepsilon(x) = 1 + \sup_{\theta \in B(\hat{\theta}_n, \varepsilon)} \|\nabla_\theta \log h_\theta(x)\|^2 + \sup_{\theta \in B(\hat{\theta}_n, \varepsilon)} \|\nabla_\theta^2 \log h_\theta(x)\|.$$

(G1) Estimators $\hat{\xi}_{n,m}^{\text{IS}}$ and $\hat{\xi}_{n,m}^{\text{NCE}}$ converge to $\hat{\xi}_n$ almost surely, and are such that

$$\nabla \ell_{n,m}^{\text{IS}}(\hat{\xi}_{n,m}^{\text{IS}}) = o(m^{-1}), \quad \nabla \ell_{n,m}^{\text{NCE}}(\hat{\xi}_{n,m}^{\text{NCE}}) = o(m^{-1}).$$

(I1) For some $\varepsilon > 0$ the following integrability condition holds:

$$\mathbb{E}_\psi \left[b_\varepsilon(X) \sup_{\theta \in B(\hat{\theta}_n, \varepsilon)} \left(\frac{h_\theta(X)}{h_\psi(X)} \right)^2 \right] < +\infty$$

$$\text{where } b_\varepsilon(x) = 1 + \sup_{\theta \in B(\hat{\theta}_n, \varepsilon)} \|\nabla_\theta \log h_\theta(x)\|.$$

Measurability of the suprema in (H2) and (I1) is ensured by the lower semi-continuity of the two first differentials in a neighbourhood of $\hat{\theta}_n$. Assumption (H2) is a regularity condition that ensures in particular that the partition function $\theta \mapsto \mathcal{Z}(\theta) = \int_{\mathcal{X}} h_\theta(x) \mu(dx)$ is twice differentiable under the integral sign, in a neighbourhood of $\hat{\theta}_n$. Following Theorem 1, Assumption (G1) is trivial as soon as Assumptions (C1) and (H1) hold and $\hat{\xi}_{n,m}^{\text{IS}}$ and $\hat{\xi}_{n,m}^{\text{NCE}}$ are exact maximisers; in that case the gradients are zero. Integrability Assumption (I1) is the critical assumption. It is essentially a (locally uniform) second moment condition on the importance weights, with $\mathbb{P}_{\hat{\theta}_n}$ as the target distribution.

Theorem 2. *Under assumptions (X1), (H2), (G1) and (I1):*

$$m \left(\hat{\xi}_{n,m}^{\text{NCE}} - \hat{\xi}_{n,m}^{\text{IS}} \right) \xrightarrow{m \rightarrow \infty} n \left(-\mathcal{H}(\hat{\xi}_n) \right)^{-1} v(\hat{\xi}_n) \quad \text{a.s.} \tag{7}$$

where $\mathcal{H}(\xi) = \nabla_\xi^2 \ell_n(\xi)$, and $v(\xi)$ is defined as follows: let $g_\xi(x) = \log h_\theta(x) + \nu$, then

$$v(\xi) = \frac{1}{n} \sum_{i=1}^n \nabla_\xi g_\xi(y_i) \left(\frac{\exp\{g_\xi(y_i)\}}{h_\psi(y_i)} \right) - \mathbb{E}_\psi \left[\nabla_\xi g_\xi(X) \left(\frac{\exp\{g_\xi(X)\}}{h_\psi(X)} \right)^2 \right].$$

Before discussing the implications of Theorem 2, it is important to consider Geyer (1994)’s result about asymptotic normality of MC-MLE, which relies on the following assumption:

(N) For some covariance matrix \mathbf{A} we have:

$$\sqrt{m} \nabla \ell_{n,m}^{\text{IS}}(\hat{\theta}_n) \xrightarrow{m \rightarrow \infty} \mathcal{N}_d(\mathbf{0}_d, \mathbf{A})$$

As noticed by Geyer (1994), asymptotics of MC-MLE are quite similar to the asymptotics of maximum likelihood, and it can be shown that under assumptions (X1), (H2), (G1) and (N),

$$\sqrt{m} \left(\hat{\theta}_{n,m}^{\text{IS}} - \hat{\theta}_n \right) \xrightarrow{m \rightarrow \infty} \mathcal{N}_d(\mathbf{0}_d, \mathbf{H}^{-1} \mathbf{A} \mathbf{H}^{-1}).$$

Theorem 2 shows that the difference between the two point estimates is $\mathcal{O}(m^{-1})$, which is negligible relative to the $\mathcal{O}_{\mathbb{P}}(m^{-1/2})$ rate of convergence to $\hat{\theta}_n$. This proves that, when n is fixed, both approaches are asymptotically equivalent when it comes to approximate the MLE. In particular, Slutsky’s lemma implies asymptotic normality of the NCE estimator with the same asymptotic variance as for MC-MLE.

Assumptions (H2) and (I1) admit a much simpler formulation when the model belongs to an exponential family. This is the point of the following Proposition.

Proposition 1. *If the parametric model is exponential, i.e. if $h_\theta(x) = \exp\{\theta^T S(x)\}$ for some statistic S , then assumptions (H2) and (I1) are equivalent to the following assumptions (H2-exp) and (I1-exp):*

(H2-exp) The Hessian matrix of the log-likelihood $\mathbf{H} = \nabla^2 \ell_n(\hat{\theta}_n)$ is invertible.

(I1-exp) The MLE $\hat{\theta}_n$ lies in the interior of $\Theta_\psi = \left\{ \theta \in \Theta : \mathbb{E}_\psi \left[\left(\frac{h_\theta(X)}{h_\psi(X)} \right)^2 \right] < +\infty \right\}$.

The set Θ_ψ is convex whenever Θ is. In particular, this is true when Θ coincides with the natural space of parameters, defined as $\tilde{\Theta} = \{ \theta \in \mathbb{R}^d : \int_{\mathcal{X}} \exp \{ \theta^T S(x) \} \mu(dx) < +\infty \}$. If $\mathbb{P}_\psi \in \{ \mathbb{P}_\theta : \theta \in \Theta \}$, then (I1-exp) holds as soon as $2\hat{\theta}_n - \psi$ lies in the interior of $\tilde{\Theta}$.

Remark 2. Condition (N) requires a \sqrt{m} -CLT (central limit theorem) for the function $\varphi : x \mapsto (\nabla_\theta \log h_\theta)(h_\theta/h_\psi)(x)$ at $\theta = \hat{\theta}_n$. There has been an extensive literature on CLT's for Markov Chains, see e.g. Roberts and Rosenthal (2004) for a review. In particular, if $(X_j)_{j \geq 1}$ is a geometrically ergodic Markov Chain with stationary distribution \mathbb{P}_ψ , then assumption (N) holds if for some $\delta > 0$, $\varphi \in \mathbb{L}_{2+\delta}(\mathbb{P}_\psi)$. This assumption is very similar to assumption (I1), especially when the model is exponential.

In practice, implications of Theorem 2 must be considered cautiously, as the Euclidean norm of the limit in (7) will typically increase with n . For several well-known un-normalised models (e.g. Ising models, Exponential Random Graph Models), n is equal to one, in which case NCE and MC-MLE will always produce very close estimates. For other models however, it is known that the two estimators may behave differently, especially when the number of actual data-points is big and when simulations have a high computational cost (see Gutmann and Hyvärinen (2012)).

To investigate to which extent both approaches provide a good approximation of the true parameter value in these models, we will require both m and n to go to infinity. As it turns out, this will also make it possible to do finer comparison between $\hat{\xi}_{n,m}^{\text{IS}}$ and $\hat{\xi}_{n,m}^{\text{NCE}}$ (and thus between $\hat{\theta}_{n,m}^{\text{NCE}}$ and $\hat{\theta}_{n,m}^{\text{IS}}$). This is the point of the next section.

4. Asymptotics of the overall error

We now assume that observations y_i are realisations of IID random variables Y_i , with probability density f_{θ^*} , for some true parameter $\theta^* \in \Theta$, while the artificial data-points $(X_j)_{j \geq 1}$ remain generated from a \mathbb{P}_ψ -ergodic process. We also assume that $(Y_i)_{i \geq 1}$ and $(X_j)_{j \geq 1}$ are independent sequences; this regime was first studied for NCE by Gutmann and Hyvärinen (2012), although the X_j 's were assumed IID in that paper.

This asymptotic regime has some drawbacks: it assumes that the model is well specified, and that \mathbb{P}_ψ is chosen independently from the data. This is rarely true in practice, as the user will generally try to choose \mathbb{P}_ψ as close as possible to the data distribution to reduce the mean square error (see Section 5). Nevertheless, allowing both m and n to go to infinity turns out to provide a better understanding of the asymptotic behaviours of NCE and MC-MLE, at least for situations where the number of actual data-points may be large.

We assume implicitly that $m = m_n$ is a non-decreasing sequence of positive integers going to infinity when n does, while $m_n/n \rightarrow \tau \in (0, +\infty)$. Every limit when n goes to infinity should be understood accordingly. Finally, $\xi^* = (\theta^*, \nu^*)$ stands for the true extended parameter, where $\nu^* = \log \{Z(\psi)/Z(\theta^*)\}$.

4.1. Consistency

Our results concerning the overall consistency (to ξ^* , as both m and $n \rightarrow \infty$) of MC-MLE and NCE rely on the following assumptions:

- (C2) The random sequences $(\hat{\xi}_{n,m}^{\text{IS}})_{n \geq 1}$ and $(\hat{\xi}_{n,m}^{\text{NCE}})_{n \geq 1}$ are approximate MC-MLE and NCE estimators, and belong to a compact set almost surely.
- (H3) The maps $\theta \mapsto h_\theta(x)$ are continuous for \mathbb{P}_ψ -almost every x , and for any $\theta \in \Theta$ there is some $\varepsilon > 0$ such that

$$\int_{\mathcal{X}} \sup_{\phi \in B(\theta, \varepsilon)} \left(\log \frac{h_\phi(x)}{h_{\theta^*}(x)} \right)_+ h_{\theta^*}(x) \mu(dx) < +\infty.$$

Theorem 3. *Under assumptions (X1), (C2) and (H3), both estimators $\hat{\xi}_{n,m}^{\text{IS}}$ and $\hat{\xi}_{n,m}^{\text{NCE}}$ converge almost surely to ξ^* as $n, m \rightarrow \infty$, while $m/n \rightarrow \tau$.*

Our proofs of NCE and MC-MLE consistency are mainly inspired from Wald (1949)’s famous proof of MLE consistency, for which the same integrability condition (H3) is required. It is noteworthy that, under this regime, MC-MLE and NCE consistency essentially rely on the same assumptions as MLE consistency.

Remark 3. *As noticed by Wald (1949), the proof does not require Θ to be a subset of \mathbb{R}^d . Theorem 3 holds as soon as Θ is a metric space.*

Proposition 2. *If the parametric model is exponential, i.e. if $h_\theta(x) = \exp \{ \theta^T S(x) \}$ for some measurable statistic S , then assumption (H3) always holds.*

4.2. Asymptotic normality

To ensure the asymptotic normality of both NCE and MC-MLE estimates, we make the following assumption.

- (X2) The sequence $(X_j)_{j \geq 1}$ is a Harris ergodic Markov chain (that is, aperiodic, ϕ -irreducible and positive Harris recurrent; for definitions see Meyn and Tweedie (2012)), with stationary distribution \mathbb{P}_ψ . The Markov kernel associated with the chain $(X_j)_{j \geq 1}$, noted $P(x, dy)$, is reversible (satisfies detailed balance) with respect to \mathbb{P}_ψ , that is

$$\mathbb{P}_\psi(dx)P(x, dy) = \mathbb{P}_\psi(dy)P(y, dx). \tag{8}$$

Moreover, the chain $(X_j)_{j \geq 1}$ is *geometrically ergodic*, i.e. there is some $\rho \in [0, 1)$ and a positive measurable function M such that for \mathbb{P}_ψ -almost every x

$$\|P^n(x, \cdot) - \mathbb{P}_\psi(\cdot)\|_{\text{TV}} \leq M(x)\rho^n \tag{9}$$

where $P^n(x; dy)$ denote the n -step Markov transition kernel corresponding to P , and $\|\cdot\|_{\text{TV}}$ stands for the total variation norm.

Under (X2), for any measurable, real-valued function φ such that $\mathbb{E}_\psi[\varphi^2] < \infty$, then a \sqrt{m} -CLT holds, i.e.

$$\sqrt{m} \left(\frac{1}{m} \sum_{j=1}^m \varphi(X_j) - \mathbb{E}_\psi[\varphi(X)] \right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_\varphi^2) \quad (10)$$

where

$$\sigma_\varphi^2 = \mathbb{V}_\psi(\varphi(X)) + 2 \sum_{i=1}^{\infty} \text{Cov}(\varphi(X_0), \varphi(X_i)).$$

In the equation above, $\text{Cov}(\varphi(X_0), \varphi(X_i))$ stands for the i -th lag autocovariance of the chain at stationarity; that is with respect to the distribution defined by $X_0 \sim \mathbb{P}_\psi$ and $X_{i+1}|X_i \sim P(X_i, \cdot)$. The sequence of artificial data-points $(X_j)_{j \geq 1}$ is not assumed stationary. Since the chain is Harris recurrent, (10) holds whenever $X_1 = x$ for any $x \in \mathcal{X}$ (see e.g. Roberts and Rosenthal (2004), especially Theorem 4 and Proposition 29).

For convenience, we choose to assume that the kernel is reversible (which is true for any Metropolis-Hastings algorithm), but the reversibility assumption (8) is not compulsory, and may be replaced by slightly stronger integrability assumptions (see e.g. Roberts and Rosenthal (2004)); in particular, if reversibility is not assumed then (10) holds whenever $\varphi \in \mathbb{L}_{2+\delta}(\mathbb{P}_\psi)$. The critical assumption is geometric ergodicity.

Geometric ergodicity is obviously stronger than assumption (X1) which only requires a law of large numbers to hold. Nevertheless, geometric ergodicity remains a state of the art condition to ensure CLT's for Markov chains (see e.g. Roberts and Rosenthal (2004) and Bradley et al. (2005)), while it can often be checked for practical MCMC samplers. We present assumption (X2) as a practical condition for ensuring CLT's when the artificial data-points are generated from a MCMC sampler, while it also covers the IID case without loss of generality. Note though that CLT's can hold under more general conditions, e.g. when the Markov chain satisfies polynomial ergodicity (Jones, 2004).

Our asymptotic normality results rely on the following assumptions:

(H4) The maps $\theta \mapsto h_\theta(x)$ are twice continuously differentiable in a neighborhood of θ^* for \mathbb{P}_ψ -almost every x ; the Fisher Information $\mathbf{I}(\theta) = \mathbb{V}_\theta(\nabla_\theta \log h_\theta(Y))$ is invertible at $\theta = \theta^*$; and for some $\varepsilon > 0$

$$\int_{\mathcal{X}} c_\varepsilon(x) \sup_{\theta \in B(\theta^*, \varepsilon)} h_\theta(x) \mu(dx) < \infty$$

$$\text{where } c_\varepsilon(x) = 1 + \sup_{\theta \in B(\theta^*, \varepsilon)} \|\nabla_\theta \log h_\theta(x)\|^2 + \sup_{\theta \in B(\theta^*, \varepsilon)} \|\nabla_\theta^2 \log h_\theta(x)\|.$$

(G2) Estimators $\hat{\xi}_{n,m}^{\text{IS}}$ and $\hat{\xi}_{n,m}^{\text{NCE}}$ converge in probability to ξ^* , and are such that

$$\nabla \ell_{n,m}^{\text{IS}}(\hat{\xi}_{n,m}^{\text{IS}}) = o_{\mathbb{P}}(n^{-1/2}), \quad \nabla \ell_{n,m}^{\text{NCE}}(\hat{\xi}_{n,m}^{\text{NCE}}) = o_{\mathbb{P}}(n^{-1/2}).$$

(I3) At $\theta = \theta^*$, the following integrability condition holds:

$$\mathbb{E}_\psi \left[d_\theta(X) \left(\frac{h_\theta(X)}{h_\psi(X)} \right)^2 \right] < \infty$$

where $d_\theta(x) = 1 + \|\nabla_\theta \log h_\theta(x)\|^2$.

Theorem 4. Under assumptions (X2), (H4) and (G2), we have

$$\sqrt{n} \left(\hat{\xi}_{n,m}^{\text{NCE}} - \xi^* \right) \xrightarrow{\mathcal{D}} \mathcal{N}_{d+1} \left(0, \mathbf{V}_\tau^{\text{NCE}}(\xi^*) \right)$$

where

$$\begin{aligned} \mathbf{V}_\tau^{\text{NCE}}(\xi) &= \mathbf{J}_\tau(\xi)^{-1} \left\{ \boldsymbol{\Sigma}_\tau(\xi) + \tau^{-1} \boldsymbol{\Gamma}_\tau(\xi) \right\} \mathbf{J}_\tau(\xi)^{-1}, \\ \mathbf{J}_\tau(\xi) &= \mathbb{E}_\theta \left[(\nabla_\xi \nabla_\xi^T g_\xi) \left(\frac{\tau f_\psi}{\tau f_\psi + f_\theta} \right) (Y) \right], \\ \boldsymbol{\Sigma}_\tau(\xi) &= \mathbb{V}_\theta \left((\nabla_\xi g_\xi) \left(\frac{\tau f_\psi}{\tau f_\psi + f_\theta} \right) (Y) \right), \\ \boldsymbol{\Gamma}_\tau(\xi) &= \mathbb{V}_\psi \left(\varphi_\xi^{\text{NCE}}(X) \right) + 2 \sum_{i=1}^{+\infty} \text{Cov} \left(\varphi_\xi^{\text{NCE}}(X_0), \varphi_\xi^{\text{NCE}}(X_i) \right), \\ \varphi_\xi^{\text{NCE}}(x) &= (\nabla_\xi g_\xi) \frac{f_\theta}{f_\psi} \left(\frac{\tau f_\psi}{\tau f_\psi + f_\theta} \right) (x). \end{aligned}$$

Moreover, under assumptions (X2), (H4), (G2) and (I3), we have

$$\sqrt{n} \left(\hat{\xi}_{n,m}^{\text{IS}} - \xi^* \right) \xrightarrow{\mathcal{D}} \mathcal{N}_{d+1} \left(0, \mathbf{V}_\tau^{\text{IS}}(\xi^*) \right)$$

where

$$\begin{aligned} \mathbf{V}_\tau^{\text{IS}}(\xi) &= \mathbf{J}(\xi)^{-1} \left\{ \boldsymbol{\Sigma}(\xi) + \tau^{-1} \boldsymbol{\Gamma}(\xi) \right\} \mathbf{J}(\xi)^{-1}, \\ \mathbf{J}(\xi) &= \mathbb{E}_\theta \left[\nabla_\xi \nabla_\xi^T g_\xi(Y) \right], \\ \boldsymbol{\Sigma}(\xi) &= \mathbb{V}_\theta \left(\nabla_\xi g_\xi(Y) \right), \\ \boldsymbol{\Gamma}(\xi) &= \mathbb{V}_\psi \left(\varphi_\xi^{\text{IS}}(X) \right) + 2 \sum_{i=1}^{+\infty} \text{Cov} \left(\varphi_\xi^{\text{IS}}(X_0), \varphi_\xi^{\text{IS}}(X_i) \right), \\ \varphi_\xi^{\text{IS}}(x) &= (\nabla_\xi g_\xi) \frac{f_\theta}{f_\psi} (x). \end{aligned}$$

Remark 4. Second moment condition (I3) is critical. It basically forbids \mathbb{P}_ψ to be chosen as a too thin tail distribution compared to \mathbb{P}_{θ^*} . Assumption (I3) is needed for establishing MC-MLE asymptotic normality, but not for NCE (inequality 21 shows that condition (H4) is enough). This already shows that, under the considered regime, NCE is more robust (to \mathbb{P}_ψ) than MC-MLE.

Assumptions (H4) and (I3) admit a simpler formulation when the model is exponential, as shown by the following proposition.

Proposition 3. *If the parametric model is exponential, i.e. if $h_\theta(x) = \exp\{\theta^T S(x)\}$ for some statistic S , then assumptions (H4) and (I3) are equivalent to the following assumptions (H4-exp) and (I3-exp):*

(H4-exp) Fisher Information $\mathbf{I}(\theta) = \mathbb{V}_\theta(\nabla_\theta \log h_\theta(Y))$ is invertible at $\theta = \theta^*$.

(I3-exp) Parameter θ^* belongs to the interior of $\Theta_\psi = \left\{ \theta : \mathbb{E}_\psi \left[\left(\frac{h_\theta(X)}{h_\psi(X)} \right)^2 \right] < \infty \right\}$.

In particular, if $\mathbb{P}_\psi \in \{\mathbb{P}_\theta\}_{\theta \in \Theta}$, then (I3-exp) holds as soon as $2\theta^* - \psi$ belongs to the interior of $\tilde{\Theta} = \left\{ \theta \in \mathbb{R}^d : \int_{\mathcal{X}} \exp\{\theta^T S(x)\} \mu(dx) < \infty \right\}$.

4.3. Comparison of asymptotic variances

Theorem 5. *If the artificial data-points $(X_j)_{j \geq 1}$ are IID, then under assumptions (H4) and (I3), $\mathbf{V}_\tau^{\text{IS}}(\xi^*) \succcurlyeq \mathbf{V}_\tau^{\text{NCE}}(\xi^*)$, i.e. $\mathbf{V}_\tau^{\text{IS}}(\xi^*) - \mathbf{V}_\tau^{\text{NCE}}(\xi^*)$ is a positive semi-definite matrix.*

Theorem 5 shows that, asymptotically, when $m/n \rightarrow \tau > 0$, and when the artificial data-points are IID, the variance of a NCE estimator is always lower than the variance of the corresponding MC-MLE estimator. This inequality is with respect to the Loewner partial order on symmetric matrices. To our knowledge, this is the first theoretical result proving that NCE dominates MC-MLE in terms of mean square error. We failed however to extend this result to correlated Markov chains.

This inequality holds for any fixed ratio $\tau \in (0, +\infty)$, and any given sampling distribution \mathbb{P}_ψ , but the sharpness of the bound remains unknown. Typically, the bigger is τ , the closer the two variances will be, as the ratio $\tau f_\psi / \tau f_\psi + f_{\theta^*}$ gets closer to one. It is also the case when the sampling distribution \mathbb{P}_ψ is close to the true data distribution \mathbb{P}_{θ^*} . Geyer (1994) noticed that MC-MLE performs better when \mathbb{P}_ψ is close to \mathbb{P}_{θ^*} . Next proposition shows that when $\mathbb{P}_\psi = \mathbb{P}_{\theta^*}$, both variances can be related to the variance of the MLE.

Proposition 4. *If the artificial data-points are IID sampled from $\mathbb{P}_\psi = \mathbb{P}_{\theta^*}$, then under assumptions (H4) and (I3) we have*

$$\mathbf{V}_\tau^{\text{NCE}}(\xi^*) = \mathbf{V}_\tau^{\text{IS}}(\xi^*) = (1 + \tau^{-1})\mathbf{V}^{\text{MLE}}(\xi^*)$$

where $\mathbf{V}^{\text{MLE}}(\xi) = \mathbf{J}(\xi)^{-1} \boldsymbol{\Sigma}(\xi) \mathbf{J}(\xi)^{-1}$.

It is straightforward to check that, under the usual conditions ensuring asymptotic normality of the MLE, the extended maximiser of the Poisson Transform ℓ_n is also asymptotically normal with variance $\mathbf{V}^{\text{MLE}}(\xi^*)$. This proposition shows what we can expect from NCE and MC-MLE in a ideal scenario where the sampling distribution is the same as the true data distribution.

5. Numerical example

This section presents a numerical example that illustrates how the variance reduction brought by NCE may vary according to the sampling distribution \mathbb{P}_ψ and the ratio τ .

We consider observations IID distributed from the multivariate Gaussian distribution $\mathcal{N}_p(\mu, \Sigma)$ truncated to $(0, +\infty)^p$; that is Y_1, \dots, Y_n are IID with the following probability density with respect to Lebesgue’s measure:

$$f_{\mu, \Sigma}(x) = \frac{1}{\mathcal{Z}(\mu, \Sigma)} \exp \left\{ -\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu) \right\} \mathbb{1}_{(0, +\infty)^p}(x)$$

where

$$\mathcal{Z}(\mu, \Sigma) = (2\pi)^{p/2} |\Sigma|^{1/2} \mathbb{P}(W \in (0, +\infty)^p), \quad W \sim \mathcal{N}_p(\mu, \Sigma).$$

The probability $\mathbb{P}(W \in (0, +\infty)^p)$ is intractable for almost every (μ, Σ) . Numerical approximations of such probabilities quickly become inefficient when p increases.

It is well known that (truncated) Gaussian densities form an exponential family under the following parametrisation: for a given $\mu \in \mathbb{R}^p$ and $\Sigma \in \mathbb{S}_p^{++}$ (the set of positive definite matrices of size p), define $\theta = (\Sigma^{-1}\mu, \text{triu}(-(1/2)\Sigma^{-1}))$, and $S(x) = (x, \text{triu}(xx^T))$, where $\text{triu}(\cdot)$ is the upper triangular part. This parametrisation is minimal and the natural parameter space is a convex open subset of \mathbb{R}^q where $q = p + p(p + 1)/2$. Indeed, under the exponential formulation, we have $\Theta = \Theta_1 \times \Theta_2$ where $\Theta_1 = \mathbb{R}^p$ and Θ_2 is on open cone of $\mathbb{R}^{p(p+1)/2}$, in bijection with \mathbb{S}_p^{++} through the function $\text{triu}(\cdot)$.

The observations are generated IID from \mathbb{P}_θ (using rejection) for some true parameter $\theta = \theta^*$, corresponding to

$$\mu^* = \begin{pmatrix} 1 \\ -1 \\ 0.5 \end{pmatrix}, \quad \Sigma^* = \begin{pmatrix} 1 & 0.5 & 1 \\ 0.5 & 1.5 & 0.3 \\ 1 & 0.3 & 2 \end{pmatrix},$$

in the usual Gaussian parametrisation. The artificial data-points are sampled IID from the distribution \mathbb{P}_ψ associated with density $f_{\mu, \Sigma}$ with $\mu = \mathbf{0}_p$ and $\Sigma = \lambda \mathbf{I}_p$ for some $\lambda > 0$. (Since Σ is diagonal, we may sample the components independently; and since $\mu = \mathbf{0}_p$, we may sample each component by taking the absolute value of a Gaussian variate.) The sample size is fixed to $n = 1000$, while m is chosen such that the ratio m/n is equal to $\tau \in \{1, 5, 20, 100\}$. The distribution \mathbb{P}_ψ is chosen as stated above for $\lambda \in [1.5, 20]$.

Figure 1 plots estimates and confidence intervals of the mean square error ratio (mean Euclidean norm square error of the estimator divided by the asymptotic variance of the MLE) of both estimators (NCE and MC-MLE), based on 1000 independent replications. (Regarding the denominator of this ratio, note that the variance of the MLE may be estimated by performing noise contrastive estimation with $\mathbb{P}_\psi = \mathbb{P}_{\theta^*}$, see Proposition 4.)




FIG 1. Estimates and confidence intervals of the Mean Square Error ratios of MC-MLE (left) and NCE (right), compared to the MLE. The MSE ratio depends both on the variance of the proposal distribution λ and the number of artificial data-points $m = \tau \times n$ ($n = 1000$). A log-scale is used for both axes.




FIG 2. Estimates and confidence intervals of the Mean Square Error ratios of MC-MLE, compared to the NCE. The MSE ratio depends both on the variance of the proposal distribution λ and the number of artificial data-points $m = \tau \times n$ ($n = 1000$). A log-scale is used for both axes.

To facilitate the direct comparison between NCE and MC-MLE, we also plot in Figure 2 estimates and confidence intervals of the MSE ratio of MC-MLE over NCE. As expected from Theorem 5, this ratio is always higher than one; it becomes larger and larger as τ decreases, or as λ moves away from its optimal value (around 4). This suggests that NCE is more robust than MC-MLE to a poor choice for the reference distribution (especially thin tails distributions, i.e. when λ goes to zero).

Finally, we discuss a technical difficulty related to the constrained nature of the parameter space Θ . In principle, both the NCE and the MC-MLE estimators should be obtained through constrained optimisation (i.e. as maximisers of their respective objective functions over Θ). However, it is much easier (here, and in many cases) to perform an unconstrained optimisation (over \mathbb{R}^q). We must check then that the so obtained solution fulfils the constraint that defines Θ (here, that the solution corresponds to a matrix Σ which is definite positive). Figure 3 plots estimates and confidence intervals of the probability that both estimators belong




FIG 3. Estimates and confidence intervals of the probability of existence of MC-MLE (left) and NCE (right) estimators. For a fixed $n = 1000$, the probability of belonging to Θ is lower for MC-MLE, especially for small values of the variance of the proposal distribution λ and the number of artificial data-points $m = \tau \times n$. A log-scale is used for both axes.

to Θ . We see that NCE (when implemented without constraints) is much more likely to produce estimates that belong to Θ .

Note also that when the considered model is an exponential family (as in this case), both functions $\ell_{n,m}^{IS}$ and $\ell_{n,m}^{NCE}$ are convex. This implies that, when the unconstrained maximiser of these functions do not fulfil the constraint that defines Θ , then the constrained maximiser does not exist. (Any solution of the constrained optimisation program lies on the boundary of the constrained set.)

6. Conclusion

The three practical conclusions we draw from our results are that: (a) NCE is as widely applicable as MC-MLE (including when the X_j 's are generated using MCMC); (b) NCE and MC-MLE are asymptotically equivalent (as $m \rightarrow \infty$) when n is fixed; (c) NCE may provide lower-variance estimates than MC-MLE when n is large (provided that $m = \mathcal{O}(n)$). The variance reduction seems to be more important when the ratio $\tau = m/n$ is small, or when the reference distribution (for generating the X_j 's) is poorly chosen. Note that we proved (c) under the assumption that the X_j 's are IID, but we conjecture it also holds when they are generated using MCMC. Proving this conjecture may be an interesting avenue for future research.

As mentioned in the introduction, another advantage of NCE is its ease of implementation. In particular, when the considered model is exponential, NCE boils down to performing a standard logistic regression. For all these reasons, it seems reasonable to recommend NCE over MC-MLE to perform inference for un-normalised models.

Acknowledgements

The research of the first author is funded by a GENES doctoral scholarship. The research of the second author is partially supported by a grant from the

French National Research Agency (ANR) as part of the Investissements d’Avenir program (ANR-11-LABEX-0047). We are grateful to Bernard Delyon for letting us include in the supplement an English translation of some technical results (and their proofs) on ergodic processes that he derived in lecture notes (available in French at <https://perso.univ-rennes1.fr/bernard.delyon/param.pdf>).

Appendix A: Main Proofs

A.1. Technical lemmas

The following lemmas are prerequisites for the proofs of our main theorems. Most of them are ‘classical’ results, but for the sake of completeness, we provide the proofs of these lemmas in the supplement (Appendix B).

All these lemma apply to a \mathbb{P}_ψ -ergodic sequence of random variables, $(X_j)_{j \geq 1}$.

First lemma is a slightly disguised version of the law of large numbers, combined with the monotone convergence of a sequence of test functions.

Lemma 1. *Let $(f_m)_{m \geq 1}$ be a non-decreasing sequence of measurable, non negative real-valued functions converging pointwise towards f . Then we have:*

$$\frac{1}{m} \sum_{j=1}^m f_m(X_j) \xrightarrow[m \rightarrow +\infty]{a.s.} \mathbb{E}_\psi[f(X)].$$

This result holds whether the expectation is finite or infinite.

Second lemma is a natural generalisation of Lemma 1 to dominated convergence.

Lemma 2. *Let $(f_m)_{m \geq 1}$, f and g be measurable, real-valued functions, such that $(f_m)_{m \geq 1}$ converges pointwise towards f ; for any $m \geq 1$, $|f_m| \leq g$; and $\mathbb{E}_\psi[g(X)] < +\infty$. Then we have:*

$$\frac{1}{m} \sum_{j=1}^m f_m(X_j) \xrightarrow[m \rightarrow +\infty]{a.s.} \mathbb{E}_\psi[f(X)].$$

Third lemma is a generalisation of Lemma 1 to the degenerate case where the expectation is infinite. In that case, Lemma 3 shows that the monotonicity assumption is unnecessary.

Lemma 3. *Let $(f_m)_{m \geq 1}$, f and g be measurable, real-valued functions, such that $(f_m)_{m \geq 1}$ converges pointwise towards f ; g is non negative, $\mathbb{E}_\psi[g(X)] < +\infty$; for any $m \geq 1$, $f_m \leq g$; and $\mathbb{E}_\psi[f(X)_-] = +\infty$ where f_- stands for the negative part of f . Then we have:*

$$\frac{1}{m} \sum_{j=1}^m f_m(X_j) \xrightarrow[m \rightarrow +\infty]{a.s.} -\infty.$$

Fourth lemma is a uniform law of large numbers. It is well known in the IID case. This result does not actually require the independence assumption.

We present a generalisation of this result to ergodic processes. The proof is due to Bernard Delyon, who made it available in an unpublished course in French (Delyon (2018)). We present an English translation of the proof in the supplement (Appendix B).

Lemma 4. *Let K a compact subset of \mathbb{R}^d ; $(\theta, x) \mapsto \varphi(\theta, x)$ a measurable function defined on $K \times \mathcal{X}$ whose values lie on \mathbb{R}^p ; and suppose that the maps $\theta \mapsto \varphi(\theta, x)$ are continuous for \mathbb{P}_ψ -almost every x . Moreover, suppose that*

$$\mathbb{E}_\psi \left[\sup_{\theta \in K} \|\varphi(\theta, X)\| \right] < +\infty.$$

Then the function $\theta \mapsto \mathbb{E}_\psi[\varphi(\theta, X)]$ defined on K is continuous, and we have

$$\sup_{\theta \in K} \left\| \frac{1}{m} \sum_{j=1}^m \varphi(\theta, X_j) - \mathbb{E}_\psi[\varphi(\theta, X)] \right\| \xrightarrow[m \rightarrow +\infty]{a.s.} 0.$$

Consequently, if there is a random sequence $(\tilde{\theta}_m)_{m \geq 1}$ converging almost surely to some parameter $\tilde{\theta} \in \Theta$. Then we have

$$\left\| \frac{1}{m} \sum_{j=1}^m \varphi(\tilde{\theta}_m, X_j) - \mathbb{E}_\psi[\varphi(\tilde{\theta}, X)] \right\| \xrightarrow[m \rightarrow \infty]{} 0 \quad a.s.$$

Fifth lemma is also a well known result. It is often used to prove the weak convergence (usually asymptotic normality) of Z-estimators.

Lemma 5. *Define any probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and let $(\ell_n(\theta, \omega))_{n \geq 1}$ be measurable real-valued functions defined on $\mathbb{R}^d \times \Omega$. Let $\theta^* \in \mathbb{R}^d$ and $\varepsilon > 0$ such that for any $n \geq 1$ and for \mathbb{P} -almost every $\omega \in \Omega$ the map $\theta \mapsto \ell_n(\theta, \omega)$ is C^2 on $B(\theta^*, \varepsilon)$. Let $(\hat{\theta}_n)_{n \geq 1}$ be a random sequence converging in probability to θ^* . Suppose also that:*

- (a) $\{\nabla_\theta \ell_n(\theta)\}_{|\theta=\hat{\theta}_n} = o_{\mathbb{P}}(n^{-1/2})$,
- (b) $\sup_{\theta \in B(\theta^*, \varepsilon)} \|\nabla_\theta^2 \ell_n(\theta) - \mathcal{H}(\theta)\| \xrightarrow{\mathbb{P}} 0$, for some $\mathbb{R}^{d \times d}$ valued function \mathcal{H} continuous at θ^* , such that $\mathcal{H}(\theta^*)$ is full rank,
- (c) $\sqrt{n}\{\nabla_\theta \ell_n(\theta)\}_{|\theta=\hat{\theta}_n} \xrightarrow{\mathcal{D}} Z$, for some random vector Z .

Then

$$\sqrt{n}(\hat{\theta}_n - \theta^*) + \mathcal{H}(\theta^*)^{-1} \sqrt{n}\{\nabla_\theta \ell_n(\theta)\}_{|\theta=\hat{\theta}_n} \xrightarrow{\mathbb{P}} 0_{\mathbb{R}^d},$$

and, consequently

$$\sqrt{n}(\hat{\theta}_n - \theta^*) \xrightarrow{\mathcal{D}} -\mathcal{H}(\theta^*)^{-1} Z.$$

Sixth lemma is a technical tool required for proving asymptotic normality of NCE. It is particularly straightforward to prove in the IID case. We present a generalisation of this result to reversible, geometrically ergodic Markov chains.

Lemma 6. *Assume that (X2) holds. Let $(f_n)_{n \geq 1}$, f and g be measurable, real-valued functions, such that $(f_n)_{n \geq 1}$ converges pointwise towards f ; for any $n \geq 1$, $|f_n| \leq g$; and $\mathbb{E}_\psi[g(X)^2] < \infty$. Then we have*

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n \{f_n(X_i) - f(X_i)\} - \mathbb{E}[f_n(X) - f(X)] \right) \xrightarrow{\mathbb{P}} 0,$$

and, consequently

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n f_n(X_i) - \mathbb{E}[f_n(X)] \right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_f^2),$$

where $\sigma_f^2 = \mathbb{V}_\psi(f(X)) + 2 \sum_{i=1}^{+\infty} \text{Cov}(f(X_0), f(X_i)) < +\infty$.

A.2. Proof of Theorem 1

A standard approach to establish consistency of M-estimators is to prove some Glivenko-Cantelli result (uniform convergence), but, to the best of our knowledge, no such result exists under the general assumption that the underlying random variables (the X_j 's in our case) are generated from an ergodic process. Instead, we follow Geyer (1994)'s approach, which relies on establishing that function $-\ell_{n,m}^{\text{NCE}}$ epiconverges to $-\ell_n$. Epiconvergence is essentially a one sided locally uniform convergence, that ensures the convergence of minimisers; for a succinct introduction to epiconvergence, see Appendix A of Geyer (1994) and Chapter 7 of Rockafellar and Wets (2009).

We follow closely Geyer (1994). In particular, Theorem 4 of Geyer (1994) shows that: if a sequence of functions $\ell_{n,m}$ hypoconverges to some function ℓ_n which has a unique maximiser $\hat{\theta}_n$ and if a random sequence $(\hat{\theta}_{n,m})_{m \geq 1}$ is an approximate maximiser of $\ell_{n,m}$ which belongs to a compact set almost surely, then $\hat{\theta}_{n,m}$ converges to $\hat{\theta}_n$ almost surely. Consequently, to prove Theorem 1, we only have to prove that $\ell_{n,m}^{\text{NCE}}$ hypoconverges to ℓ_n (i.e. that $-\ell_{n,m}^{\text{NCE}}$ epiconverges to $-\ell_n$); that is

$$\ell_n(\theta, \nu) \leq \inf_{B \in \mathcal{N}(\theta, \nu)} \liminf_{m \rightarrow +\infty} \sup_{(\phi, \mu) \in B} \{\ell_{n,m}^{\text{NCE}}(\phi, \mu)\} \quad (11)$$

$$\ell_n(\theta, \nu) \geq \inf_{B \in \mathcal{N}(\theta, \nu)} \limsup_{m \rightarrow +\infty} \sup_{(\phi, \mu) \in B} \{\ell_{n,m}^{\text{NCE}}(\phi, \mu)\} \quad (12)$$

where $\mathcal{N}(\theta, \nu)$ denotes the set of neighborhoods of the point (θ, ν) .

Since $\Xi = \Theta \times \mathbb{R}$ is a separable metric space, there exists a countable base $\mathcal{B} = \{B_1, B_2, \dots\}$ for the considered topology. For any point (θ, ν) define the countable base of neighborhoods $\mathcal{N}_c(\theta, \nu) = \mathcal{B} \cap \mathcal{N}(\theta, \nu)$ which can replace $\mathcal{N}(\theta, \nu)$ in the infima of the preceding inequalities. Choose a countable dense subset $\Gamma_c = \{(\theta_1, \nu_1), (\theta_2, \nu_2), \dots\}$ as follows. For each k let (θ_k, ν_k) be a point of B_k such that:

$$\ell_n(\theta_k, \nu_k) \geq \sup_{(\phi, \mu) \in B_k} \{\ell_n(\phi, \mu)\} - \frac{1}{k}.$$

The proof is very similar to Theorem 1 of Geyer (1994). However, in this slightly different proof, we will need

$$\lim_{m \rightarrow +\infty} \left[\frac{1}{m} \sum_{j=1}^m \log \left\{ \left(1 + e^\nu \frac{nh_\theta(X_j)}{mh_\psi(X_j)} \right)^{\frac{m}{n}} \right\} \right] = \mathbb{E}_\psi \left[e^\nu \frac{h_\theta(X)}{h_\psi(X)} \right] = e^\nu \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)} \tag{13}$$

and

$$\lim_{m \rightarrow +\infty} \frac{1}{m} \sum_{j=1}^m \log \left(1 + \frac{n}{m} \inf_{(\phi, \mu) \in B} \left[e^\mu \frac{h_\phi(X_j)}{h_\psi(X_j)} \right] \right)^{\frac{m}{n}} = \mathbb{E}_\psi \left[\inf_{(\phi, \mu) \in B} \left\{ e^\mu \frac{h_\phi(X)}{h_\psi(X)} \right\} \right] \tag{14}$$

to hold simultaneously with probability one for any $(\theta, \nu) \in \Gamma_c$ and any $B \in \mathcal{B}$. For any fixed (θ, ν) , Lemma 1 applies to the maps $x \mapsto (1 + \frac{x}{m})^m$, and since any countable union of null sets is still a null set, convergence holds simultaneously for every element of Γ_c and \mathcal{B} with probability one. One may note that infima in the last equation are measurable under (H1) (in that case, an infima over any set $B \in \mathcal{B}$ can be replaced by an infima over the countable dense subset $B \cap \Gamma_c$).

Proving inequality (11) is straightforward:

$$\forall B \in \mathcal{B}, \quad \forall (\theta, \nu) \in B \cap \Gamma_c, \\ \ell_n(\theta, \nu) = \lim_{m \rightarrow +\infty} \ell_{n,m}^{\text{NCE}}(\theta, \nu) \leq \lim_{m \rightarrow +\infty} \inf \sup_{(\phi, \mu) \in B} \{ \ell_{n,m}^{\text{NCE}}(\phi, \mu) \}$$

and thus

$$\inf_{B \in \mathcal{N}_c(\theta, \nu)} \sup_{(\phi, \mu) \in B \cap \Gamma_c} \{ \ell_n(\phi, \mu) \} \leq \inf_{B \in \mathcal{N}_c(\theta, \nu)} \lim_{m \rightarrow +\infty} \inf \sup_{(\phi, \mu) \in B} \{ \ell_{n,m}^{\text{NCE}}(\phi, \mu) \}.$$

(Geyer, 1994) proved that $\theta \mapsto \mathcal{Z}(\theta)$ is lower semi-continuous (cf Theorem 1). This result directly implies that $(\theta, \nu) \mapsto \ell_n(\theta, \nu)$ is upper semi-continuous as a sum of upper semi-continuous functions. Thus the left hand side is equal to $l(\theta, \nu)$ by construction of Γ_c .

The proof of the second inequality also follows closely Geyer (1994):

$$\begin{aligned} & \inf_{B \in \mathcal{N}(\theta, \nu)} \limsup_{m \rightarrow +\infty} \sup_{(\phi, \mu) \in B} \{ \ell_{n,m}^{\text{NCE}}(\phi, \mu) \} \\ & \leq \inf_{B \in \mathcal{N}(\theta, \nu)} \left\{ \sup_{(\phi, \mu) \in B} \left[\frac{1}{n} \sum_{i=1}^n \log \left\{ \frac{h_\phi(y_i)}{h_\psi(y_i)} \right\} + \mu \right] \right. \\ & \quad \left. - \liminf_{m \rightarrow +\infty} \inf_{(\phi, \mu) \in B} \left[\frac{1}{n} \sum_{i=1}^n \log \left\{ 1 + \frac{n}{m} e^\mu \frac{h_\phi(y_i)}{h_\psi(y_i)} \right\} \right] \right. \\ & \quad \left. - \liminf_{m \rightarrow +\infty} \frac{1}{m} \sum_{j=1}^m \log \left(1 + \frac{n}{m} \inf_{(\phi, \mu) \in B} \left[e^\mu \frac{h_\phi(X_j)}{h_\psi(X_j)} \right] \right)^{\frac{m}{n}} \right\} \\ & = \frac{1}{n} \sum_{i=1}^n \log \left\{ \frac{h_\theta(y_i)}{h_\psi(y_i)} \right\} + \nu - \sup_{B \in \mathcal{N}(\theta, \nu)} \mathbb{E}_\psi \left[\inf_{(\phi, \mu) \in B} \left\{ e^\mu \frac{h_\phi(X)}{h_\psi(X)} \right\} \right]. \end{aligned}$$

The inequality follows directly from superadditivity of the supremum (and sub-additivity of the infimum) and the continuity and monotonicity of the maps $x \mapsto \log(1 + \frac{x}{m})^{\frac{m}{n}}$. The last equality holds because the infimum over $\mathcal{N}(\theta, \nu)$ can be replaced by the infimum over the countable set $\mathcal{A}_c(\theta, \nu)$: the set of open balls centered on (θ, ν) of radius k^{-1} , $k \geq 1$, which means the infimum is also the limit of a decreasing sequence, which can be splitted into three terms. The second term converges deterministically to zero, while convergences (13) and (14) apply for the first and third terms.

To conclude, apply the monotone convergence theorem to the remaining term:

$$\begin{aligned} \sup_{B \in \mathcal{A}_c(\theta, \nu)} \mathbb{E}_\psi \left[\inf_{(\phi, \mu) \in B} \left\{ e^\mu \frac{h_\phi(X)}{h_\psi(X)} \right\} \right] &= \mathbb{E}_\psi \left[\sup_{B \in \mathcal{A}_c(\theta, \nu)} \inf_{(\phi, \mu) \in B} \left\{ e^\mu \frac{h_\phi(X)}{h_\psi(X)} \right\} \right] \\ &= \mathbb{E}_\psi \left[e^\nu \frac{h_\theta(X)}{h_\psi(X)} \right] = e^\nu \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)}. \end{aligned}$$

A.3. Proof of Theorem 2

Define $g_\xi(x) = \log h_\theta(x) + \nu$, and the following gradients (dropping n and m in the notation for convenience):

$$\begin{aligned} \Psi^{\text{NCE}}(\xi) = \nabla \ell_{n,m}^{\text{NCE}}(\xi) &= \frac{1}{n} \sum_{i=1}^n \nabla_\xi g_\xi(y_i) \left(\frac{m h_\psi(y_i)}{m h_\psi(X_j) + n \exp\{g_\xi(y_i)\}} \right) \\ &\quad - \frac{1}{m} \sum_{j=1}^m \nabla_\xi g_\xi(X_j) \left(\frac{m \exp\{g_\xi(X_j)\}}{m h_\psi(X_j) + n \exp\{g_\xi(X_j)\}} \right), \\ \Psi^{\text{IS}}(\xi) = \nabla \ell_{n,m}^{\text{IS}}(\xi) &= \frac{1}{n} \sum_{i=1}^n \nabla_\xi g_\xi(y_i) - \frac{1}{m} \sum_{j=1}^m \nabla_\xi g_\xi(X_j) \left(\frac{\exp\{g_\xi(X_j)\}}{h_\psi(X_j)} \right). \end{aligned}$$

By Taylor-Lagrange, for any component k , $1 \leq k \leq d + 1$, there exists (a random variable) $\xi_m^{(k)} \in [\hat{\xi}_{n,m}^{\text{IS}}; \hat{\xi}_{n,m}^{\text{NCE}}]$ such that

$$\Psi_k^{\text{IS}}(\hat{\xi}_{n,m}^{\text{IS}}) = \Psi_k^{\text{IS}}(\hat{\xi}_{n,m}^{\text{NCE}}) + \left\{ \nabla \Psi_k^{\text{IS}}(\xi_m^{(k)}) \right\}^T \left(\hat{\xi}_{n,m}^{\text{IS}} - \hat{\xi}_{n,m}^{\text{NCE}} \right)$$

where $\Psi_k^{\text{IS}}(\xi)$ denotes the k -th component of $\Psi^{\text{IS}}(\xi)$, and $[\hat{\xi}_{n,m}^{\text{IS}}; \hat{\xi}_{n,m}^{\text{NCE}}]$ denotes the line segment in \mathbb{R}^{d+1} which joins $\hat{\xi}_{n,m}^{\text{IS}}$ and $\hat{\xi}_{n,m}^{\text{NCE}}$.

By assumption (G1), the left hand side is $o_{\mathbb{P}}(m^{-1})$. The matrix form yields:

$$\begin{aligned} o_{\mathbb{P}}(m^{-1}) &= \Psi^{\text{IS}}(\hat{\xi}_{n,m}^{\text{NCE}}) + \mathbf{H}_m^{\text{IS}} \left(\hat{\xi}_{n,m}^{\text{IS}} - \hat{\xi}_{n,m}^{\text{NCE}} \right), \\ \mathbf{H}_m^{\text{IS}} &= \begin{pmatrix} \left\{ \nabla \Psi_1^{\text{IS}}(\xi_m^{(1)}) \right\}^T \\ \vdots \\ \left\{ \nabla \Psi_{d+1}^{\text{IS}}(\xi_m^{(d+1)}) \right\}^T \end{pmatrix}. \end{aligned}$$

Let us prove first the convergence of the Hessian matrix. Lemma 4 can be applied to each row component of the following matrix-valued function, the uniform norm of which is \mathbb{P}_ψ -integrable under (H2):

$$\begin{aligned} \varphi_h : (\xi, x) \mapsto & \left(\frac{1}{n} \sum_{i=1}^n \nabla_\xi^2 g_\xi(y_i) \right) \\ & - \left(\nabla_\xi^2 g_\xi(x) + \nabla_\xi g_\xi(x) \{ \nabla_\xi g_\xi(x) \}^T \right) \left(\frac{\exp\{g_\xi(x)\}}{h_\psi(x)} \right). \end{aligned}$$

Convergences of the $d + 1$ rows of \mathbf{H}_m^{IS} can be combined to get the following result:

$$\left\| \mathbf{H}_m^{\text{IS}} - \mathcal{H}(\hat{\xi}_n) \right\|_{m \rightarrow \infty} \rightarrow 0 \quad \text{a.s.}$$

where

$$\mathcal{H}(\xi) = \mathbb{E}_\psi [\varphi_h(\xi, X)] = \nabla_\xi^2 \ell_n(\xi).$$

It turns out that $\mathcal{H}(\hat{\xi}_n)$ is invertible as soon as (H2) holds. This is the point of the following lemma. This implies in particular that \mathbf{H}_m^{IS} is eventually invertible with probability one.

Lemma 7. *Assume (H2) holds. At the point $\xi = \hat{\xi}_n$, the Hessian matrix of the Poisson Transform $\nabla_\xi^2 \ell_n(\xi)$ is negative definite if and only if the Hessian of the log-likelihood $\nabla_\theta^2 \ell_n(\theta)$ is definite negative.*

The proof of Lemma 7 follows from a direct block matrix computation (using Schur’s complement). For the sake of completeness, we present a proof in the supplement (Appendix B).

Now, let us prove the convergence of the gradient. By assumption (G1), we can write $\Psi^{\text{IS}}(\hat{\xi}_{n,m}^{\text{NCE}}) = \Delta_m + o(m^{-1})$, where:

$$\begin{aligned} \Delta_m = & \Psi^{\text{IS}}(\hat{\xi}_{n,m}^{\text{NCE}}) - \Psi^{\text{NCE}}(\hat{\xi}_{n,m}^{\text{NCE}}) \\ = & \left\{ \frac{1}{n} \sum_{i=1}^n \nabla_\xi g_\xi(y_i) \left(\frac{n \exp\{g_\xi(y_i)\}}{mh_\psi(y_i) + n \exp\{g_\xi(y_i)\}} \right) \right. \\ & \left. - \frac{1}{m} \sum_{j=1}^m \nabla_\xi g_\xi(X_j) \left(\frac{\exp\{g_\xi(X_j)\}}{h_\psi(X_j)} \right) \left(\frac{n \exp\{g_\xi(X_j)\}}{mh_\psi(X_j) + n \exp\{g_\xi(X_j)\}} \right) \right\} \Big|_{\xi = \hat{\xi}_{n,m}^{\text{NCE}}} \end{aligned}$$

hence

$$\begin{aligned} \frac{m}{n} \Delta_m = & \left\{ \frac{1}{n} \sum_{i=1}^n \nabla_\xi g_\xi(y_i) \left(\frac{\exp\{g_\xi(y_i)\}}{h_\psi(y_i)} \right) \left(1 - \frac{n \exp\{g_\xi(y_i)\}}{mh_\psi(y_i) + n \exp\{g_\xi(y_i)\}} \right) \right. \\ & - \frac{1}{m} \sum_{j=1}^m \nabla_\xi g_\xi(X_j) \left(\frac{\exp\{g_\xi(X_j)\}}{h_\psi(X_j)} \right)^2 \\ & \left. \times \left(1 - \frac{n \exp\{g_\xi(X_j)\}}{mh_\psi(X_j) + n \exp\{g_\xi(X_j)\}} \right) \right\} \Big|_{\xi = \hat{\xi}_{n,m}^{\text{NCE}}} \end{aligned}$$

$$\begin{aligned}
&= \left\{ \frac{1}{n} \sum_{i=1}^n \nabla_{\xi} g_{\xi}(y_i) \left(\frac{\exp\{g_{\xi}(y_i)\}}{h_{\psi}(y_i)} \right) - \frac{1}{m} \sum_{j=1}^m \nabla_{\xi} g_{\xi}(X_j) \left(\frac{\exp\{g_{\xi}(X_j)\}}{h_{\psi}(X_j)} \right)^2 \right. \\
&\quad - \frac{1}{n} \sum_{i=1}^n \nabla_{\xi} g_{\xi}(y_i) \left(\frac{\exp\{g_{\xi}(y_i)\}}{h_{\psi}(y_i)} \right) \left(\frac{n \exp\{g_{\xi}(y_i)\}}{m h_{\psi}(y_i) + n \exp\{g_{\xi}(y_i)\}} \right) \\
&\quad + \frac{1}{m} \sum_{j=1}^m \nabla_{\xi} g_{\xi}(X_j) \left(\frac{\exp\{g_{\xi}(X_j)\}}{h_{\psi}(X_j)} \right)^2 \\
&\quad \left. \times \left(\frac{n \exp\{g_{\xi}(X_j)\}}{m h_{\psi}(X_j) + n \exp\{g_{\xi}(X_j)\}} \right) \right\} \Big|_{\xi = \hat{\xi}_{n,m}^{\text{NCE}}}.
\end{aligned}$$

The two last terms of the right hand side are residuals for which we want to bound the uniform norm over the ball $B(\hat{\theta}_n, \varepsilon)$. The sup norm of the second term is eventually bounded by:

$$\frac{1}{m} \sup_{\xi \in B(\hat{\xi}_n, \varepsilon)} \sum_{i=1}^n \|\nabla_{\xi} g_{\xi}(y_i)\| \left(\frac{\exp\{g_{\xi}(y_i)\}}{h_{\psi}(y_i)} \right)^2 \xrightarrow{m \rightarrow \infty} 0.$$

The sup norm of the third term is eventually bounded by $\frac{1}{m} \sum_{j=1}^m f_m(X_j)$ where

$$f_m(x) = \sup_{\xi \in B(\hat{\xi}_n, \varepsilon)} \|\nabla_{\xi} g_{\xi}(x)\| \left(\frac{\exp\{g_{\xi}(x)\}}{h_{\psi}(x)} \right)^2 \left(\frac{n \exp\{g_{\xi}(x)\}}{m h_{\psi}(x) + n \exp\{g_{\xi}(x)\}} \right)$$

and Lemma 2 applies under (I1) to the sequence $(f_m)_{m \geq 1}$ converging pointwise towards 0, and dominated by the integrable function $g(x) = \sup_{\xi \in B(\hat{\xi}_n, \varepsilon)} \|\nabla_{\xi} g_{\xi}(x)\| \times \left(\frac{\exp\{g_{\xi}(x)\}}{h_{\psi}(x)} \right)^2$.

The limit of $(m/n)\Delta_m$ is thus dictated by the behaviour of the first term. We apply Lemma 4 to the following vector-valued function, whose uniform norm is integrable under (I1) and under the continuity of the deterministic part assumed in (H2):

$$\varphi_g : (\xi, x) \mapsto \left(\frac{1}{n} \sum_{i=1}^n \nabla_{\xi} g_{\xi}(y_i) \frac{\exp\{g_{\xi}(y_i)\}}{h_{\psi}(y_i)} \right) - \nabla_{\xi} g_{\xi}(x) \left(\frac{\exp\{g_{\xi}(x)\}}{h_{\psi}(x)} \right)^2.$$

Lemma 4 yields $(m/n)\Delta_m \xrightarrow{m \rightarrow +\infty} v(\hat{\xi}_n)$ a.s. where

$$v(\xi) = \frac{1}{n} \sum_{i=1}^n \nabla_{\xi} g_{\xi}(y_i) \left(\frac{\exp\{g_{\xi}(y_i)\}}{h_{\psi}(y_i)} \right) - \mathbb{E}_{\psi} \left[\nabla_{\xi} g_{\xi}(X) \left(\frac{\exp\{g_{\xi}(X)\}}{h_{\psi}(X)} \right)^2 \right].$$

Combination of these facts ensure that on a set of probability one, we have eventually:

$$\frac{m}{n} \left(\hat{\xi}_{n,m}^{\text{IS}} - \hat{\xi}_{n,m}^{\text{NCE}} \right) = o(1) + (-\mathbf{H}_m^{\text{IS}})^{-1} \left(\frac{m}{n} \Delta_m + o(1) \right) \xrightarrow{m \rightarrow \infty} \left(-\mathcal{H}(\hat{\xi}_n) \right)^{-1} v(\hat{\xi}_n).$$

A.4. Proof of Theorem 3

The proof of MC-MLE consistency under the considered regime is a very straightforward adaptation of Wald’s proof of consistency for the MLE. We thus choose to present in appendix only the proof of NCE consistency, which is slightly more technical, although the sketch is similar. For the sake of completeness, a proof of MC-MLE consistency is presented in the supplement (Appendix B).

A.4.1. NCE consistency

For convenience, we choose to analyse a slightly different objective function (sharing the same maximiser with $\ell_{n,m}^{\text{NCE}}$), defined as:

$$M_n^{\text{NCE}}(\theta, \nu) = \frac{1}{n} \sum_{i=1}^n \left\{ \varphi_{(\theta, \nu)}(Y_i) - \zeta_{(\theta, \nu)}^{(n)}(Y_i) \right\} - \left(\frac{m_n}{n} \right) \times \frac{1}{m_n} \sum_{j=1}^{m_n} \zeta_{(\theta, \nu)}^{(n)}(X_j) \quad (15)$$

where $\varphi_{(\theta, \nu)}(x) = \log \left\{ \frac{e^\nu h_\theta(x)}{e^{\nu^*} h_{\theta^*}(x)} \right\}$ and $\zeta_{(\theta, \nu)}^{(n)}(x) = \log \left\{ \frac{\frac{m_n}{n} h_\psi(x) + e^\nu h_\theta(x)}{\frac{m_n}{n} h_\psi(x) + e^{\nu^*} h_{\theta^*}(x)} \right\}$.

We begin our proof with the following lemma.

Lemma 8. *For any fixed (θ, ν) , almost surely, $M_n^{\text{NCE}}(\theta, \nu) \xrightarrow[n \rightarrow \infty]{} \mathcal{M}_\tau^{\text{NCE}}(\theta, \nu)$, where:*

$$\begin{aligned} \mathcal{M}_\tau^{\text{NCE}}(\theta, \nu) = & \mathbb{E}_{\theta^*} \left[\log \left\{ \frac{e^\nu h_\theta}{e^{\nu^*} h_{\theta^*}} \right\} - \log \left\{ \frac{\tau h_\psi + e^\nu h_\theta}{\tau h_\psi + e^{\nu^*} h_{\theta^*}} \right\} \right] \\ & - \tau \mathbb{E}_\psi \left[\log \left\{ \frac{\tau h_\psi + e^\nu h_\theta}{\tau h_\psi + e^{\nu^*} h_{\theta^*}} \right\} \right] \end{aligned}$$

Moreover, (θ^*, ν^*) is the unique maximiser of $\mathcal{M}_\tau^{\text{NCE}}(\theta, \nu)$.

Proof. For any fixed (θ, ν) , the sequence $\zeta_{(\theta, \nu)}^{(n)}$ is eventually dominated (by a \mathbb{P}_ψ -integrable function), since for any $c > 0$ (in particular for $c = \tau \pm \varepsilon$) we have by Jensen’s inequality:

$$\mathbb{E}_\psi \left[\log \left\{ \frac{c h_\psi + e^\nu h_\theta}{c h_\psi + e^{\nu^*} h_{\theta^*}} \right\} \right] \geq \mathbb{E}_\psi \left[\log \left\{ \frac{f_\psi}{f_\psi + \frac{1}{c} f_{\theta^*}} \right\} \right] \geq -\log \left(1 + \frac{1}{c} \right) \quad (16)$$

$$\mathbb{E}_\psi \left[\log \left\{ \frac{c h_\psi + e^\nu h_\theta}{c h_\psi + e^{\nu^*} h_{\theta^*}} \right\} \right] \leq \mathbb{E}_\psi \left[\log \left\{ \frac{f_\psi + \frac{e^\nu \mathcal{Z}(\theta)}{c \mathcal{Z}(\psi)} f_\theta}{f_\psi} \right\} \right] \leq \log \left(1 + \frac{e^\nu \mathcal{Z}(\theta)}{c \mathcal{Z}(\psi)} \right) \quad (17)$$

Moreover, $\zeta_{(\theta, \nu)}^{(n)}$ converges pointwise to $\zeta_{(\theta, \nu)}^\infty(x) = \log \left\{ \frac{\tau h_\psi(x) + e^\nu h_\theta(x)}{\tau h_\psi(x) + e^{\nu^*} h_{\theta^*}(x)} \right\}$, thus Lemma 2 applies: the second empirical average in (15) converges almost surely to $\mathbb{E}_\psi [\zeta_{(\theta, \nu)}^\infty(X)]$.

Now, the sequence $\{\varphi_{(\theta, \nu)} - \zeta_{(\theta, \nu)}^{(n)}\}$ is upper bounded by the positive part of $\varphi_{(\theta, \nu)}$ which is \mathbb{P}_{θ^*} -integrable. In particular, if $\mathbb{E}_{\theta^*} [(\varphi_{(\theta, \nu)} - \zeta_{(\theta, \nu)}^\infty)_-] = +\infty$,

then Lemma 3 applies and the first empirical average in (15) converges towards $-\infty$.

Conversely, suppose that $\mathbb{E}_{\theta^*} [(\varphi_{(\theta,\nu)} - \zeta_{(\theta,\nu)}^\infty)_-] < \infty$. The law of large numbers would apply directly if the sequence m_n/n was exactly equal to τ . To handle this technical issue, we can consider the two following inequalities. Note that for any $a \geq b > 0$:

$$\begin{aligned} \log \left\{ \frac{ah_\psi(x) + e^\nu h_\theta(x)}{ah_\psi(x) + e^{\nu^*} h_{\theta^*}(x)} \right\} &\leq \log \left\{ \frac{\frac{a}{b}bh_\psi(x) + \frac{a}{b}e^\nu h_\theta(x)}{bh_\psi(x) + e^{\nu^*} h_{\theta^*}(x)} \right\} \\ \log \left\{ \frac{bh_\psi(x) + e^\nu h_\theta(x)}{bh_\psi(x) + e^{\nu^*} h_{\theta^*}(x)} \right\} &\leq \log \left\{ \frac{ah_\psi(x) + e^\nu h_\theta(x)}{\frac{b}{a}ah_\psi(x) + \frac{b}{a}e^{\nu^*} h_{\theta^*}(x)} \right\} \end{aligned}$$

This yields a useful uniform bound for any $a, b > 0$:

$$\left| \log \left\{ \frac{ah_\psi(x) + e^\nu h_\theta(x)}{ah_\psi(x) + e^{\nu^*} h_{\theta^*}(x)} \right\} - \log \left\{ \frac{bh_\psi(x) + e^\nu h_\theta(x)}{bh_\psi(x) + e^{\nu^*} h_{\theta^*}(x)} \right\} \right| \leq |\log a - \log b| \quad (18)$$

Thus, if $\mathbb{E}_{\theta^*} [(\varphi_{(\theta,\nu)} - \zeta_{(\theta,\nu)}^\infty)_-] < +\infty$, then the uniform bound (18) also ensures that:

$$\mathbb{E}_{\theta^*} \left[\left(\varphi_{(\theta,\nu)} - \log \left\{ \frac{ch_\psi + e^\nu h_\theta}{ch_\psi + e^{\nu^*} h_{\theta^*}} \right\} \right)_- \right] < +\infty$$

for any positive $c > 0$. The sequence can now be easily dominated and Lemma 2 applies; the first empirical average in (15) converges to $\mathbb{E}_{\theta^*} [\varphi_{(\theta,\nu)}(Y) - \zeta_{(\theta,\nu)}^\infty(Y)]$.

Finally, let us prove that (θ^*, ν^*) is the unique maximiser of $\mathcal{M}_\tau^{\text{NCE}}$. We have:

$$\begin{aligned} \mathcal{M}_\tau^{\text{NCE}}(\theta, \nu) &= \frac{1}{\mathcal{Z}(\psi)} \left[\int_{\mathcal{X}} -\log \left\{ \frac{e^{\nu^*} h_{\theta^*}(x)}{e^\nu h_\theta(x)} \right\} e^{\nu^*} h_{\theta^*}(x) \right. \\ &\quad \left. + \log \left\{ \frac{\tau h_\psi(x) + e^{\nu^*} h_{\theta^*}(x)}{\tau h_\psi(x) + e^\nu h_\theta(x)} \right\} (\tau h_\psi(x) + e^{\nu^*} h_{\theta^*}(x)) \lambda(dx) \right] \\ &\leq \frac{1}{\mathcal{Z}(\psi)} \left[\int_{\mathcal{X}} -\log \left\{ \frac{e^{\nu^*} h_{\theta^*}(x)}{e^\nu h_\theta(x)} \right\} e^{\nu^*} h_{\theta^*}(x) \right. \\ &\quad \left. + \log \left\{ \frac{\tau h_\psi(x)}{\tau h_\psi(x)} \right\} \tau h_\psi(x) + \log \left\{ \frac{e^{\nu^*} h_{\theta^*}(x)}{e^\nu h_\theta(x)} \right\} e^{\nu^*} h_{\theta^*}(x) \lambda(dx) \right] \\ &= 0 \end{aligned}$$

by the log-sum inequality, which applies with equality if and only if $e^\nu h_\theta(x) = e^{\nu^*} h_{\theta^*}(x)$ for \mathbb{P}_{θ^*} almost every x . This occurs if and only if ν and θ are chosen such that $f_{\theta^*}(x) = \frac{e^\nu}{\mathcal{Z}(\psi)} h_\theta(x)$. The model being identifiable, there is only one choice for both the unnormalized density and the normalizing constant; $\theta = \theta^*$ and $\nu = \nu^*$. \square

We now prove that the NCE estimator converges almost surely to this unique maximiser. Let $\eta > 0$, and define $K_\eta = \{\xi \in K : d(\xi, \xi^*) \geq \eta\}$ where K is the compact set defined in (C2).

Under (H3), monotone convergence ensures that for any $\xi \in K_\eta$:

$$\lim_{\varepsilon \downarrow 0} \mathbb{E}_{\theta^*} \left[\sup_{\beta \in B(\xi, \varepsilon)} \left(\varphi_\beta(Y) - \zeta_\beta^\infty(Y) \right) \right] = \mathbb{E}_{\theta^*} \left[\varphi_\xi(Y) - \zeta_\xi^\infty(Y) \right]$$

and

$$\lim_{\varepsilon \downarrow 0} \mathbb{E}_\psi \left[\inf_{\beta \in B(\xi, \varepsilon)} \zeta_\beta^\infty(X) \right] = \mathbb{E}_\psi \left[\zeta_\xi^\infty(X) \right].$$

Indeed, since maps $\theta \mapsto h_\theta(x)$ are continuous, the two previous expectations (on the left hand side) are respectively bounded from above for ε small enough, and bounded from below for any ε .

Thus, for any $\xi \in K_\eta$ and any $\gamma > 0$ we can find $\varepsilon_\xi > 0$ such that simultaneously:

$$\begin{aligned} \mathbb{E}_{\theta^*} \left[\sup_{\beta \in B(\xi, \varepsilon_\xi)} \left(\varphi_\beta(Y) - \zeta_\beta^\infty(Y) \right) \right] &\leq \mathbb{E}_{\theta^*} \left[\varphi_\xi(Y) - \zeta_\xi^\infty(Y) \right] + \frac{\gamma}{2} \\ \mathbb{E}_\psi \left[\inf_{\beta \in B(\xi, \varepsilon_\xi)} \zeta_\beta^\infty(X) \right] &\geq \mathbb{E}_\psi \left[\zeta_\xi^\infty(X) \right] - \frac{\gamma}{2\tau}. \end{aligned}$$

The compactness assumption ensures that there is a finite set $\{\xi_1, \dots, \xi_p\} \subset K_\eta$ such that $K_\eta \subset \bigcup_{k=1}^p B(\xi_k, \varepsilon_{\xi_k})$. This yields the following inequality:

$$\begin{aligned} \sup_{\xi \in K_\eta} M_n^{\text{NCE}}(\xi) &\leq \max_{k=1, \dots, p} \left\{ \frac{1}{n} \sum_{i=1}^n \sup_{q \geq n} \sup_{\beta \in B(\xi_k, \varepsilon_{\xi_k})} \left(\varphi_\beta(Y_i) - \zeta_\beta^{(q)}(Y_i) \right) \right. \\ &\quad \left. - \left(\frac{m_n}{n} \right) \times \frac{1}{m_n} \sum_{j=1}^{m_n} \inf_{\beta \in B(\xi_k, \varepsilon_{\xi_k})} \zeta_\beta^{(n)}(X_j) \right\} \end{aligned}$$

Choose any x for which the map $\theta \mapsto h_\theta(x)$ is continuous, and any $\xi \in K_\eta$. From the definition of $\zeta_\beta^{(n)}$, the following convergence is trivial:

$$\inf_{\beta \in B(\xi, \varepsilon_\xi)} \left(\zeta_\beta^{(n)}(x) \right) \xrightarrow{n \rightarrow +\infty} \inf_{\beta \in B(\xi, \varepsilon_\xi)} \left(\zeta_\beta^\infty(x) \right).$$

Moreover, using inequalities (16) et (17), one can easily show that the sequence $\left\{ \inf_{\beta \in B(\xi, \varepsilon_\xi)} \zeta_\beta^{(n)} \right\}$ is dominated (by a \mathbb{P}_ψ -integrable function). Lemma 2 applies:

$$\frac{1}{m_n} \sum_{j=1}^{m_n} \inf_{\beta \in B(\xi_k, \varepsilon_{\xi_k})} \zeta_\beta^{(n)}(X_j) \xrightarrow{n \rightarrow +\infty} \mathbb{E}_\psi \left[\inf_{\beta \in B(\xi, \varepsilon_\xi)} \zeta_\beta^\infty(X) \right] \quad \text{a.s.}$$

Now, subadditivity of the supremum and inequality (18) yield

$$\left| \sup_{\beta \in B(\xi, \varepsilon_\xi)} \left(\varphi_\beta(x) - \zeta_\beta^{(n)}(x) \right) - \sup_{\beta \in B(\xi, \varepsilon_\xi)} \left(\varphi_\beta(x) - \zeta_\beta^\infty(x) \right) \right|$$

$$\leq \sup_{\beta \in B(\xi, \varepsilon_\xi)} \left| \zeta_\beta^{(n)}(x) - \zeta_\beta^\infty(x) \right| \leq \left| \log \frac{m_n}{n} - \log \tau \right| \xrightarrow{n \rightarrow +\infty} 0$$

while monotonicity ensures that

$$\sup_{\beta \in B(\xi, \varepsilon_\xi)} \left(\varphi_\beta - \zeta_\beta^\infty \right) \leq \sup_{q \geq n} \sup_{\beta \in B(\xi, \varepsilon_\xi)} \left(\varphi_\beta - \zeta_\beta^{(q)} \right) \leq \sup_{\beta \in B(\xi, \varepsilon_\xi)} \left(\varphi_\beta \right)_+.$$

In the last inequality, the right hand side is \mathbb{P}_{θ^*} -integrable under (H3), and the sequence (in the middle) converges pointwise towards its lower bound whose negative part has either finite or infinite expectation. In both cases, either Lemma 2 or Lemma 3 can be applied and ensures that, almost surely:

$$\frac{1}{n} \sum_{i=1}^n \sup_{q \geq n} \sup_{\beta \in B(\xi_k, \varepsilon_{\xi_k})} \left(\varphi_\beta(Y_i) - \zeta_\beta^{(q)}(Y_i) \right) \xrightarrow{n \rightarrow +\infty} \mathbb{E}_{\theta^*} \left[\sup_{\beta \in B(\xi, \varepsilon_\xi)} \left(\varphi_\beta(Y) - \zeta_\beta^\infty(Y) \right) \right]$$

Combining these convergences simultaneously on a finite set, we get almost surely:

$$\begin{aligned} \limsup_{n \rightarrow +\infty} \sup_{\xi \in K_\eta} M_n^{\text{NCE}}(\xi) &\leq \max_{k=1, \dots, p} \left\{ \mathbb{E}_{\theta^*} \left[\sup_{\beta \in B(\xi_k, \varepsilon_{\xi_k})} \left(\varphi_\beta(Y) - \zeta_\beta^\infty(Y) \right) \right] \right. \\ &\quad \left. - \tau \mathbb{E}_\psi \left[\inf_{\beta \in B(\xi_k, \varepsilon_{\xi_k})} \zeta_\beta^\infty(X) \right] \right\} \\ &\leq \sup_{\xi \in K_\eta} \mathcal{M}_\tau^{\text{NCE}}(\xi) + \gamma \end{aligned}$$

This leads to the following inequality since γ is arbitrary small:

$$\limsup_{n \rightarrow +\infty} \sup_{\xi \in K_\eta} M_n^{\text{NCE}}(\xi) \leq \sup_{\xi \in K_\eta} \mathcal{M}_\tau^{\text{NCE}}(\xi) \quad \text{a.s.} \tag{19}$$

This last inequality is the heart of the proof. To conclude, we need only to show that the right hand side is negative, this is the aim of the following lemma.

Lemma 9. *Under (H3), the map $\xi \mapsto \mathcal{M}_\tau^{\text{NCE}}(\xi)$ is upper semi continuous.*

The proof of Lemma 9 is straightforward. For the sake of completeness, we present a proof in the supplement (Appendix B).

Since an upper semi continuous function achieves its maximum on any compact set, this lemma proves in particular that $\sup_{\xi \in K_\eta} \mathcal{M}_\tau^{\text{NCE}}(\xi) < 0$.

Thus inequality (19) implies that we can always find some $\alpha < 0$ such that eventually $\sup_{\xi \in K_\eta} M_n^{\text{NCE}}(\xi) < \alpha$, while (C2) implies that $M_n^{\text{NCE}}(\hat{\xi}_{n,m}^{\text{IS}}) \geq \sup_{\xi \in \Xi} M_n^{\text{NCE}}(\xi) - \delta_n$ where $\delta_n \rightarrow 0$, and where

$$\sup_{\xi \in \Xi} M_n^{\text{NCE}}(\xi) \geq M_n^{\text{NCE}}(\xi^*) \xrightarrow[n \rightarrow +\infty]{\text{a.s.}} \mathcal{M}^{\text{NCE}}(\xi^*) = 0.$$

Combination of these facts show that with probability one we have eventually:

$$M_n^{\text{NCE}}(\hat{\xi}_{n,m}^{\text{IS}}) > \alpha > \sup_{\xi \in K_\eta} M_n^{\text{NCE}}(\xi).$$

This is enough to prove strong consistency. Indeed, with probability one, $\hat{\xi}_{n,m}^{\text{NCE}}$ eventually escapes from K_η (otherwise there would be a contradiction with the inequality above). Since the sequence belongs to K by assumption, the sequence has no choice but to stay eventually in the ball of radius η . Thus with probability one, for any $\eta > 0$, we have eventually $d(\hat{\xi}_{n,m}^{\text{NCE}}, \xi^*) < \eta$. This is the definition of almost sure convergence.

A.5. Proof of Theorem 4

The proof of MC-MLE asymptotic normality is entirely classical. We choose to present in appendix only the proof of NCE asymptotic normality, which follows the same sketch but is slightly more technical. For the sake of completeness, a proof of MC-MLE asymptotic normality is presented in the supplement (Appendix B).

A.5.1. NCE asymptotic normality

Let $G_n^{\text{NCE}}(\xi) = \nabla_\xi \ell_{n,m}^{\text{NCE}}(\xi)$ and $\mathbf{H}_n^{\text{NCE}}(\xi) = \nabla_\xi^2 \ell_{n,m}^{\text{NCE}}(\xi)$. We have:

$$\begin{aligned} G_n^{\text{NCE}}(\xi) &= \frac{1}{n} \sum_{i=1}^n \nabla_\xi g_\xi(Y_i) \left(\frac{m_n h_\psi}{m_n h_\psi + n \exp\{g_\xi\}} \right) (Y_i) \\ &\quad - \frac{1}{m_n} \sum_{j=1}^{m_n} \nabla_\xi g_\xi(X_j) \frac{\exp\{g_\xi(X_j)\}}{h_\psi(X_j)} \left(\frac{m_n h_\psi}{m_n h_\psi + n \exp\{g_\xi\}} \right) (X_j) \\ \mathbf{H}_n^{\text{NCE}}(\xi) &= \frac{1}{n} \sum_{i=1}^n \nabla_\xi^2 g_\xi(Y_i) \left(\frac{m_n h_\psi}{m_n h_\psi + n \exp\{g_\xi\}} \right) (Y_i) \\ &\quad - \frac{1}{n} \sum_{i=1}^n \nabla_\xi \nabla_\xi^T g_\xi(Y_i) \left(\frac{m_n h_\psi n \exp\{g_\xi\}}{(m_n h_\psi + n \exp\{g_\xi\})^2} \right) (Y_i) \\ &\quad - \frac{1}{m_n} \sum_{j=1}^{m_n} \left\{ (\nabla_\xi^2 + \nabla_\xi \nabla_\xi^T) g_\xi(X_j) \right\} \\ &\quad \times \frac{\exp\{g_\xi(X_j)\}}{h_\psi(X_j)} \left(\frac{m_n h_\psi}{m_n h_\psi + n \exp\{g_\xi\}} \right) (X_j) \\ &\quad + \frac{1}{m_n} \sum_{j=1}^{m_n} \nabla_\xi \nabla_\xi^T g_\xi(X_j) \frac{\exp\{g_\xi(X_j)\}}{h_\psi(X_j)} \left(\frac{m_n h_\psi n \exp\{g_\xi\}}{(m_n h_\psi + n \exp\{g_\xi\})^2} \right) (X_j) \end{aligned}$$

We firstly show that the study can be reduced to the following random sequences:

$$\begin{aligned}
 G_n^\tau(\xi) &= \frac{1}{n} \sum_{i=1}^n \nabla_\xi g_\xi(Y_i) \left(\frac{\tau h_\psi}{\tau h_\psi + \exp\{g_\xi\}} \right) (Y_i) \\
 &\quad - \frac{1}{m_n} \sum_{j=1}^{m_n} \nabla_\xi g_\xi(X_j) \frac{\exp\{g_\xi(X_j)\}}{h_\psi(X_j)} \left(\frac{\tau h_\psi}{\tau h_\psi + \exp\{g_\xi\}} \right) (X_j) \\
 \mathbf{H}_n^\tau(\xi) &= \frac{1}{n} \sum_{i=1}^n \nabla_\xi^2 g_\xi(Y_i) \left(\frac{\tau h_\psi}{\tau h_\psi + \exp\{g_\xi\}} \right) (Y_i) \\
 &\quad - \frac{1}{n} \sum_{i=1}^n \nabla_\xi \nabla_\xi^T g_\xi(Y_i) \left(\frac{\tau h_\psi \exp\{g_\xi\}}{(\tau h_\psi + \exp\{g_\xi\})^2} \right) (Y_i) \\
 &\quad - \frac{1}{m_n} \sum_{j=1}^{m_n} \left\{ (\nabla_\xi^2 + \nabla_\xi \nabla_\xi^T) g_\xi(X_j) \right\} \frac{\exp\{g_\xi(X_j)\}}{h_\psi(X_j)} \left(\frac{\tau h_\psi}{\tau h_\psi + \exp\{g_\xi\}} \right) (X_j) \\
 &\quad + \frac{1}{m_n} \sum_{j=1}^{m_n} \nabla_\xi \nabla_\xi^T g_\xi(X_j) \frac{\exp\{g_\xi(X_j)\}}{h_\psi(X_j)} \left(\frac{\tau h_\psi \exp\{g_\xi\}}{(\tau h_\psi + \exp\{g_\xi\})^2} \right) (X_j)
 \end{aligned}$$

To do so, we show that almost surely $\sup_{\xi \in B(\xi^*, \varepsilon)} \|\mathbf{H}_n^{\text{NCE}}(\xi) - \mathbf{H}_n^\tau(\xi)\| \xrightarrow[n \rightarrow \infty]{} 0$.

Splitting the uniform norm into four parts yields:

$$\begin{aligned}
 &\sup_{\xi \in B(\xi^*, \varepsilon)} \left\| \mathbf{H}_n^{\text{NCE}}(\xi) - \mathbf{H}_n^\tau(\xi) \right\| \\
 &\leq \frac{1}{n} \sum_{i=1}^n \sup_{\xi \in B(\xi^*, \varepsilon)} \left\| \nabla_\xi^2 g_\xi(Y_i) \right\| \eta_n^\tau(Y_i) \\
 &\quad + \frac{1}{n} \sum_{i=1}^n \sup_{\xi \in B(\xi^*, \varepsilon)} \left\| \nabla_\xi \nabla_\xi^T g_\xi(Y_i) \right\| \Gamma_n^\tau(Y_i) \\
 &\quad + \frac{1}{m_n} \sum_{j=1}^{m_n} \sup_{\xi \in B(\xi^*, \varepsilon)} \left\| (\nabla_\xi^2 + \nabla_\xi \nabla_\xi^T) g_\xi(X_j) \right\| \frac{\exp\{g_\xi(X_j)\}}{h_\psi(X_j)} \eta_n^\tau(X_j) \\
 &\quad + \frac{1}{m_n} \sum_{j=1}^{m_n} \sup_{\xi \in B(\xi^*, \varepsilon)} \left\| \nabla_\xi \nabla_\xi^T g_\xi(X_j) \right\| \frac{\exp\{g_\xi(X_j)\}}{h_\psi(X_j)} \gamma_n^\tau(X_j) \tag{20}
 \end{aligned}$$

where the sequences of functions

$$\eta_n^\tau = \sup_{\xi \in B(\xi^*, \varepsilon)} \left| \frac{m_n h_\psi}{m_n h_\psi + n \exp\{g_\xi\}} - \frac{\tau h_\psi}{\tau h_\psi + \exp\{g_\xi\}} \right|$$

and

$$\gamma_n^\tau = \sup_{\xi \in B(\xi^*, \varepsilon)} \left| \frac{m_n h_\psi n \exp\{g_\xi\}}{(m_n h_\psi + n \exp\{g_\xi\})^2} - \frac{\tau h_\psi \exp\{g_\xi\}}{(\tau h_\psi + \exp\{g_\xi\})^2} \right|$$

are both upper bounded by 1 and converge pointwise (for any $x \in \mathcal{X}$) to 0 (use the continuity of $\xi \mapsto g_\xi(x)$).

Lemma 2 applies to each empirical average in (20) (every integrability condition holds under (H4)). The sum converges to 0 almost surely.

Now, we prove that $\forall a \in \mathbb{R}^{d+1} \quad a^T \sqrt{n} (G_n^{\text{NCE}}(\xi^*) - G_n^\tau(\xi^*)) \xrightarrow{\mathbb{P}} 0$.

Define $\eta_{\theta,\tau}^{(n)} = \frac{m_n f_\psi}{m_n f_\psi + n f_\theta} - \frac{\tau f_\psi}{\tau f_\psi + f_\theta}$. At the point $\xi = \xi^*$ we have:

$$\begin{aligned} & \sqrt{n} (G_n^{\text{NCE}}(\xi^*) - G_n^\tau(\xi^*)) \\ &= \sqrt{n} \left\{ \frac{1}{n} \sum_{i=1}^n (\nabla_\xi g_\xi) \eta_{\theta,\tau}^{(n)}(Y_i) - \mathbb{E}_\theta \left[(\nabla_\xi g_\xi) \eta_{\theta,\tau}^{(n)}(Y) \right] \right\} \Big|_{\xi=\xi^*} \\ & \quad - \sqrt{\frac{n}{m_n}} \times \sqrt{m_n} \left\{ \frac{1}{m_n} \sum_{j=1}^{m_n} (\nabla_\xi g_\xi) \frac{f_\theta}{f_\psi} \eta_{\theta,\tau}^{(n)}(X_j) - \mathbb{E}_\theta \left[(\nabla_\xi g_\xi) \eta_{\theta,\tau}^{(n)}(Y) \right] \right\} \Big|_{\xi=\xi^*} \end{aligned}$$

The sequence $|\eta_{\theta,\tau}^{(n)}|$ is upper bounded by 1 and converges pointwise towards 0. Moreover, for any $c > \tau$, the sequence $|\eta_{\theta,\tau}^{(n)}|$ is also eventually upper bounded by $2 \frac{c f_\psi}{c f_\psi + f_\theta}$. This ensures that both second moment conditions required holds under (H4) since:

$$\begin{aligned} & \int_{\mathcal{X}} \|\nabla_\xi g_\xi\|^2 \left(\frac{f_\theta}{f_\psi}\right)^2 \left(\frac{c f_\psi}{c f_\psi + f_\theta}\right)^2 f_\psi d\mu \\ &= c \int_{\mathcal{X}} \|\nabla_\xi g_\xi\|^2 \left(\frac{c f_\psi}{c f_\psi + f_\theta}\right) \left(\frac{f_\theta}{c f_\psi + f_\theta}\right) f_\theta d\mu \\ &\leq c \times \mathbb{E}_\theta [\|\nabla_\xi g_\xi\|^2] < +\infty \end{aligned} \tag{21}$$

We can thus apply Lemma 6:

$$\begin{aligned} & \sqrt{n} \left\{ \frac{1}{n} \sum_{i=1}^n (a^T \nabla_\xi g_\xi) \eta_{\theta,\tau}^{(n)}(Y_i) - \mathbb{E}_\theta \left[(a^T \nabla_\xi g_\xi) \eta_{\theta,\tau}^{(n)}(Y) \right] \right\} \Big|_{\xi=\xi^*} \xrightarrow{\mathbb{P}} 0 \\ & \sqrt{m_n} \left\{ \frac{1}{m_n} \sum_{j=1}^{m_n} (a^T \nabla_\xi g_\xi) \frac{f_\theta}{f_\psi} \eta_{\theta,\tau}^{(n)}(X_j) - \mathbb{E}_\theta \left[(a^T \nabla_\xi g_\xi) \eta_{\theta,\tau}^{(n)}(Y) \right] \right\} \Big|_{\xi=\xi^*} \xrightarrow{\mathbb{P}} 0 \end{aligned}$$

Finally, Cramér-Wold's device applies: $\sqrt{n} (G_n^{\text{NCE}}(\xi^*) - G_n^\tau(\xi^*)) \xrightarrow{\mathbb{P}} 0_{\mathbb{R}^{d+1}}$.

Now, we can work directly with G_n^τ and \mathbf{H}_n^τ , which is much easier. Indeed, Lemma 4 yields $\sup_{\xi \in B(\xi^*, \varepsilon)} \|\mathbf{H}_n^\tau(\xi) - \mathbf{H}_\tau(\xi)\| \xrightarrow[n \rightarrow \infty]{} 0$ almost surely, where:

$$\begin{aligned} \mathbf{H}_\tau(\xi) &= \mathbb{E}_{\theta^*} \left[\nabla_\xi^2 g_\xi(Y) \left(\frac{\tau h_\psi}{\tau h_\psi + \exp\{g_\xi\}} \right)(Y) \right] \\ & \quad - \mathbb{E}_{\theta^*} \left[\nabla_\xi \nabla_\xi^T g_\xi(Y) \left(\frac{\tau h_\psi \exp\{g_\xi\}}{(\tau h_\psi + \exp\{g_\xi\})^2} \right)(Y) \right] \\ & \quad - \mathbb{E}_\psi \left[\left\{ (\nabla_\xi^2 + \nabla_\xi \nabla_\xi^T) g_\xi(X) \right\} \frac{\exp\{g_\xi(X)\}}{h_\psi(X)} \left(\frac{\tau h_\psi}{\tau h_\psi + \exp\{g_\xi\}} \right)(X) \right] \\ & \quad + \mathbb{E}_\psi \left[\nabla_\xi \nabla_\xi^T g_\xi(X) \frac{\exp\{g_\xi(X)\}}{h_\psi(X)} \left(\frac{\tau h_\psi \exp\{g_\xi\}}{(\tau h_\psi + \exp\{g_\xi\})^2} \right)(X) \right] \end{aligned}$$

The only condition required is that the supremum norm of each integrand is integrable, which is satisfied under (H4) (bound the ratios by one).

Note also that, at the point $\xi = \xi^*$, functions \mathbf{H}_τ and $-\mathbf{J}_\tau$ coincide, where:

$$\mathbf{J}_\tau(\xi) = \mathbb{E}_\theta \left[(\nabla_\xi \nabla_\xi^T g_\xi) \left(\frac{\tau f_\psi}{\tau f_\psi + f_\theta} \right) (Y) \right]$$

A quick block matrix calculation shows that Schur's complement in $-\mathbf{J}_\tau(\xi)$ is proportional to:

$$\mathbf{I}_\tau(\theta) = \mathbb{V}_{X \sim \mathbb{Q}_\tau} \left(\nabla_\theta \log h_\theta(X) \right)$$

where \mathbb{Q}_τ refers to the probability measure whose density with respect to μ is defined as $q_\tau(x) \propto \frac{\tau f_\psi(x) f_\theta(x)}{\tau f_\psi(x) + f_\theta(x)}$. Note that $\mathbb{P}_\theta \ll \mathbb{Q}_\tau$ since the model is dominated by \mathbb{P}_ψ .

In particular, $\mathbf{J}_\tau(\xi^*)$ is invertible if and only if $\mathbf{I}_\tau(\theta^*)$ is invertible. Since $\mathbf{I}_\tau(\theta)$ is a covariance matrix, if it is not full rank, then $\nabla_\theta \log h_\theta(X)$ belongs to a hyperplane \mathbb{Q}_τ -almost surely (and thus \mathbb{P}_θ -almost surely). This contradicts assumption (H4) since the Fisher Information is full rank. Thus $\mathbf{I}_\tau(\theta^*)$ and $\mathbf{J}_\tau(\xi^*)$ are both invertible.

Now, we prove the weak convergence of the gradient:

$$\begin{aligned} & \sqrt{n} G_n^\tau(\xi^*) \\ &= \sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n (\nabla_\xi g_\xi) \left(\frac{\tau f_\psi}{\tau f_\psi + f_\theta} \right) (Y_i) - \mathbb{E}_\theta \left[(\nabla_\xi g_\xi) \left(\frac{\tau f_\psi}{\tau f_\psi + f_\theta} \right) (Y) \right] \right) \Big|_{\xi=\xi^*} \\ & \quad - \sqrt{\frac{n}{m_n}} \sqrt{m_n} \left(\frac{1}{m_n} \sum_{j=1}^{m_n} (\nabla_\xi g_\xi) \frac{f_\theta}{f_\psi} \left(\frac{\tau f_\psi}{\tau f_\psi + f_\theta} \right) (X_j) \right. \\ & \quad \left. - \mathbb{E}_\theta \left[(\nabla_\xi g_\xi) \left(\frac{\tau f_\psi}{\tau f_\psi + f_\theta} \right) (Y) \right] \right) \Big|_{\xi=\xi^*} \end{aligned}$$

Slutsky's lemma applies as follows. It is noteworthy that second moment conditions hold under (H4) only (use inequality (21)), whereas assumption (I3) is necessary for proving MC-MLE asymptotic normality (see B.3.2).

$$\sqrt{n} G_n^\tau(\xi^*) \xrightarrow{\mathcal{D}} \mathcal{N} \left(0_{\mathbb{R}^{d+1}}, \boldsymbol{\Sigma}_\tau(\xi^*) + \tau^{-1} \boldsymbol{\Gamma}_\tau(\xi^*) \right)$$

where

$$\begin{aligned} \boldsymbol{\Sigma}_\tau(\xi) &= \mathbb{V}_\theta \left((\nabla_\xi g_\xi) \left(\frac{\tau f_\psi}{\tau f_\psi + f_\theta} \right) (Y) \right), \\ \boldsymbol{\Gamma}_\tau(\xi) &= \mathbb{V}_\psi \left(\varphi_\xi^{\text{NCE}}(X) \right) + 2 \sum_{i=1}^{+\infty} \text{Cov} \left(\varphi_\xi^{\text{NCE}}(X_0), \varphi_\xi^{\text{NCE}}(X_i) \right), \\ \varphi_\xi^{\text{NCE}} &= (\nabla_\xi g_\xi) \frac{f_\theta}{f_\psi} \left(\frac{\tau f_\psi}{\tau f_\psi + f_\theta} \right). \end{aligned}$$

Finally, Lemma 5 applies:

$$\sqrt{n} \left(\hat{\xi}_{n,m}^{\text{NCE}} - \xi^* \right) \xrightarrow{\mathcal{D}} \mathcal{N}_{d+1} \left(0, \mathbf{V}_\tau^{\text{NCE}}(\xi^*) \right)$$

where $\mathbf{V}_\tau^{\text{NCE}}(\xi) = \mathbf{J}_\tau(\xi)^{-1} \{ \boldsymbol{\Sigma}_\tau(\xi) + \tau^{-1} \boldsymbol{\Gamma}_\tau(\xi) \} \mathbf{J}_\tau(\xi)^{-1}$.

A.6. Proof of Theorem 5

For convenience, we will use some shorthand notations. Define the real-valued measurable functions $Q = f_\theta / f_\psi$ and $R = \tau f_\psi / (\tau f_\psi + f_\theta)$. Note that we have the relationship $QR = \tau(1 - R)$. In the following, assume that $\xi = \xi^*$, and for any measurable function φ , note that $\mathbb{E}_\theta[\varphi]$ stands for the expectation of $\varphi(X)$ where $X \sim \mathbb{P}_\theta$, and that $\nabla \nabla^T g_\xi$ stands for the measurable matrix-valued function $x \mapsto \nabla_\xi g_\xi(x) (\nabla_\xi g_\xi(x))^T$. We begin with the following computations:

$$\begin{aligned} \mathbf{J}(\xi) &= \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi \right], \\ \boldsymbol{\Sigma}(\xi) &= \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi \right] - \mathbb{E}_\theta \left[\nabla g_\xi \right] \mathbb{E}_\theta \left[\nabla^T g_\xi \right], \\ \boldsymbol{\Gamma}(\xi) &= \mathbb{E}_\psi \left[\nabla \nabla^T g_\xi Q^2 \right] - \mathbb{E}_\psi \left[\nabla g_\xi Q \right] \mathbb{E}_\psi \left[\nabla^T g_\xi Q \right] \\ &= \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi (R^{-1} - 1) \right] \times \tau - \mathbb{E}_\theta \left[\nabla g_\xi \right] \mathbb{E}_\theta \left[\nabla^T g_\xi \right], \\ \mathbf{J}_\tau(\xi) &= \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi R \right], \\ \boldsymbol{\Sigma}_\tau(\xi) &= \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi R^2 \right] - \mathbb{E}_\theta \left[\nabla g_\xi R \right] \mathbb{E}_\theta \left[\nabla^T g_\xi R \right], \\ \boldsymbol{\Gamma}_\tau(\xi) &= \mathbb{E}_\psi \left[\nabla \nabla^T g_\xi Q^2 R^2 \right] - \mathbb{E}_\psi \left[\nabla g_\xi Q R \right] \mathbb{E}_\psi \left[\nabla^T g_\xi Q R \right] \\ &= \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi R (1 - R) \right] \times \tau - \mathbb{E}_\theta \left[\nabla g_\xi R \right] \mathbb{E}_\theta \left[\nabla^T g_\xi R \right]. \end{aligned}$$

Fortunately, the expression of the asymptotic variances simplify, as shown by the following lemma.

Lemma 10. *Let Z be any real-valued, non-negative measurable function such that $\mathbb{E}_\theta[\nabla \nabla^T g_\xi Z]$ is finite and invertible. Then:*

$$\mathbf{M} := \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi Z \right]^{-1} \mathbb{E}_\theta \left[\nabla g_\xi Z \right] \mathbb{E}_\theta \left[\nabla^T g_\xi Z \right] \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi Z \right]^{-1} = \begin{pmatrix} 0_{\mathbb{R}^d \times \mathbb{R}^d} & 0_{\mathbb{R}^d} \\ 0_{\mathbb{R}^d}^T & 1 \end{pmatrix}.$$

The proof of Lemma 10 follows from a direct block matrix computation. For the sake of completeness, we present a proof in the supplement (Appendix B).

Let \mathbf{M} be defined as in Lemma 10, matrix calculations yield

$$\begin{aligned} \mathbf{J}(\xi)^{-1} \boldsymbol{\Sigma}(\xi) \mathbf{J}(\xi)^{-1} &= \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi \right]^{-1} - \mathbf{M}, \\ \mathbf{J}_\tau(\xi)^{-1} \boldsymbol{\Sigma}_\tau(\xi) \mathbf{J}_\tau(\xi)^{-1} &= \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi R \right]^{-1} \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi R^2 \right] \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi R \right]^{-1} - \mathbf{M}, \end{aligned}$$

$$\begin{aligned} \mathbf{J}(\xi)^{-1}\mathbf{\Gamma}(\xi)\mathbf{J}(\xi)^{-1} &= \tau\mathbb{E}_\theta\left[\nabla\nabla^T g_\xi\right]^{-1}\mathbb{E}_\theta\left[\nabla\nabla^T g_\xi(R^{-1}-1)\right] \\ &\quad \times \mathbb{E}_\theta\left[\nabla\nabla^T g_\xi\right]^{-1} - \mathbf{M}, \\ \mathbf{J}_\tau(\xi)^{-1}\mathbf{\Gamma}_\tau(\xi)\mathbf{J}_\tau(\xi)^{-1} &= \tau\mathbb{E}_\theta\left[\nabla\nabla^T g_\xi R\right]^{-1}\mathbb{E}_\theta\left[\nabla\nabla^T g_\xi R(1-R)\right] \\ &\quad \times \mathbb{E}_\theta\left[\nabla\nabla^T g_\xi R\right]^{-1} - \mathbf{M}. \end{aligned}$$

Summing up these expressions we finally get

$$\begin{aligned} \mathbf{V}_\tau^{\text{IS}}(\xi) &= \mathbb{E}_\theta\left[\nabla\nabla^T g_\xi\right]^{-1}\mathbb{E}_\theta\left[\nabla\nabla^T g_\xi R^{-1}\right]\mathbb{E}_\theta\left[\nabla\nabla^T g_\xi\right]^{-1} - (1+\tau^{-1})\mathbf{M}, \\ \mathbf{V}_\tau^{\text{NCE}}(\xi) &= \mathbb{E}_\theta\left[\nabla\nabla^T g_\xi R\right]^{-1} - (1+\tau^{-1})\mathbf{M}. \end{aligned}$$

Now, to compare these variances, the idea is the following: $(x, y) \mapsto x^2/y$ is a convex function on \mathbb{R}^2 , which means Jensen's inequality ensures that for any random variables X, Y such that the following expectations exist we have $\mathbb{E}[X^2/Y] \geq \mathbb{E}[X]^2/\mathbb{E}[Y]$. Here the variances are matrices, but it turns out that it is possible to use a generalization of Jensen's inequality to the Loewner partial order on matrices. We introduce the following notations:

$\mathbb{M}_{n,m}$ is the set of $n \times m$ matrices,

\mathbb{S}_n is the set of $n \times n$ symmetric matrices,

\mathbb{S}_n^+ is the set of $(n \times n)$ symmetric positive semi-definite matrices,

\mathbb{S}_n^{++} is the set of $(n \times n)$ symmetric positive definite matrices,

$\mathcal{R}(A)$ is the range of A ,

$$\Delta_{n,m} = \left\{ (A, B) \in \mathbb{S}_n^+ \times \mathbb{M}_{n,m} : \mathcal{R}(B) \subset \mathcal{R}(A) \right\},$$

A^\dagger denotes the Moore-Penrose pseudo-inverse of A ,

\succcurlyeq denotes the Loewner partial order ($A_1 \succcurlyeq A_2$ iff $A_1 - A_2 \in \mathbb{S}_n^+$).

Lemma 11. *Let A, B be random matrices such that $(A, B) \in \Delta_{n,m}$ with probability one for some positive integers n, m . Let $\varphi : (A, B) \mapsto B^T A^\dagger B$ defined on $\Delta_{n,m}$. Then $\mathbb{E}[\varphi(A, B)] \succcurlyeq \varphi(\mathbb{E}[A], \mathbb{E}[B])$ provided that the three expectations exist.*

Proof. We just have to prove that f is convex with respect to \succcurlyeq , i.e. that for any $\lambda \in [0, 1]$, and any $(A_1, B_1), (A_2, B_2) \in \Delta_{n,m}$ we have $\lambda\varphi(A_1, B_1) + (1-\lambda)\varphi(A_2, B_2) \succcurlyeq \varphi(\lambda(A_1, B_1) + (1-\lambda)(A_2, B_2))$. Indeed, if this convex relationship on matrices is satisfied then for any $x \in \mathbb{R}^m$, the real-valued map $q : (A, B) \mapsto x^T \varphi(A, B)x$ is necessarily convex on $\Delta_{n,m}$. Consequently, Jensen's inequality applies, i.e. for any random $(A, B) \in \Delta_{n,m}$ a.s. and any $x \in \mathbb{R}^m$ we have

$$x^T \mathbb{E}[\varphi(A, B)]x = \mathbb{E}[q(A, B)] \geq q(\mathbb{E}[A], \mathbb{E}[B]) = x^T \varphi(\mathbb{E}[A], \mathbb{E}[B])x$$

which is the claim of the lemma.

Now, to prove that φ is convex with respect to \succcurlyeq , we use a property of the generalized Schur’s complement in positive semi-definite matrices (see Boyd and Vandenberghe (2004) p.651): let $A \in \mathbb{S}_n, B \in \mathbb{M}_{n,m}, C \in \mathbb{S}_m$, and consider the block symmetric matrix

$$D = \begin{pmatrix} A & B \\ B^T & C \end{pmatrix}.$$

Then we have

$$D \succcurlyeq 0 \iff A \succcurlyeq 0, \mathcal{R}(B) \subset \mathcal{R}(A), C - B^T A^\dagger B \succcurlyeq 0.$$

This leads to a straightforward proof of the convexity of φ . To our knowledge, the following trick is due to Ando (1979), whose original proof was restricted to positive definite matrices. We use the generalized Schur’s complement to extend this result to any $(A, B) \in \Delta_{n,m}$: let $\lambda \in [0, 1]$, and $(A_1, B_1), (A_2, B_2) \in \Delta_{n,m}$. The sum of two positive semi definite matrices is positive semi-definite thus we have

$$\lambda \begin{pmatrix} A_1 & B_1 \\ B_1^T & B_1^T A_1^\dagger B_1 \end{pmatrix} + (1 - \lambda) \begin{pmatrix} A_2 & B_2 \\ B_2^T & B_2^T A_2^\dagger B_2 \end{pmatrix} \succcurlyeq 0$$

which is the same as

$$\begin{pmatrix} \lambda A_1 + (1 - \lambda)A_2 & \lambda B_1 + (1 - \lambda)B_2 \\ \lambda B_1^T + (1 - \lambda)B_2^T & \lambda B_1^T A_1^\dagger B_1 + (1 - \lambda)B_2^T A_2^\dagger B_2 \end{pmatrix} \succcurlyeq 0.$$

Consequently, the generalised Schur’s complement in this last block matrix is also positive semi-definite, i.e.

$$\begin{aligned} & \lambda B_1^T A_1^\dagger B_1 + (1 - \lambda)B_2^T A_2^\dagger B_2 \\ & \succcurlyeq (\lambda B_1 + (1 - \lambda)B_2)^T [\lambda A_1 + (1 - \lambda)A_2]^\dagger (\lambda B_1 + (1 - \lambda)B_2) \end{aligned}$$

which proves the convexity of φ with respect to \succcurlyeq , and thus the claim of the lemma. \square

Finally, we compare the asymptotic variances of the two estimators. Note that for any $(A, B) \in \mathbb{S}_n \times \mathbb{S}_n^{++}$, and for every $x \in \mathbb{R}^n$, we have

$$x^T A x \geq 0 \iff x^T B A B x \geq 0.$$

Indeed, if A is semi definite positive then for some integer k we can find $P \in \mathbb{M}_{k,n}$ such that $A = P^T P$, moreover, B being symmetric we have $x^T B A B x = \|P B x\|^2 \geq 0$. The direct implication is enough since $B^{-1} \in \mathbb{S}_n^{++}$.

Consequently, the relation $\mathbf{V}_\tau^{\text{IS}}(\xi) \succcurlyeq \mathbf{V}_\tau^{\text{NCE}}(\xi)$ is equivalent to the relation

$$\mathbb{E}_\theta \left[\nabla \nabla^T g_\xi R^{-1} \right] \succcurlyeq \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi \right] \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi R \right]^{-1} \mathbb{E}_\theta \left[\nabla \nabla^T g_\xi \right]. \tag{22}$$

Inequality (22) is a direct application of Lemma 11 (let $B = \nabla \nabla^T g_\xi, A = B R$; note that $(A, B) \in \Delta_{d+1,d+1}$ almost surely; and use basic properties of the pseudo-inverse).

Appendix B: Supplementary Proofs

B.1. Proofs of technical lemmas

B.1.1. Proof of Lemma 1

For all $k \in \mathbb{N}$, eventually (for any $m \geq k$) we have

$$\frac{1}{m} \sum_{j=1}^m f_k(X_j) \leq \frac{1}{m} \sum_{j=1}^m f_m(X_j) \leq \frac{1}{m} \sum_{j=1}^m f(X_j).$$

Moreover, since $(X_j)_{j \geq 1}$ is \mathbb{P}_ψ -ergodic, the law of large numbers applies (even if the expectations are infinite, since f_k and f are non-negative):

$$\frac{1}{m} \sum_{j=1}^m f_k(X_j) \xrightarrow[m \rightarrow +\infty]{a.s.} \mathbb{E}_\psi[f_k(X)] \quad \text{and} \quad \frac{1}{m} \sum_{j=1}^m f(X_j) \xrightarrow[m \rightarrow +\infty]{a.s.} \mathbb{E}_\psi[f(X)].$$

Thus, there is a set of probability one on which for every $k \in \mathbb{N}$,

$$\mathbb{E}_\psi[f_k(X)] \leq \liminf_{m \rightarrow +\infty} \frac{1}{m} \sum_{j=1}^m f_m(X_j) \leq \limsup_{m \rightarrow +\infty} \frac{1}{m} \sum_{j=1}^m f_m(X_j) \leq \mathbb{E}_\psi[f(X)].$$

Since the inequality holds for any $k \in \mathbb{N}$, it also holds for the supremum over k :

$$\sup_{k \in \mathbb{N}} \mathbb{E}_\psi[f_k(X)] \leq \liminf_{m \rightarrow +\infty} \frac{1}{m} \sum_{j=1}^m f_m(X_j) \leq \limsup_{m \rightarrow +\infty} \frac{1}{m} \sum_{j=1}^m f_m(X_j) \leq \mathbb{E}_\psi[f(X)].$$

Finally, the monotone convergence theorem yields

$$\sup_{k \in \mathbb{N}} \mathbb{E}_\psi[f_k(X)] = \lim_{k \rightarrow +\infty} \mathbb{E}_\psi[f_k(X)] = \mathbb{E}_\psi \left[\lim_{k \rightarrow +\infty} f_k(X) \right] = \mathbb{E}_\psi[f(X)].$$

Consequently, the lower and upper limits are both equal to $\mathbb{E}_\psi[f(X)]$ almost surely.

B.1.2. Proof of Lemma 2

Since $(X_j)_{j \geq 1}$ is \mathbb{P}_ψ -ergodic and f is dominated by the integrable function g , the law of large numbers applies to function f . Thus, we just need to prove that

$$\left| \frac{1}{m} \sum_{j=1}^m \{f_m(X_j) - f(X_j)\} \right| \xrightarrow[m \rightarrow +\infty]{a.s.} 0.$$

To do so, use the fact that

$$\left| \frac{1}{m} \sum_{j=1}^m \{f_m(X_j) - f(X_j)\} \right| \leq \frac{1}{m} \sum_{j=1}^m |f_m(X_j) - f(X_j)|$$

$$\leq \frac{1}{m} \sum_{j=1}^m \sup_{k \geq m} |f_k(X_j) - f(X_j)|.$$

Define $h_m = 2g - \sup_{k \geq m} |f_k - f|$ and note that $(h_m)_{m \geq 1}$ is a non-decreasing sequence of non negative functions converging pointwise towards $2g$. Lemma 1 yields

$$\frac{1}{m} \sum_{j=1}^m h_m(X_j) \xrightarrow[m \rightarrow +\infty]{a.s.} \mathbb{E}_\psi[2g(X)].$$

Finally g is integrable, thus the remainder converges almost surely towards zero:

$$\frac{1}{m} \sum_{j=1}^m \sup_{k \geq m} |f_k(X_j) - f(X_j)| = \frac{2}{m} \sum_{j=1}^m g(X_j) - \frac{1}{m} \sum_{j=1}^m h_m(X_j) \xrightarrow[m \rightarrow +\infty]{a.s.} 0.$$

B.1.3. Proof of Lemma 3

Since g is integrable and $(X_j)_{j \geq 1}$ is \mathbb{P}_ψ -ergodic we have

$$\frac{1}{m} \sum_{j=1}^m g(X_j) \xrightarrow[m \rightarrow +\infty]{a.s.} \mathbb{E}_\psi[g(X)] < +\infty.$$

Thus we only need to show that

$$\frac{1}{m} \sum_{j=1}^m \{g(X_j) - f_m(X_j)\} \xrightarrow[m \rightarrow +\infty]{a.s.} +\infty.$$

Define $h_m = g - \sup_{k \geq m} f_k$, an increasing sequence of non negative functions converging pointwise to $g - f$. Lemma 1 applies whether $g - f$ is integrable or not:

$$\frac{1}{m} \sum_{j=1}^m (g(X_j) - f_m(X_j)) \geq \frac{1}{m} \sum_{j=1}^m h_m(X_j) \xrightarrow[m \rightarrow +\infty]{a.s.} \mathbb{E}_\psi[g(X) - f(X)].$$

The following inequality shows that the expectation is indeed infinite:

$$\mathbb{E}_\psi[g(X) - f(X)] = \mathbb{E}_\psi[(g(X) - f(X)_+) + f(X)_-] \geq \mathbb{E}_\psi[f(X)_-] = +\infty.$$

B.1.4. Proof of Lemma 4

To begin, note that measurability of the supremum is ensured by the lower semi-continuity of the maps $\theta \mapsto \varphi(\theta, x)$ on a set of probability one that does not depend on θ .

For every $\theta \in K$, consider the following function:

$$f_\theta(\eta) = \mathbb{E}_\psi \left[\sup_{\phi \in B(\theta, \eta)} \|\varphi(\phi, X) - \varphi(\theta, X)\| \right].$$

Dominated convergence implies that $f_\theta(\eta)$ converges to zero when η goes to zero. This is enough to ensure the continuity of the map $\theta \mapsto \mathbb{E}_\psi[\varphi(\theta, X)]$ because of the following inequality:

$$\sup_{\phi \in B(\theta, \eta)} \|\mathbb{E}_\psi[\varphi(\phi, X) - \varphi(\theta, X)]\| \leq f_\theta(\eta).$$

Let $\varepsilon > 0$. For every $\theta \in K$, we can always find $\eta_{(\theta, \varepsilon)} > 0$ small enough such that $f_\theta(\eta_{(\theta, \varepsilon)}) < \varepsilon$. Note that the open balls centered on $\theta \in K$ of radius $\eta_{(\theta, \varepsilon)}$, form an open cover of K , from which we can extract a finite subcover thanks to the compactness assumption. Thus we can build a finite set $\{\phi_1, \dots, \phi_I\} \subset K$ (centers of the balls) such that

$$K \subset \bigcup_{i=1}^I B_i, \quad B_i = B(\phi_i, \eta_{(\phi_i, \varepsilon)}).$$

Now, for any $\theta \in K$ define i_θ as the smallest integer $i \in \{1, \dots, I\}$ such that $\theta \in B_i$, and consider the following equality:

$$\begin{aligned} \frac{1}{m} \sum_{j=1}^m \varphi(\theta, X_j) - \mathbb{E}_\psi[\varphi(\theta, X)] &= \frac{1}{m} \sum_{j=1}^m \left\{ \varphi(\theta, X_j) - \varphi(\phi_{i_\theta}, X_j) \right\} \\ &\quad + \frac{1}{m} \sum_{j=1}^m \varphi(\phi_{i_\theta}, X_j) - \mathbb{E}_\psi[\varphi(\phi_{i_\theta}, X)] \\ &\quad + \mathbb{E}_\psi[\varphi(\phi_{i_\theta}, X)] - \mathbb{E}_\psi[\varphi(\theta, X)] \end{aligned}$$

The three last terms are functions of θ for which we want to bound the uniform norm.

The uniform norm of the third term is lower than ε since $\forall \theta \in K, d(\theta, \phi_{i_\theta}) \leq \eta_{(\phi_{i_\theta}, \varepsilon)}$. The second term converges to zero by the law of large number since $\{\phi_1, \dots, \phi_I\}$ is finite. Finally, the uniform norm of the second term can be bounded by

$$U_m = \max_{1 \leq i \leq I} \left\{ \frac{1}{m} \sum_{j=1}^m \sup_{\theta \in B_i} \|\varphi(\theta, X_j) - \varphi(\phi_i, X_j)\| \right\}.$$

The supremum are integrable by assumption, thus the law of large numbers applies:

$$U_m \xrightarrow[m \rightarrow +\infty]{a.s.} \max_{1 \leq i \leq I} f_{\phi_i}(\eta_{(\phi_i, \varepsilon)}) < \varepsilon.$$

To sum up, we have just proven that for any $\varepsilon > 0$, almost surely,

$$\limsup_{m \rightarrow +\infty} \left\{ \sup_{\theta \in K} \left\| \frac{1}{m} \sum_{j=1}^m \varphi(\theta, X_j) - \mathbb{E}_\psi [\varphi(\theta, X)] \right\| \right\} < 2\varepsilon.$$

Since ε is arbitrary small, we get the first claim of the lemma.

Now, if $\tilde{\theta}_m \rightarrow \tilde{\theta}$, we have eventually $\|\tilde{\theta}_m - \tilde{\theta}\| \leq \varepsilon$ with probability one. This yields the following inequality for m large enough:

$$\begin{aligned} & \left\| \frac{1}{m} \sum_{j=1}^m \varphi(\tilde{\theta}_m, X_j) - \mathbb{E}_\psi [\varphi(\tilde{\theta}, X)] \right\| \\ & \leq \sup_{\theta \in B(\tilde{\theta}_m, \varepsilon)} \left\| \frac{1}{m} \sum_{j=1}^m \varphi(\theta, X_j) - \mathbb{E}_\psi [\varphi(\theta, X)] \right\| \\ & \quad + \left\| \mathbb{E}_\psi [\varphi(\tilde{\theta}_m, X)] - \mathbb{E}_\psi [\varphi(\tilde{\theta}, X)] \right\|. \end{aligned}$$

The first term converges to zero since the first claim of the lemma applies to the compact closure of $B(\tilde{\theta}, \varepsilon)$. The continuity of the map $\theta \mapsto \mathbb{E}_\psi [\varphi(\theta, X)]$ ensures that the second term also goes to zero, proving the second claim of the Lemma.

B.1.5. Proof of Lemma 5

Let $\varepsilon > 0$, and $G_n(\theta, \omega) = \nabla_\theta \ell_n(\theta, \omega)$ defined on $B(\theta^*, \varepsilon)$. Define also $g_k^{(n)}(\theta)$ as the k -th component of $G_n(\theta)$. By assumption, for any $\delta > 0$,

$$\left\{ \omega \in \Omega : \max \left(\|\hat{\theta}_n - \theta^*\|, \|\sqrt{n}G_n(\hat{\theta}_n)\|, \sup_{\theta \in B(\theta^*, \varepsilon)} \|\nabla_\theta^2 \ell_n(\theta) - \mathcal{H}(\theta)\| \right) \leq \delta \right\}$$

defines a sequence of sets whose probability goes to one.

On any of these sets (for a fixed ω), Taylor Lagrange’s theorem ensures that we can find $(\tilde{\theta}_j^{(n)})_{j=1, \dots, d}$ on the segment line $[\theta^*, \hat{\theta}_n]$ such that

$$G_n(\hat{\theta}_n) = G_n(\theta^*) + \mathbf{H}_n(\hat{\theta}_n - \theta^*), \quad \mathbf{H}_n = \begin{pmatrix} (\nabla_\theta g_1^{(n)}(\tilde{\theta}_1^{(n)}))^T \\ \vdots \\ (\nabla_\theta g_d^{(n)}(\tilde{\theta}_d^{(n)}))^T \end{pmatrix}.$$

In particular, for any $\delta \in]0, \varepsilon[$,

$$\|\mathbf{H}_n - \mathcal{H}(\theta^*)\| \leq d \sup_{\theta \in B(\theta^*, \varepsilon)} \|\nabla_\theta^2 \ell_n(\theta) - \mathcal{H}(\theta)\| + \sum_{j=1}^d \|\mathcal{H}(\tilde{\theta}_j^{(n)}) - \mathcal{H}(\theta^*)\|.$$

For any $j = 1, \dots, d$, the distance between $\tilde{\theta}_j^{(n)}$ and θ^* is at most δ , and \mathcal{H} is continuous, thus δ can always be chosen small enough such that \mathbf{H}_n is invertible. We thus have:

$$\begin{aligned} \hat{\theta}_n - \theta^* &= \mathbf{H}_n^{-1} \{G_n(\hat{\theta}_n) - G_n(\theta^*)\} \\ \sqrt{n}(\hat{\theta}_n - \theta^*) + \mathcal{H}(\theta^*)^{-1} \sqrt{n}G_n(\theta^*) &= \mathbf{H}_n^{-1} \sqrt{n}G_n(\hat{\theta}_n) - \{\mathbf{H}_n^{-1} - \mathcal{H}(\theta^*)^{-1}\} \sqrt{n}G_n(\theta^*) \end{aligned}$$

The right hand side converges in probability to zero because $G_n(\hat{\theta}_n) = o_{\mathbb{P}}(n^{-1/2})$ by assumption, and because $\sqrt{n}G_n(\theta^*)$ converges in distribution and is thus bounded in probability. The last equality being true on a sequence of sets whose probability goes to one, this implies that the left hand side must also converge to zero in probability.

The last conclusion follows from Slutsky's lemma.

B.1.6. Proof of Lemma 6

Before proving the lemma, we recall a powerful result from Jones (2004). Under (X2), the chain $(X_j)_{j \geq 1}$ is asymptotically uncorrelated with exponential decay, i.e. there is some $\gamma > 0$ such that

$$\rho(n) = \sup \left\{ \text{Corr}(U, V), U \in L^2(\mathcal{F}_1^k), V \in L^2(\mathcal{F}_{k+n}^\infty), k \geq 1 \right\} = \mathcal{O}(e^{-\gamma n})$$

where \mathcal{F}_k^m is the sigma-algebra generated by X_k, \dots, X_m .

Let $h_n = f_n - f$, and note that $\mathbb{V}_\psi(h_n(X_0)) \leq \mathbb{E}_\psi[(h_n(X_0))^2] \xrightarrow{n \rightarrow \infty} 0$ by dominated convergence. Combined with the previous result, this implies that

$$\begin{aligned} \frac{1}{n} \mathbb{V} \left(\sum_{i=1}^n h_n(X_i) \right) &= \mathbb{V}_\psi(h_n(X_0)) \\ &\quad \times \left\{ 1 + 2 \sum_{i=1}^n \frac{n-i}{n} \text{Corr}(h_n(X_0), h_n(X_i)) \right\} \xrightarrow{n \rightarrow \infty} 0 \end{aligned}$$

since

$$\left| \sum_{i=1}^n \frac{n-i}{n} \text{Corr}(h_n(X_0), h_n(X_i)) \right| \leq \sum_{i=1}^{+\infty} \rho(i) < +\infty.$$

The first claim of the lemma follows from Chebyshev's inequality, since for any $\varepsilon > 0$

$$\mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n h_n(X_i) - \mathbb{E}[h_n(X)] \geq \frac{\varepsilon}{\sqrt{n}} \right) \leq \frac{1}{n\varepsilon^2} \mathbb{V} \left(\sum_{i=1}^n h_n(X_i) \right) \xrightarrow{n \rightarrow \infty} 0.$$

Finally, under (X2) a \sqrt{n} -CLT holds for f dominated by g

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n f(X_i) - \mathbb{E}[f(X)] \right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_f^2).$$

An application of Slutsky's lemma yields the second claim of the lemma.

B.2. Proofs of the remaining lemmas

B.2.1. Proof of Lemma 7

Assumption (H2) ensures in particular that the partition function $\theta \mapsto \mathcal{Z}(\theta)$ is differentiable in a neighborhood of $\hat{\theta}_n$. Write the Hessian of the Poisson Transform as the following block matrix:

$$\nabla_{(\theta, \nu)}^2 \ell_n(\theta, \nu) = \begin{pmatrix} A & b \\ b^T & c \end{pmatrix}$$

where

$$\begin{aligned} A &= \nabla_{\theta}^2 \ell_n(\theta, \nu) = \frac{1}{n} \sum_{i=1}^n \nabla_{\theta}^2 \log h_{\theta}(y_i) - e^{\nu} \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)} \frac{\nabla_{\theta}^2 \mathcal{Z}(\theta)}{\mathcal{Z}(\theta)}, \\ b &= \nabla_{\theta} \frac{\partial}{\partial \nu} \ell_n(\theta, \nu) = -e^{\nu} \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)} \frac{\nabla_{\theta} \mathcal{Z}(\theta)}{\mathcal{Z}(\theta)}, \\ c &= \frac{\partial^2}{\partial \nu^2} \ell_n(\theta, \nu) = -e^{\nu} \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)} < 0. \end{aligned}$$

The Hessian of the Poisson transform is negative definite if and only if Schur’s complement of c in the Hessian also is. Use the following equality to compute it:

$$\begin{aligned} \nabla_{\theta}^2 \log \mathcal{Z}(\theta) &= \frac{\nabla_{\theta}^2 \mathcal{Z}(\theta)}{\mathcal{Z}(\theta)} - \frac{\nabla_{\theta} \mathcal{Z}(\theta) (\nabla_{\theta} \mathcal{Z}(\theta))^T}{\mathcal{Z}(\theta)^2}, \\ A - bc^{-1}b^T &= \frac{1}{n} \sum_{i=1}^n \nabla_{\theta}^2 \log h_{\theta}(y_i) - e^{\nu} \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)} \nabla_{\theta}^2 \log \mathcal{Z}(\theta). \end{aligned}$$

At the point $\xi = \hat{\xi}_n$, Schur’s complement of c is also the Hessian of the log likelihood:

$$\nabla_{\theta}^2 \ell_n(\theta) = \frac{1}{n} \sum_{i=1}^n \nabla_{\theta}^2 \log h_{\theta}(y_i) - \nabla_{\theta}^2 \log \mathcal{Z}(\theta), \quad \left\{ e^{\nu} \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)} \right\}_{\xi=\hat{\xi}_n} = 1.$$

B.2.2. Proof of Lemma 9

Let $\xi_n \rightarrow \xi$, we have

$$\lim_{n \rightarrow +\infty} \sup_{k \geq n} \mathcal{M}_{\tau}^{\text{NCE}}(\xi_k) \leq \frac{1}{\mathcal{Z}(\psi)} \lim_{n \rightarrow +\infty} \left\{ \int_{\mathcal{X}} \sup_{k \geq n} \varphi_k(x) \mu(dx) \right\},$$

where

$$\begin{aligned} \varphi_k(x) &= \log \left\{ \frac{e^{\nu_k} h_{\theta_k}(x)}{e^{\nu^*} h_{\theta^*}(x)} \right\} e^{\nu^*} h_{\theta^*}(x) \\ &\quad + \log \left\{ \frac{\tau h_{\psi}(x) + e^{\nu^*} h_{\theta^*}(x)}{\tau h_{\psi}(x) + e^{\nu_k} h_{\theta_k}(x)} \right\} (\tau h_{\psi}(x) + e^{\nu^*} h_{\theta^*}(x)). \end{aligned}$$

The sequence $\{\sup_{k \geq n} \varphi_k\}$ is a decreasing sequence converging pointwise. It may be bounded from above thanks to the log-sum inequality, since for any k we have

$$\varphi_k \leq \log \left\{ \frac{e^{\nu_k} h_{\theta_k}}{e^{\nu^*} h_{\theta^*}} \right\} e^{\nu^*} h_{\theta^*} + \log \left\{ \frac{\tau h_\psi}{\tau h_\psi} \right\} \tau h_\psi + \log \left\{ \frac{e^{\nu^*} h_{\theta^*}}{e^{\nu_k} h_{\theta_k}} \right\} e^{\nu^*} h_{\theta^*} = 0.$$

Monotone convergence theorem applies:

$$\lim_{n \rightarrow +\infty} \sup_{k \geq n} \mathcal{M}_\tau^{\text{NCE}}(\xi_k) \leq \frac{1}{\mathcal{Z}(\psi)} \int_{\mathcal{X}} \lim_{n \rightarrow +\infty} \varphi_n(x) \mu(dx) = \mathcal{M}_\tau^{\text{NCE}}(\xi).$$

B.2.3. Proof of Lemma 10

Without loss of generality, we may suppose that $\mathbb{E}_\theta[Z] = 1$. Recall the following expressions:

$$\nabla g_\xi = \begin{pmatrix} \nabla g_\theta \\ 1 \end{pmatrix} \quad \nabla \nabla^T g_\xi = \begin{pmatrix} \nabla \nabla^T g_\theta & \nabla g_\theta \\ \nabla^T g_\theta & 1 \end{pmatrix}.$$

We thus have

$$\begin{aligned} \mathbb{E}_\theta[\nabla g_\xi Z] \mathbb{E}_\theta[\nabla^T g_\xi Z] &= \begin{pmatrix} \mathbb{E}_\theta[\nabla g_\theta Z] \mathbb{E}_\theta[\nabla^T g_\theta Z] & \mathbb{E}_\theta[\nabla g_\theta Z] \\ \mathbb{E}_\theta[\nabla^T g_\theta Z] & 1 \end{pmatrix}, \\ \mathbb{E}_\theta[\nabla \nabla^T g_\xi Z] &= \begin{pmatrix} \mathbb{E}_\theta[\nabla \nabla^T g_\theta Z] & \mathbb{E}_\theta[\nabla g_\theta Z] \\ \mathbb{E}_\theta[\nabla^T g_\theta Z] & 1 \end{pmatrix}. \end{aligned}$$

We use the following decomposition

$$\mathbb{E}_\theta[\nabla g_\xi Z] \mathbb{E}_\theta[\nabla^T g_\xi Z] = \mathbb{E}_\theta[\nabla \nabla^T g_\xi Z] - \begin{pmatrix} A_Z & 0 \\ 0 & 0 \end{pmatrix}$$

where Schur's complement $A_Z = \mathbb{E}_\theta[\nabla \nabla^T g_\theta Z] - \mathbb{E}_\theta[\nabla g_\theta Z] \mathbb{E}_\theta[\nabla^T g_\theta Z]$ is definite positive.

So we can re-write the matrix \mathbf{M} as:

$$\mathbf{M} = \mathbb{E}_\theta[\nabla \nabla^T g_\xi Z]^{-1} - \mathbb{E}_\theta[\nabla \nabla^T g_\xi Z]^{-1} \begin{pmatrix} A_Z & 0 \\ 0 & 0 \end{pmatrix} \mathbb{E}_\theta[\nabla \nabla^T g_\xi Z]^{-1}.$$

Now, on the one hand, an inverse block matrix calculation yields

$$\mathbb{E}_\theta[\nabla \nabla^T g_\xi Z]^{-1} = \begin{pmatrix} A_Z^{-1} & -A_Z^{-1} \mathbb{E}_\theta[\nabla g_\theta Z] \\ -\mathbb{E}_\theta[\nabla^T g_\theta Z] A_Z^{-1} & 1 + \mathbb{E}_\theta[\nabla^T g_\theta Z] A_Z^{-1} \mathbb{E}_\theta[\nabla g_\theta Z] \end{pmatrix},$$

while, on the other hand, a direct computation yields

$$\begin{aligned} \mathbb{E}_\theta[\nabla \nabla^T g_\xi Z]^{-1} \begin{pmatrix} A_Z & 0 \\ 0 & 0 \end{pmatrix} \mathbb{E}_\theta[\nabla \nabla^T g_\xi Z]^{-1} \\ = \begin{pmatrix} A_Z^{-1} & -A_Z^{-1} \mathbb{E}_\theta[\nabla g_\theta Z] \\ -\mathbb{E}_\theta[\nabla^T g_\theta Z] A_Z^{-1} & \mathbb{E}_\theta[\nabla^T g_\theta Z] A_Z^{-1} \mathbb{E}_\theta[\nabla g_\theta Z] \end{pmatrix}. \end{aligned}$$

The matrix \mathbf{M} being the difference between these two quantities, we get the claim of the lemma.

B.3. Proofs of MC-MLE consistency and asymptotic normality

B.3.1. MC-MLE consistency

The following proof is a straightforward adaptation of Wald’s proof of consistency for the MLE (Wald (1949)). The sketch of proof is mainly inspired from Geyer (2012), which has the merit of giving a very accessible presentation of this technical proof.

To begin, define the opposite of the Kullback-Leibler divergence:

$$\lambda(\theta) = \mathbb{E}_{\theta^*} \left[\log \frac{f_\theta(Y)}{f_{\theta^*}(Y)} \right] \leq 0.$$

Since the model is identifiable, λ has a unique maximizer achieved at θ^* . It may be $-\infty$ for some values of θ , but this does not pose problems in the following proof.

For convenience, we choose to analyse the MC-MLE objective function through the following translational motion (sharing the same maximiser with $\ell_{n,m}^{\text{IS}}$):

$$M_n^{\text{IS}}(\theta, \nu) = \frac{1}{n} \sum_{i=1}^n \log \left\{ \frac{e^\nu h_\theta(Y_i)}{e^{\nu^*} h_{\theta^*}(Y_i)} \right\} + 1 - e^\nu \frac{1}{r_n} \sum_{j=1}^{r_n} \frac{h_\theta(X_j)}{h_\psi(X_j)}.$$

For any $\xi \in \Xi = \Theta \times \mathbb{R}$, the law of large numbers yields $M_n^{\text{IS}}(\xi) \xrightarrow[n \rightarrow +\infty]{a.s.} \mathcal{M}^{\text{IS}}(\xi)$ where

$$\mathcal{M}^{\text{IS}}(\theta, \nu) = \lambda(\theta) + \nu + \log \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)} + 1 - e^\nu \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)} \leq 0.$$

Note that by construction \mathcal{M}^{IS} also has a unique maximiser at $\xi^* = (\theta^*, \nu^*)$.

Let $\eta > 0$. Define $K_\eta = \{\xi \in K : d(\xi, \xi^*) \geq \eta\}$ where K is the compact set defined in (C2). Under (H3), continuity of the maps $\theta \mapsto h_\theta(x)$ and monotone convergence ensure that for any $\xi \in K_\eta$,

$$\lim_{\varepsilon \downarrow 0} \mathbb{E}_{\theta^*} \left[\sup_{(\phi, \mu) \in B(\xi, \varepsilon)} \log \frac{e^\mu h_\phi(Y)}{e^{\nu^*} h_{\theta^*}(Y)} \right] = \mathbb{E}_{\theta^*} \left[\log \frac{e^\nu h_\theta(Y)}{e^{\nu^*} h_{\theta^*}(Y)} \right],$$

while dominated convergence ensures that

$$\lim_{\varepsilon \downarrow 0} \mathbb{E}_\psi \left[\inf_{(\phi, \mu) \in B(\xi, \varepsilon)} e^\mu \frac{h_\phi(X)}{h_\psi(X)} \right] = e^\nu \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)}.$$

Thus for any $\xi \in K_\eta$ and $\gamma > 0$, we can always find $\varepsilon_\xi > 0$ such that simultaneously:

$$\mathbb{E}_{\theta^*} \left[\sup_{(\phi, \mu) \in B(\xi, \varepsilon_\xi)} \log \frac{e^\mu h_\phi(Y)}{e^{\nu^*} h_{\theta^*}(Y)} \right] \leq \mathbb{E}_{\theta^*} \left[\log \frac{e^\nu h_\theta(Y)}{e^{\nu^*} h_{\theta^*}(Y)} \right] + \frac{\gamma}{2},$$

and

$$\mathbb{E}_\psi \left[\inf_{(\phi, \mu) \in B(\xi, \varepsilon_\xi)} e^\mu \frac{h_\phi(X)}{h_\psi(X)} \right] \geq e^\nu \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)} - \frac{\gamma}{2}.$$

The set of open balls $\{B(\xi, \varepsilon_\xi) : \xi \in K\}$ form an open cover of K_η from which we can extract a finite subcover by compactness, i.e. we can build a finite set $\{\xi_1, \dots, \xi_p\} \subset K_\eta$ such that $K_\eta \subset \bigcup_{k=1}^p B(\xi_k, \varepsilon_{\xi_k})$. This yields the following inequality:

$$\sup_{\xi \in K_\eta} M_n^{\text{IS}}(\xi) \leq \max_{k=1, \dots, p} \left\{ \frac{1}{n} \sum_{i=1}^n \sup_{(\phi, \mu) \in B(\xi_k, \varepsilon_{\xi_k})} \left(\log \frac{e^\mu h_\phi(Y_i)}{e^{\nu^*} h_{\theta^*}(Y_i)} \right) + 1 - \frac{1}{r_n} \sum_{j=1}^{r_n} \inf_{(\phi, \mu) \in B(\xi_k, \varepsilon_{\xi_k})} \left(e^\mu \frac{h_\phi(X_j)}{h_\psi(X_j)} \right) \right\}.$$

The right hand side converges almost surely as the law of large numbers applies simultaneously on a finite set. We can thus bound the upper limit:

$$\limsup_{n \rightarrow +\infty} \sup_{\xi \in K_\eta} M_n^{\text{IS}}(\xi) \leq \max_{k=1, \dots, p} \left\{ \mathbb{E}_{\theta^*} \left[\sup_{(\phi, \mu) \in B(\xi_k, \varepsilon_{\xi_k})} \left(\log \frac{e^\mu h_\phi(Y)}{e^{\nu^*} h_{\theta^*}(Y)} \right) \right] + 1 - \mathbb{E}_\psi \left[\inf_{(\phi, \mu) \in B(\xi_k, \varepsilon_{\xi_k})} \left(e^\mu \frac{h_\phi(X)}{h_\psi(X)} \right) \right] \right\},$$

$$\limsup_{n \rightarrow +\infty} \sup_{\xi \in K_\eta} M_n^{\text{IS}}(\xi) \leq \max_{k=1, \dots, p} \mathcal{M}^{\text{IS}}(\xi_k) + \gamma \leq \sup_{\xi \in K_\eta} \mathcal{M}^{\text{IS}}(\xi) + \gamma.$$

Moreover γ is arbitrary small, thus the inequality still holds when γ is zero:

$$\limsup_{n \rightarrow +\infty} \sup_{\xi \in K_\eta} M_n^{\text{IS}}(\xi) \leq \sup_{\xi \in K_\eta} \mathcal{M}^{\text{IS}}(\xi) \quad \text{a.s.} \quad (23)$$

To conclude, let us prove that the right hand side is negative. Indeed, subadditivity of the supremum yields

$$\sup_{\xi \in K_\eta} \mathcal{M}^{\text{IS}}(\theta, \nu) \leq \sup_{\xi \in K_\eta} \lambda(\theta) + \sup_{\xi \in K_\eta} \left(\nu + \log \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)} + 1 - e^\nu \frac{\mathcal{Z}(\theta)}{\mathcal{Z}(\psi)} \right)$$

where the second term is non positive by construction. Under (H3), it is easy to check that λ is upper semi continuous, which implies in particular that λ achieves its maximum on any compact set. Consequently: $\sup_{\xi \in K_\eta} \mathcal{M}^{\text{IS}}(\xi) \leq \sup_{\xi \in K_\eta} \lambda(\theta) < 0$.

The last part of the proof is the same as for NCE consistency (see the appendix).

B.3.2. MC-MLE asymptotic normality

For convenience, for any $\xi = (\theta, \nu)$, we note $g_\xi(x) = \nu + \log h_\theta(x)$.

Let $G_n^{\text{IS}}(\xi) = \nabla_{\xi} \ell_{n,m}^{\text{IS}}(\xi)$ and $\mathbf{H}_n^{\text{IS}}(\xi) = \nabla_{\xi}^2 \ell_{n,m}^{\text{IS}}(\xi)$. We have

$$G_n^{\text{IS}}(\xi) = \frac{1}{n} \sum_{i=1}^n \nabla_{\xi} g_{\xi}(Y_i) - \frac{1}{m_n} \sum_{j=1}^{m_n} \nabla_{\xi} g_{\xi}(X_j) \frac{\exp\{g_{\xi}(X_j)\}}{h_{\psi}(X_j)},$$

$$\mathbf{H}_n^{\text{IS}}(\xi) = \frac{1}{n} \sum_{i=1}^n \nabla_{\xi}^2 g_{\xi}(Y_i) - \frac{1}{m_n} \sum_{j=1}^{m_n} \left\{ (\nabla_{\xi}^2 + \nabla_{\xi} \nabla_{\xi}^T) g_{\xi}(X_j) \right\} \frac{\exp\{g_{\xi}(X_j)\}}{h_{\psi}(X_j)}. \quad (24)$$

We start by proving that, almost surely,

$$\sup_{\xi \in B(\xi^*, \varepsilon)} \|\mathbf{H}_n^{\text{IS}}(\xi) - \mathbf{H}(\xi)\| \xrightarrow{n \rightarrow \infty} 0, \quad (25)$$

where

$$\mathbf{H}(\xi) = \mathbb{E}_{\theta^*} \left[\nabla_{\xi}^2 g_{\xi}(Y) \right] - \mathbb{E}_{\psi} \left[\left\{ (\nabla_{\xi}^2 + \nabla_{\xi} \nabla_{\xi}^T) g_{\xi}(X) \right\} \frac{\exp\{g_{\xi}(X)\}}{h_{\psi}(X)} \right].$$

To prove (25), split the supremum norm in two and apply Lemma 4 to both empirical averages in definition (24). Both supremum norms are integrable under (H4), this is proven in the following.

$$\nabla_{\xi}^2 g_{\xi}(x) = \begin{pmatrix} \nabla_{\theta}^2 \log h_{\theta}(x) & 0 \\ 0 & 0 \end{pmatrix}$$

$$\nabla_{\xi} \nabla_{\xi}^T g_{\xi}(x) = \begin{pmatrix} \nabla_{\theta} \nabla_{\theta}^T \log h_{\theta}(x) & \nabla_{\theta} \log h_{\theta}(x) \\ \nabla_{\theta}^T \log h_{\theta}(x) & 1 \end{pmatrix}$$

First supremum norm is integrable under (H4), since

$$\int_{\mathcal{X}} \sup_{\xi \in B(\xi^*, \varepsilon)} \|\nabla_{\xi}^2 g_{\xi}(x)\| h_{\theta^*} \mu(dx)$$

$$\leq \int_{\mathcal{X}} \sup_{\theta \in B(\theta^*, \varepsilon)} \|\nabla_{\theta}^2 \log h_{\theta}(Y)\| \sup_{\theta \in B(\theta^*, \varepsilon)} h_{\theta}(x) \mu(dx) < +\infty.$$

For the second one, use the following decomposition:

$$\|(\nabla_{\xi}^2 + \nabla_{\xi} \nabla_{\xi}^T) g_{\xi}(x)\|_1 = \|(\nabla_{\theta}^2 + \nabla_{\theta} \nabla_{\theta}^T) \log h_{\theta}(x)\|_1 + 2\|\nabla_{\theta} \log h_{\theta}(x)\|_1 + 1,$$

$$\|(\nabla_{\theta}^2 + \nabla_{\theta} \nabla_{\theta}^T) \log h_{\theta}(x)\|_1 \leq \|\nabla_{\theta}^2 \log h_{\theta}(x)\|_1 + \|\nabla_{\theta} \log h_{\theta}(x)\|_1^2,$$

$$\|\nabla_{\theta} \log h_{\theta}(x)\|_1 \leq 1 + \|\nabla_{\theta} \log h_{\theta}(x)\|_1^2.$$

This yields a finite upper bound under (H4), since

$$\int_{\mathcal{X}} \sup_{\xi \in B(\xi^*, \varepsilon)} \|(\nabla_{\xi}^2 + \nabla_{\xi} \nabla_{\xi}^T) g_{\xi}(x)\|_1 \exp\{g_{\xi}(x)\} \mu(dx)$$

$$\leq e^{\nu^* + \varepsilon} \int_{\mathcal{X}} \sup_{\theta \in B(\theta^*, \varepsilon)} \left(\|\nabla_{\theta}^2 \log h_{\theta}(x)\|_1 + 3\|\nabla_{\theta} \log h_{\theta}(x)\|_1^2 + 3 \right)$$

$$\times \sup_{\theta \in B(\theta^*, \varepsilon)} h_{\theta}(x) \mu(dx) < +\infty.$$

Note also that, at the point $\xi = \xi^*$, functions \mathbf{H} and $-\mathbf{J}$ coincide, where

$$\mathbf{J}(\xi) = \mathbb{E}_\theta [\nabla_\xi \nabla_\xi^T g_\xi(Y)] = \begin{pmatrix} \mathbb{E}_\theta [\nabla_\theta \nabla_\theta^T \log h_\theta(Y)] & \mathbb{E}_\theta [\nabla_\theta \log h_\theta(Y)] \\ \mathbb{E}_\theta [\nabla_\theta^T \log h_\theta(Y)] & 1 \end{pmatrix}.$$

In particular, the matrix $\mathbf{J}(\xi^*)$ is definite positive, since Schur's complement is also the Fisher Information, definite positive by assumption:

$$\mathbb{E}_\theta [\nabla_\theta \nabla_\theta^T g_\theta(Y)] - \mathbb{E}_\theta [\nabla_\theta g_\theta(Y)] \mathbb{E}_\theta [\nabla_\theta^T g_\theta(Y)] = \mathbb{V}_\theta (\nabla_\theta \log h_\theta(Y)) = \mathbf{I}(\theta).$$

Now we establish the weak convergence of the gradient. We have

$$\begin{aligned} \sqrt{n} G_n^{\text{IS}}(\xi^*) &= \sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n \nabla_\xi g_\xi(Y_i) - \mathbb{E}_\theta [\nabla_\xi g_\xi(Y)] \right)_{|\xi=\xi^*} \\ &\quad - \sqrt{\frac{n}{m_n}} \sqrt{m_n} \left(\frac{1}{m_n} \sum_{j=1}^{m_n} \nabla_\xi g_\xi(X_j) \frac{f_\theta(X_j)}{f_\psi(X_j)} - \mathbb{E}_\theta [\nabla_\xi g_\xi(Y)] \right)_{|\xi=\xi^*}. \end{aligned}$$

Simulations and observations are assumed independent, thus Slutsky's lemma yields the following. Second moment conditions hold under (I3).

$$\sqrt{n} G_n^{\text{IS}}(\xi^*) \xrightarrow{\mathcal{D}} \mathcal{N}_{d+1} \left(0, \boldsymbol{\Sigma}(\xi^*) + \tau^{-1} \boldsymbol{\Gamma}(\xi^*) \right),$$

where

$$\boldsymbol{\Sigma}(\xi) = \mathbb{V}_\theta (\nabla_\xi g_\xi(Y)) = \begin{pmatrix} \mathbf{I}(\theta) & 0 \\ 0 & 0 \end{pmatrix},$$

and

$$\boldsymbol{\Gamma}(\xi) = \mathbb{V}_\psi (\varphi_\xi^{\text{IS}}(X)) + 2 \sum_{i=1}^{+\infty} \text{Cov}(\varphi_\xi^{\text{IS}}(X_0), \varphi_\xi^{\text{IS}}(X_i)), \quad \varphi_\xi^{\text{IS}} = (\nabla_\xi g_\xi) \frac{f_\theta}{f_\psi}.$$

Finally, Lemma 5 applies:

$$\sqrt{n} (\hat{\xi}_{n,m}^{\text{IS}} - \xi^*) \xrightarrow{\mathcal{D}} \mathcal{N} \left(0_{\mathbb{R}^{d+1}}, \mathbf{V}_\tau^{\text{IS}}(\xi^*) \right)$$

where $\mathbf{V}_\tau^{\text{IS}}(\xi) = \mathbf{J}(\xi)^{-1} \{ \boldsymbol{\Sigma}(\xi) + \tau^{-1} \boldsymbol{\Gamma}(\xi) \} \mathbf{J}(\xi)^{-1}$.

B.4. Proofs related to exponential families

The following calculations are entirely classical. For the sake of completeness, we present the few tricks required for proving Propositions 1, 2 and 3.

To begin, define $b(x) = \text{sgn}(S(x))$, the vector composed by the signs of each component of $S(x)$. Note that for any $\theta \in \Theta$, the following supremum is necessarily achieved on the boundary of the 1-ball, in the direction of the sign vector:

$$\sup_{\|\phi - \theta\|_1 \leq \varepsilon} \exp \{ \phi^T S(x) \} = \exp \{ (\theta + \varepsilon b(x))^T S(x) \}. \quad (26)$$

Since $\|S(x)\|_1 = b(x)^T S(x)$, we have (for the 1-norm for instance):

$$\sup_{\phi \in B(\theta, \varepsilon)} \left(\log \frac{h_\phi(x)}{h_{\theta^*}(x)} \right) = (\theta - \theta^*)^T S(x) + \varepsilon \|S(x)\| \leq (\|\theta - \theta^*\| + \varepsilon) \|S(x)\|,$$

which proves the claim of Proposition 2, since

$$\begin{aligned} & \int_{\mathcal{X}} \sup_{\phi \in B(\theta, \varepsilon)} \left(\log \frac{h_\phi(x)}{h_{\theta^*}(x)} \right)_+ h_{\theta^*}(x) \mu(dx) \\ & \leq (\|\theta - \theta^*\| + \varepsilon) \int_{\mathcal{X}} \|S(x)\| h_{\theta^*}(x) \mu(dx) < +\infty. \end{aligned}$$

For Propositions 1 and 3, use also the fact that $\|S(x)\|_1 = b(x)^T S(x)$ and that $y \leq e^y$ for any $y \in \mathbb{R}$. We have

$$\|S(x)\|_1^2 \leq \varepsilon^{-2} \exp \{ 2\varepsilon b(x)^T S(x) \}. \tag{27}$$

Equations (26) and (27) can be combined as follows:

$$\begin{aligned} \int_{\mathcal{X}} (1 + \|S(x)\|^2) \sup_{\phi \in B(\theta, \varepsilon)} h_\phi(x) \mu(dx) & \leq \sum_{b \in \{-1, 1\}^d} \int_{\mathcal{X}} \exp \{ (\theta + b\varepsilon)^T S(x) \} \mu(dx) \\ & \quad + \varepsilon^{-2} \sum_{b \in \{-1, 1\}^d} \int_{\mathcal{X}} \exp \{ (\theta + 3b\varepsilon)^T S(x) \} \mu(dx), \end{aligned}$$

and

$$\begin{aligned} & \mathbb{E}_\psi \left[(1 + \|S(X)\|^2) \sup_{\phi \in B(\theta, \varepsilon)} \left(\frac{h_\phi(X)}{h_\psi(X)} \right)^2 \right] \\ & \leq \sum_{b \in \{-1, 1\}^d} \mathbb{E}_\psi \left[\left(\frac{\exp \{ (\theta + b\varepsilon)^T S(X) \}}{h_\psi(X)} \right)^2 \right] \\ & \quad + \varepsilon^{-2} \sum_{b \in \{-1, 1\}^d} \mathbb{E}_\psi \left[\left(\frac{\exp \{ (\theta + 2b\varepsilon)^T S(X) \}}{h_\psi(X)} \right)^2 \right]. \end{aligned}$$

Choosing $\theta = \hat{\theta}_n$ in the preceding inequalities yields Proposition 1, while choosing $\theta = \theta^*$ yields Proposition 3.

References

Ando, T. (1979). Concavity of certain maps on positive definite matrices and applications to hadamard products. *Linear Algebra and its Applications*, 26:203–241. [MR0535686](#)

Barthelme, S. and Chopin, N. (2015). The Poisson transform for unnormalised statistical models. *Stat. Comput.*, 25(4):767–780. [MR3360491](#)

- Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press. [MR2061575](#)
- Bradley, R. C. et al. (2005). Basic properties of strong mixing conditions. a survey and some open questions. *Probability surveys*, 2:107–144. [MR2178042](#)
- Caimo, A. and Friel, N. (2011). Bayesian inference for exponential random graph models. *Social Networks*, 33(1):41–55. [MR2873466](#)
- Delyon, B. (2018). Estimation paramétrique. *Unpublished lecture notes*.
- Geyer, C. J. (1994). On the convergence of Monte Carlo maximum likelihood calculations. *J. Roy. Statist. Soc. Ser. B*, 56(1):261–274. [MR1257812](#)
- Geyer, C. J. (2012). The Wald consistency theorem. *Unpublished lecture notes*.
- Gu, M. G. and Zhu, H.-T. (2001). Maximum likelihood estimation for spatial models by Markov chain Monte Carlo stochastic approximation. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 63(2):339–355. [MR1841419](#)
- Gutmann, M. U. and Hyvärinen, A. (2012). Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *J. Mach. Learn. Res.*, 13:307–361. [MR2913702](#)
- Jones, G. L. (2004). On the Markov chain central limit theorem. *Probab. Surv.*, 1:299–320. [MR2068475](#)
- Meyn, S. P. and Tweedie, R. L. (2012). *Markov chains and stochastic stability*. Springer Science & Business Media. [MR1287609](#)
- Mnih, A. and Teh, Y. W. (2012). A fast and simple algorithm for training neural probabilistic language models. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 1751–1758.
- Roberts, G. O. and Rosenthal, J. S. (2004). General state space Markov chains and MCMC algorithms. *Probab. Surv.*, 1:20–71. [MR2095565](#)
- Rockafellar, R. T. and Wets, R. J.-B. (2009). *Variational analysis*, volume 317. Springer Science & Business Media.
- Salakhutdinov, R. and Hinton, G. (2009). Deep boltzmann machines. In *Artificial Intelligence and Statistics*, pages 448–455.
- Wald, A. (1949). Note on the consistency of the maximum likelihood estimate. *Ann. Math. Statistics*, 20:595–601. [MR0032169](#)
- Walker, S. G. (2011). Posterior sampling when the normalizing constant is unknown. *Comm. Statist. Simulation Comput.*, 40(5):784–792. [MR2783887](#)
- Wang, C., Komodakis, N., and Paragios, N. (2013). Markov random field modeling, inference & learning in computer vision & image understanding: A survey. *Computer Vision and Image Understanding*, 117(11):1610–1627.